No. 2024-2203

United States Court of Appeals for the Federal Circuit

NETLIST, INC.,

Plaintiff-Appellee,

v.

SAMSUNG ELECTRONICS CO., LTD., SAMSUNG ELECTRONICS AMERICA, INC., SAMSUNG SEMICONDUCTOR, INC.,

Defendants-Appellants.

APPEAL FROM UNITED STATES DISTRICT COURT FOR THE EASTERN DISTRICT OF TEXAS IN CASE NO. 2:21-cv-00463-JRG

DEFENDANTS-APPELLANTS' NON-CONFIDENTIAL PRINCIPAL BRIEF

Ruffin B. Cordell
Lauren A. Degnan
Michael J. McKeon
Michael J. Ballanco
Benjamin J. Christoff
FISH & RICHARDSON P.C.
1000 Maine Avenue S.W., Suite 1000
Washington, DC 20024
Tel: (202) 783-5070
cordell@fr.com

Francis J. Albert
FISH & RICHARDSON P.C.
12860 El Camino Real, Suite 400
San Diego, CA 92130
Tel: (858) 678-5070

Alexander Pechette
FISH & RICHARDSON P.C.
One Marina Park Drive
Suite 1700
Boston, MA 02210
Tel: (617) 542-5070

December 16, 2024

Attorneys for Defendants-Appellants

Claim 1 of U.S. Patent No. 8,787,060 (Appx696-719)

1. A memory package, comprising:

a plurality of input/output terminals via which the memory package communicates data and control/address signals with one or more external devices;

a plurality of stacked array dies including a first group of array dies and a second group of at least one array die, each array die having data ports;

at least a first die interconnect and a second die interconnect, the first die interconnect in electrical communication with the first group of array dies and not in electrical communication with the second group of at least one array die, the second die interconnect in electrical communication with the second group of at least one array die and not in electrical communication with the first group of array dies; and

a control die comprising at least a first data conduit between the first die interconnect and a first terminal of the plurality of input/output terminals, and at least a second data conduit between the second die interconnect and the first terminal, the first terminal being a data terminal, the control die further comprising a control circuit to control respective states of the first data conduit and the second data conduit in response to control signals received via one or more second terminals of the plurality of terminals.

Claim 1 of U.S. Patent No. 9,318,160 (Appx720-743)

1. A memory package, comprising:

data terminals and control terminals;

stacked array dies including a first group of array dies and a second group of at least one array die;

first die interconnects and second die interconnects, the first die interconnects in electrical communication with the first group of array dies and not in electrical communication with the second group of at least one array die, the second die interconnects in electrical communication with the second group of at least one array die and not in electrical communication with the first group of array dies; and

a control die comprising first data conduits between the first die interconnects and the data terminals, and second data conduits between the second die interconnects and the data terminals, the first data conduit including first drivers each having a first driver size and configured to drive a data signal from a corresponding data terminal to the first group of array dies, the second data conduit including second drivers each having a second driver size and configured to drive a data signal from a corresponding data terminal to the second group of at least one array die, the second driver size being different from the first driver size.

Claim 16 of U.S. Patent No. 11,016,918 (Appx600-647)

16. A memory module comprising:

a printed circuit board (PCB) having an interface configured to fit into a corresponding slot connector of a host system, the interface including a plurality of edge connections configured to couple power, data, address and control signals between the memory module and the host system;

first, second, and third buck converters configured to receive a pre-regulated input voltage and to produce first, second and third regulated voltages, respectively;

a converter circuit configured to reduce the pre-regulated input voltage to provide a fourth regulated voltage, wherein the first, second, third and fourth regulated voltages have first, second, third, and fourth voltage amplitudes, respectively;

a plurality of components coupled to the PCB, the plurality of components including a plurality of synchronous dynamic random access memory (SDRAM) devices, each component of the plurality of components coupled to one or more regulated voltages of the first, second, third and fourth regulated voltages; and

a voltage monitor circuit configured to monitor an input voltage received via a first portion of the plurality of edge connections, the voltage monitor circuit configured to produce a signal in response to the input voltage having a voltage amplitude that is greater than a first threshold voltage.

Claim 16 of U.S. Patent No. 11,232,054 (Appx648-695)

16. A memory module comprising:

a printed circuit board (PCB) having an interface configured to fit into a corresponding slot connector of a host system;

a voltage conversion circuit coupled to the PCB and configured to provide a plurality of regulated voltages, wherein the voltage conversion circuit includes three buck converters each of which is configured to produce a regulated voltage of the plurality of regulated voltages;

a plurality of components coupled to the PCB, each component of the plurality of components coupled to at least one regulated voltage of the plurality of regulated voltages, the plurality of components including a plurality of synchronous dynamic random access memory (SDRAM) devices, the plurality of SDRAM devices coupled to a first regulated voltage of the plurality of regulated voltages; and

a controller coupled to the PCB, the controller including a voltage monitor circuit coupled to an input voltage received from the host system via the interface, the voltage monitor circuit configured to detect an amplitude change in the input voltage, wherein, in response to the voltage monitor detecting an amplitude change in the input voltage, the memory module transitions from a first operable state to a second operable state.

Claim 1 of U.S. Patent No. 10,949,339 (Appx564-599)

1. A N-bit-wide memory module mountable in a memory socket of a computer system and configurable to communicate with a memory controller of the computer system via address and control signal lines and N-bit wide data signal lines, the N-bit wide data signal lines including a plurality of sets of data signal lines, each set of data signal lines is a byte wide, the memory module comprising:

a printed circuit board (PCB) having an edge connector comprising a plurality of electrical contacts which are positioned on an edge of the PCB and configured to be releasably coupled to corresponding contacts of the memory socket;

double data rate dynamic random access memory (DDR DRAM) devices coupled to the PCB and arranged in multiple N-bit-wide ranks;

a module controller coupled to the PCB and operatively coupled to the DDR DRAM devices, wherein the module controller is configurable to receive from the memory controller via the address and control signal lines input address and control signals for a memory write operation to write N-bit-wide write data from the memory controller into a first N-bit-wide rank of the multiple N-bit-wide ranks, and to output registered address and control signals in response to receiving the input address and control signals, wherein the registered address and control signals cause the first N-bit-wide rank to perform the memory write operation by receiving the N-bit-wide write data, wherein the module controller is further configurable to output module control signals in response to at least some of the input address and control signals; and

a plurality of byte-wise buffers coupled to the PCB and configured to receive the module control signals, wherein each respective byte-wise buffer of the plurality of byte-wise buffers has a first side configured to be operatively coupled to a respective set of data signal lines, a second side that is operatively coupled to at least one respective DDR DRAM device in each of the multiple N-bit-wide ranks via respective module data lines, and a byte-wise data path between the first side and the second side, wherein the each respective byte-wise buffer is disposed on the PCB at a respective position corresponding to the respective set of the plurality of sets of data signal lines;

wherein the each respective byte-wise buffer further includes logic configurable to control the byte-wise data path in response to the module control signals, wherein the byte-wise data path is enabled for a first time period in accordance with a latency parameter to actively drive a respective byte-wise section of the N-bit wide write data associated with the memory operation from the first side to the second side during the first time period; and

wherein the byte-wise data path includes first tristate buffers, and the logic in response to the module control signals is configured to enable the first tristate buffers to drive the respective byte-wise section of the N-bit wide write data to the respective module data lines during the first time period.

CERTIFICATE OF INTEREST

Counsel for Defendants-Appellants Samsung Electronics Co., Ltd., Samsung Electronics America, Inc., and Samsung Semiconductor, Inc. (collectively, "Samsung" or "Defendants-Appellants") certify the following:

1. Provide the full names of all entities represented by undersigned counsel in this case.

   Samsung Electronics Co., Ltd.; Samsung Electronics America, Inc.; and Samsung Semiconductor, Inc.

2. Provide the full names of all real parties in interest for the entities. Do not list the real parties if they are the same as the entities.

   None/Not Applicable

3. Provide the full names of all parent corporations for the entities and all publicly held companies that own 10% or more stock in the entities.

   Samsung Electronics Co., Ltd.; Samsung Electronics America, Inc.

4. List all law firms, partners, and associates that (a)appeared for the entities in the originating court or agency or (b) are expected to appear in this court for the entities. Do not include those who have already entered an appearance in this court. Fed. Cir. R. 47.4(a)(4).

   Fish & Richardson P.C.: Brian Livedalen, Sara Fish, Anthony V. Nguyen, Thomas H. Reger, II, Katherine Reardon, Matt Colvin, Jeffrey H. Burton, Karolina Jesien, Matthew Mosteller, Daniel A. Tishman

   Gillam & Smith LLP: Melissa Richards Smith, Andrew T. Gorham, James Travis Underwood

   Covington & Burling LLP: Brian R. Nester, Peter A. Swanson, R. Jason Fowler, Alice J. Ahn

5. Other than the originating case(s) for this case, are there related or prior cases that meet the criteria under Fed. Cir. R. 47.5(a)?

   Yes (*see* Samsung's Notice of Related Case Information filed at Dkt. 6)

Samsung Electronics Co., Ltd.
Ex. 1077, p. 6

6. Provide any information required under Fed. R. App. P. 26.1(b) (organizational victims in criminal cases) and 26.1(c) (bankruptcy case debtors and trustees). Fed. Cir. R. 47.4(a)(6).

None/Not Applicable

Dated: December 16, 2024

/s/ Ruffin B. Cordell
Ruffin B. Cordell

Samsung Electronics Co., Ltd.
Ex. 1077, p. 7

TABLE OF CONTENTS

CERTIFICATE OF INTEREST i

TABLE OF AUTHORITIES viii

STATEMENT OF RELATED CASES 1

JURISDICTIONAL STATEMENT 2

STATEMENT OF THE ISSUES 2

INTRODUCTION 4

STATEMENT OF THE CASE 4

I. The ’060 and ’160 Patents and Samsung’s Accused HBM Products 4

II. The ’054 and ’918 Patents and Samsung’s Accused DIMM Products 5

III. The ’339 Patent and Samsung’s Accused DIMM Products 7

IV. Damages 8

V. Procedural History 9

A. *Daubert* Motions 9

B. Trial 9

C. Post-Trial Briefing 10

1. ’060 and ’160 Patents 10

2. ’054 Patent 10

3. ’918 Patent 11

4. ’339 Patent 11

5. Damages 12

SUMMARY OF THE ARGUMENT 12

Samsung Electronics Co., Ltd.
Ex. 1077, p. 8

STANDARD OF REVIEW 15

ARGUMENT 16

I. Substantial Evidence Does Not Support the Jury's Infringement Verdict 16

A. The '060 and '160 Patents: Substantial Evidence Does Not Support the Jury's Finding That Samsung's Accused DRAM Dies Satisfy the District Court's Construction of "Array Die" 16

1. Netlist Lost at Claim Construction, Then Acquiesced to the District Court's Ruling 16

2. This Court Should Reject Netlist's Relitigation of the District Court's Disclaimer Holding 18

a. The District Court Did Not Invite Netlist To Rehash Prosecution Disclaimer Before the Jury 20

b. Netlist's Untimely Prosecution Disclaimer Argument Is Contrary to the Prosecution History, the '060 and '160 Patents, and Netlist's Own Admissions 21

3. Netlist Did Not Present Substantial Evidence That Samsung's Accused DRAM Dies Are "Different from a DRAM Circuit" Under the District Court's Construction 24

B. The '054 and '918 Patents 28

1. Substantial Evidence Does Not Support the Jury's Finding That Samsung's Accused Products Satisfy the '054 Patent Claims' "Second Operable State" Requirement 28

2. Substantial Evidence Does Not Support the Jury's Finding That Samsung's Accused Products Satisfy the '918 Patent's "Converter Circuit" Limitation 31

Samsung Electronics Co., Ltd.
Ex. 1077, p. 9

a. Netlist's Literal Infringement Theory Applied a Purely Functional Definition of "Converter Circuit" 31

b. The Record Lacks Substantial Evidence To Support Netlist's Assertion of Literal Infringement Under the Plain Meaning of "Converter Circuit" 33

c. Substantial Evidence Does Not Support Netlist's Alternative, Conclusory DOE Argument 35

C. The '339 Patent: Substantial Evidence Does Not Support the Jury's Finding That Samsung's Accused Products Satisfy the Multiple "Byte-Wise Data Paths" Limitation 37

1. Netlist's Faulty Interpretation of "Byte-Wise" Does Not Provide a Basis To Sustain the Infringement Verdict 38

2. The District Court Erroneously Eviscerated the "Byte-Wise" Limitation 40

II. Substantial Evidence Does Not Support the Damages Verdict 42

A. Mr. Kennedy Did Not Calculate Samsung's Incremental Profits and Gave Netlist 100% of the Alleged Incremental Revenue 43

1. Mr. Kennedy Calculated Revenues, Not Profits 43

2. Mr. Kennedy Applied a 100%-of-Revenue Rule of Thumb 46

B. The Record Lacks Evidence of Proper Apportionment 48

1. The Alleged Technical Benefit of the '054 and '918 Patents Includes Unpatented and/or Prior-Art Features 49

Samsung Electronics Co., Ltd.
Ex. 1077, p. 10

2. The Alleged Technical Benefit of the ’339 Patent Is Untethered to the Patent and Was in the Prior Art 51

3. The Alleged Technical Benefit of the ’060 and ’160 Patents Is Untethered to the Patents 52

4. Netlist Failed To Present Evidence That Mr. Kennedy’s Noninfringing Alternatives Were Available and Acceptable 53

C. The Court Should Not Give Netlist a “Do-Over” 55

III. Allowing Mr. Kennedy To Rely on a Non-Comparable License and a Flawed Regression Model Was an Abuse of Discretion 58

A. The District Court Should Have Excluded Mr. Kennedy’s Comparable License Analysis 58

1. Netlist’s Technical Expert Did Not Analyze 99.4% of the Licensed Rambus Patents 58

2. Mr. Kennedy’s Alleged Reason for Ignoring Licenses to the Asserted Patents Applies Equally to the Rambus License 61

B. Netlist’s Reliance on a Non-Analogous, Unreliable Regression Analysis Requires a New Damages Trial for the ’054 and ’918 Patents 62

1. Mr. Kennedy Was Merely a Mouthpiece for Dr. Groehn, Who Did Not Testify 62

2. The Regression Analysis Says Nothing About DDR5 Prices 64

3. The Regression Analysis’s Variables Are Incomplete and Collinear 65

CONCLUSION 68

ADDENDUM 70

Samsung Electronics Co., Ltd.
Ex. 1077, p. 11

CERTIFICATE OF SERVICE AND FILING i

CERTIFICATE OF COMPLIANCE ii

CERTIFICATE OF CONFIDENTIAL INFORMATION iii

Under Federal Circuit Rule 25.1(B), Samsung states as follows: The material redacted from Appx7, Appx12-16, Appx18, Appx23, Appx24, Appx28, Appx32, Appx33, Appx35-37, Appx51, Appx52, and Appx55 is technical information regarding the accused Samsung products that was filed under seal before the district court. The material redacted from Appx 42, Appx44, Appx52, and Appx53 is licensing information that was filed under seal before the district court.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 12

TABLE OF AUTHORITIES

Page(s)

Cases

ActiveVideo Networks, Inc. v. Verizon Commc'ns, Inc.,
694 F.3d 1312 (Fed. Cir. 2012) 17, 31

Alexsam, Inc. v. IDT Corp.,
715 F.3d 1336 (Fed. Cir. 2013) 30

Applied Med. Res. Corp. v. U.S. Surgical Corp.,
312 F. App'x 326 (Fed. Cir. 2009) 36

Aqua Shield v. Inter Pool Cover Team,
774 F.3d 766 (Fed. Cir. 2014) 43

Aro Mfg. Co. v. Convertible Top Replacement Co.,
377 U.S. 476 (1964).....48

Banks v. United States,
721 F. App'x 928 (Fed. Cir. 2017) 65

Baran v. Med. Device Techs., Inc.,
616 F.3d 1309 (Fed. Cir. 2010) 41

Becton, Dickinson & Co. v. Tyco Healthcare Grp., LP,
616 F.3d 1249 (Fed. Cir. 2010) 30

Broadcom Corp. v. Qualcomm Inc.,
543 F.3d 683 (Fed. Cir. 2008) 20, 31, 33

Commil USA, LLC v. Cisco Sys., Inc.,
813 F.3d 994 (Fed. Cir. 2015) 30

CommScope Techs. LLC v. Dali Wireless Inc.,
10 F.4th 1289 (Fed. Cir. 2021) 34

Contour IP Holding, LLC v. GoPro, Inc.,
2020 WL 5106845 (N.D. Cal. Aug. 31, 2020) 47

Samsung Electronics Co., Ltd.
Ex. 1077, p. 13

Cordis Corp. v. Bos. Sci. Corp.,
561 F.3d 1319 (Fed. Cir. 2009) 20

Cordis Corp. v. Bos. Sci. Corp.,
658 F.3d 1347 (Fed. Cir. 2011) 15

Cox Commc'ns, Inc. v. Sprint Commc'n Co. LP,
838 F.3d 1224 (Fed. Cir. 2016) 32

In re Ecast, Inc.,
96 F. App'x 710 (Fed. Cir. 2004) 54, 55

Echeverry v. Jazz Casino Co.,
988 F.3d 221 (5th Cir. 2021) 15

Ericsson, Inc. v. D-Link Sys., Inc.,
773 F.3d 1201 (Fed. Cir. 2014) 31

Exergen Corp. v. Wal-Mart Stores, Inc.,
575 F.3d 1312 (Fed. Cir. 2009) 25

Factory Mut. Ins. Co. v. Alon USA L.P.,
705 F.3d 518 (5th Cir. 2013) 63

Finjan, Inc. v. Blue Coat Sys., Inc.,
879 F.3d 1299 (Fed. Cir. 2018) 57

Garretson v. Clark,
111 U.S. 120 (1884).....49, 50, 51

Gen. Elec. Co. v. Wabash Appliance Corp.,
304 U.S. 364 (1938).....31

Georgia-Pacific Corp. v. U.S. Plywood Corp.,
318 F. Supp. 1116 (S.D.N.Y. 1970) 9, 42, 58, 61

Georgia-Pacific Corp. v. U.S. Plywood-Champion Papers, Inc.,
446 F.2d 295 (2d Cir. 1971) 46

Grain Processing Corp. v. Am. Maize-Prods. Co.,
185 F.3d 1341 (Fed. Cir. 1999) 54, 55

Samsung Electronics Co., Ltd.
Ex. 1077, p. 14

Hewlett-Packard Co. v. Mustek Sys., Inc.,
340 F.3d 1314 (Fed. Cir. 2003) 35

Holland Furniture Co. v. Perkins Glue Co.,
277 U.S. 245 (1928).....31

Hughes Tool Co. v. Dresser Indus., Inc.,
816 F.2d 1549 (Fed. Cir. 1987) 46

Intel Corp. v. Qualcomm Inc.,
21 F.4th 801 (Fed. Cir. 2021) 39

Johnson v. United States,
318 U.S. 189 (1943).....18

Kinetic Concepts, Inc. v. Blue Sky Med. Grp., Inc.,
554 F.3d 1010 (Fed. Cir. 2009) 29

Labyrinth Optical Techs. LLC v. Alcatel-Lucent USA, Inc.,
2015 WL 12696081 (C.D. Cal. Sept. 2, 2015) 47

LaserDynamics, Inc. v. Quanta Comput., Inc.,
694 F.3d 51 (Fed. Cir. 2012) .....*passim*

Lindemann Maschinenfabrik GmbH v. Am. Hoist & Derrick Co.,
895 F.2d 1403 (Fed. Cir. 1990) 46

Liquid Dynamics Corp. v. Vaughan Co.,
449 F.3d 1209 (Fed. Cir. 2006) 25, 27

Looksmart Grp., Inc. v. Microsoft Corp.,
2019 WL 4009263 (N.D. Cal. Aug. 5, 2019) 47

Lucent Techs., Inc. v. Gateway, Inc.,
580 F.3d 1301 (Fed. Cir. 2009) 56, 58

Mike's Train House, Inc. v. Lionel, L.L.C.,
472 F.3d 398 (6th Cir. 2006) 63

Mirror Worlds, LLC v. Apple Inc.,
692 F.3d 1351 (Fed. Cir. 2012) 15

Samsung Electronics Co., Ltd.
Ex. 1077, p. 15

MLC Intellectual Property, LLC v. Micron Technology, Inc.,
10 F.4th 1358 (Fed. Cir. 2021) 59

MobileMedia Ideas LLC v. Apple Inc.,
780 F.3d 1159 (Fed. Cir. 2015) 32, 39

MTD Prods. Inc. v. Iancu,
933 F. 3d 1336 (Fed. Cir. 2019) 33, 34

Netlist, Inc. v. Samsung Elecs. Co.,
No. 2:22-CV-00293-JRG, Dkt. 818 (E.D. Tex. Nov. 6, 2024) 59

Network Prot. Scis., LLC v. Fortinet, Inc.,
2013 WL 5402089 (N.D. Cal. Sept. 26, 2013).....56

Nordock Inc. v. Sys. Inc.,
2013 WL 989864 (E.D. Wis. Mar. 13, 2013).....47

Omega Pats., LLC v. CalAmp Corp.,
13 F.4th 1361 (Fed. Cir. 2021) 50, 52, 53

Philp v. Nock,
84 U.S. 460 (1873).....56, 57

Promega Corp. v. Life Technologies Corp.,
875 F.3d 651 (Fed. Cir. 2017) 55, 57

Realcomp II, Ltd. v. F.T.C.,
635 F.3d 815 (6th Cir. 2011) 66, 67

ResQNet.com, Inc. v. Lansa, Inc.,
594 F.3d 860 (Fed. Cir. 2010) 44, 51, 58

Robroy Indus.-Tex., LLC v. Thomas & Betts Corp.,
2017 WL 1319553 (E.D. Tex. Apr. 10, 2017).....62, 63

Sheehan v. Purolator, Inc.,
839 F.2d 99 (2d Cir. 1988) 66, 67

Smith v. Transworld Drilling Co.,
773 F.2d 610 (5th Cir. 1985) 15

Samsung Electronics Co., Ltd.
Ex. 1077, p. 16

TEK Global, S.R.L. v. Sealant Sys. Int'l, Inc.,
920 F.3d 777 (Fed. Cir. 2019) 21, 24

Tex. Instruments Inc. v. Cypress Semiconductor Corp.,
90 F.3d 1558 (Fed. Cir. 1996) 37

Transclean Corp. v. Bridgewood Servs., Inc.,
290 F.3d 1364 (Fed. Cir. 2002) 45

Trell v. Marlee Electronics Corp.,
912 F.2d 1443 (Fed. Cir. 1990) 59, 60

U.S. Steel Grp. v. United States,
225 F.3d 1284 (Fed. Cir. 2000) 43

Uniloc USA, Inc. v. Microsoft Corp.,
632 F.3d 1292 (Fed. Cir. 2011) 45, 47, 60

United States v. Grey Bear,
883 F.2d 1382 (8th Cir. 1989) 63

United States v. Musquiz,
45 F.3d 927 (5th Cir. 1995) 18

Unwired Planet, LLC v. Apple Inc.,
2017 WL 589195 (N.D. Cal. Feb. 14, 2017) 56

VirnetX, Inc. v. Cisco Sys., Inc.,
767 F.3d 1308 (Fed. Cir. 2014) 36, 47

VLSI Tech. LLC v. Intel Corp.,
87 F.4th 1332 (Fed. Cir. 2023) 43, 45, 49

Warner-Jenkinson Co. v. Hilton Davis Chem. Co.,
520 U.S. 17 (1997).....32, 36

Yoon Ja Kim v. ConAgra Foods, Inc.,
465 F.3d 1312 (Fed. Cir. 2006) 39

Statutes

35 U.S.C. §112, ¶6 .....13, 32

35 U.S.C. § 284 48

Samsung Electronics Co., Ltd.
Ex. 1077, p. 17

Other Authorities

Fed. R. Evid. 703 63

William F. Lee & Mark A. Lemley, *The Broken Balance: How "Built-In Apportionment" and the Failure to Apply* Daubert *Have Distorted Patent Infringement Damages*, 37 Harv. J.L. & Tech. 255 (2024).....67

Samsung Electronics Co., Ltd.
Ex. 1077, p. 18

STATEMENT OF RELATED CASES

Pursuant to Federal Circuit Rule 47.5, Samsung states that no appeal from the same civil action has been previously brought before this or any other appellate court, and that the following are cases that may directly affect or be directly affected by the Court's decision in this appeal: *Netlist, Inc. v. Samsung Electronics Co., Ltd.*, CAFC No. 24-1707, originating from Patent Trial & Appeal Board ("PTAB") Nos. IPR2022-00639, IPR2023-00204; *Netlist, Inc. v. Samsung Electronics Co., Ltd.*, CAFC Nos. 24-1859, 24-1863, originating from PTAB Nos. IPR2022-00996, IPR2023-00406, IPR2022-00999, and IPR2023-00405; *Netlist, Inc. v. Samsung Electronics Co., Ltd.*, CAFC No. 24-2240, originating from PTAB No. IPR2022-01427; *Netlist, Inc. v. Samsung Electronics Co., Ltd.*, CAFC No. 24-2241, originating from PTAB No. IPR2022-01428; *Samsung Electronics Co., Ltd. v. Netlist Inc.*, No. 1:21-cv-01453-RGA-JLH (D. Del.); *Netlist Inc. v. Samsung Electronics Co., Ltd.*, No. 8:20-cv-00993-MCS-ADS (C.D. Cal.), and eponymous appeal Nos. 22-55209, 22-55247 (9th Cir.); *Netlist, Inc. v. Micron Technology, Inc.*, No. 2-23-cv-00628 (E.D. Tex.); *Micron Technology, Inc. v. Netlist, Inc.*, No. IPR2023-00882 (PTAB); *Micron Technology, Inc. v. Netlist, Inc.*, No. IPR2023-00883 (PTAB).

JURISDICTIONAL STATEMENT

The district court had jurisdiction under 28 U.S.C. §§1331 and 1338(a); entered final judgment on August 11, 2023; and decided the last post-judgment motion on July 15, 2024. This appeal, noticed August 9, 2024, is timely. *See* 28 U.S.C. §2107(a); Fed. R. App. P. 4. This Court has jurisdiction under 28 U.S.C. §1295(a)(1).

STATEMENT OF THE ISSUES

1. Does substantial evidence support the jury's infringement verdict for the '060 and '160 patents when Netlist's infringement evidence for the "array die" term is based entirely on a prosecution-history-disclaimer argument Netlist lost at claim construction?

2. Does substantial evidence support the jury's infringement verdict for the '054 patent when the accused memory module's alleged "second ***operable*** state" is an ***in***operable state?

3. Does substantial evidence support the jury's infringement verdict for the '918 patent when Netlist asserted only a purely functional theory that "converter circuit" covers anything that alters a voltage—even a simple wire, adduced no evidence of infringement under the plain meaning of "converter circuit," and ignored evidence that a skilled artisan would understand "converter

Samsung Electronics Co., Ltd.
Ex. 1077, p. 20

circuit" as a particular class of circuits distinct from the Samsung component Netlist says meets the limitation?

4. Does substantial evidence support the jury's infringement verdict for the '339 patent when Netlist's theory eviscerated the claim requirement for multiple byte-wise data paths by alleging that a single byte-wise data path satisfies that limitation?

5. Should this Court instruct the district court to enter JMOL of no damages because Netlist (a) based its damages theory on Samsung's incremental ***revenue*** allegedly associated with the patented features, demanded ***100%*** of such revenue, and ignored Samsung's associated incremental costs; and (b) failed to apportion such revenue between patented and unpatented features?

6. Is a new trial on damages necessary because the district court abused its discretion by allowing Netlist's damages expert to (a) present a $1.1-billion non-comparable license to the jury, while ignoring actual licenses for the asserted patents, and (b) parrot the flawed conclusions of a non-testifying regression expert?

Samsung Electronics Co., Ltd.
Ex. 1077, p. 21

INTRODUCTION

Netlist won an Eastern District of Texas jury trial by presenting legally untenable infringement and damages theories. On infringement, Netlist ignored the district court's claim constructions and the plain language of claim terms, asserting instead its own unsupported and impermissible interpretations. On damages, Netlist advanced an unprecedented ***100%-of-Revenue Rule***, under which Samsung would have agreed to give Netlist 100% of the incremental revenue (the revenue allegedly attributable to the patented features) in a hypothetical negotiation. Netlist also failed to present adequate evidence on apportionment. Each of these errors requires this Court's intervention.

STATEMENT OF THE CASE

The jury found that certain Samsung memory products infringe U.S. Patents 8,787,060; 9,318,160; 10,949,339, 11,016,918; and 11,232,054. The PTAB later held ***all*** asserted claims of these patents invalid in IPR Final Written Decisions. Netlist's appeals of those IPR decisions are companion cases to this appeal. *Netlist, Inc. v. Samsung Electronics Co.*, Nos. 24-1707, 24-1859, 24-1863, 24-2240, 24-2241.

I. The '060 and '160 Patents and Samsung's Accused HBM Products

The '060 and '160 patents disclose memory packages having multiple memory dies (called "array dies") that send signals to a ***subset*** of those dies,

Samsung Electronics Co., Ltd.
Ex. 1077, p. 22

supposedly to reduce the package drivers' electrical load. Appx698-699; Appx706 (1:18-21).[1] The patents describe memory packages having any memory type, including DRAM, SDRAM, SRAM, flash memory, or DDR SDRAM. Appx708 (5:1-5); Appx716 (22:2-4). The array dies may employ various interconnects, including "through-silicon vias (TSV), conducting rods, wire bonds, and pins," Appx708 (5:51-54), or "a wire, a conducting rod, or a conducting trace," Appx709 (8:22-27). Based on prosecution disclaimer, the district court construed "'array die' as 'array die that is ***different from a DRAM circuit***.'"[2] Appx1380-1381.

Netlist accused various Samsung HBM products of infringing the '060 patent (claims 1, 5, and 6) and '160 patent (claim 1). Appx5479; Appx5484. The accused HBM memory arrays contain millions of the DRAM circuits that Netlist disclaimed during prosecution, and so the accused products lack the claimed "array die." Appx5489-5490; Appx5512; Appx5553-5554. Samsung argued noninfringement under this disclaimer.

II. The '054 and '918 Patents and Samsung's Accused DIMM Products

The shared '054 and '918 patent specification[3] describes a hybrid memory module that transfers data between volatile DRAM memory and non-volatile flash

[1] The '060 and '160 patents share a specification; this brief cites the '060 specification. *See* Appx696; Appx720.

[2] All emphasis added.

[3] This brief cites the '918 specification.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 23

memory on the module if power begins to fail. Appx627-628 (1:65-4:27); Appx605 (Fig. 3A). DRAM memory requires power to maintain stored information; flash memory does not. Appx628 (3:53-61). To avoid data loss when a module loses power, the patents describe using residual power to transfer data from DRAM to flash memory. *See* Appx628 (3:59-62); Appx638-639 (24:60-25:7). This "last gasp" state is key to these patents. *See* Appx638-639 (24:60-25:7).

Each asserted '054 patent claim requires the module to "transition[] from a first operable state to a ***second operable state***," reflecting the shift from normal to power-failure mode. Appx694 (40:19-20); Appx687 (25:3-7). The parties agreed "second operable state" means a "state in which the ***memory module is operated*** after transition." Appx1241.

Each '918 patent asserted claim requires a "***converter circuit*** configured to provide a fourth regulated voltage." *E.g.*, Appx645 (38:31-32). The patents disclose three converter circuits: buck, boost, and buck-boost converters. Appx641 (29:23-27); Appx620. Each disclosed converter circuit includes a switch for "continually switching on and off" the voltage "to provide charge to an inductor." Appx5909-5910.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 24

Netlist accused Samsung's SODIMM, UDIMM, and RDIMM products[4] of infringing the '054 patent's claims 16-17 and the '918 patent's claims 16 and 18-19, and Samsung's RDIMM products of infringing the '918 patent's claims 1, 5, and 13. Appx5336; Appx5343-5345; Appx5383-5384. The accused products neither transfer data to onboard backup storage under power failure, nor have a "second operable state." Further, each accused product contains a low drop-out ("LDO") regulator, which "receives the pre-regulated input voltage and produces a regulated output voltage." Appx5338; Appx5667; Appx5970. LDO regulators adjust voltages by discarding energy as heat rather than using reactive buck/boost technologies, and so the accused products lack the claimed "converter circuit." Appx5909; Appx5911.

III. The '339 Patent and Samsung's Accused DIMM Products

The '339 patent discloses memory modules having multiple memory devices arranged into different "ranks." Appx564. Asserted claims 1, 8, and 9 recite a memory module comprising "a plurality of byte-wise buffers," referring to eight-bit wide memory buffers. Appx595 (19:40-52); Appx592 (13:43-53, 14:15-18); Appx5446 (443:9-12). The specification expressly distinguishes eight-bit from

[4] "SODIMM" is small-outline dual in-line memory module (DIMM); "UDIMM" is unbuffered DIMM; "RDIMM" is registered DIMM. Appx5199 (199:11-17); Appx5146 (146:16-18).

Samsung Electronics Co., Ltd.
Ex. 1077, p. 25

four-bit data paths. Appx592 (13:57-14:8); Appx594 (17:30-32); Appx595 (20:6-14).

The claims recite a "byte-wise buffer" having a "first side" coupled to the computer and "second side" coupled to multiple "ranks" of DDR DRAM devices. Appx595 (19:40-61). When writing data to memory, the byte-wise buffer "drive[s] a respective byte-wise section" of the write data from the first to the second side of the buffer. *Id.*; Appx1356. The district court construed "drive" to mean "enabl[e] only one of the data paths while the other possible paths are disabled." Appx1359. Thus, the district court's construction requires multiple byte-wise write data paths.

Samsung's accused DDR4 LRDIMM products include only a single byte-wise write data path—not multiple paths—through the buffer, and all eight bits are driven down a single path. Appx5451-5452; Appx5454; Appx5944 (924:9-12). As such, the accused products do not have the required multiple byte-wise (i.e., eight-bit) data paths.

IV. Damages

Applying the hypothetical-negotiation framework, Netlist's damages expert, David Kennedy, opined that Samsung would have agreed to pay Netlist ***100%*** of the ***revenue*** associated with the accused features, without allowing deductions for Samsung's costs or risk. Appx5781; *see also* Appx5704-5708. Mr. Kennedy did not divide any supposed value associated with the patents between the parties

Samsung Electronics Co., Ltd.
Ex. 1077, p. 26

using the *Georgia-Pacific* factors. Instead, he calculated the alleged revenue attributable to the patented features and "gave all of that money to Netlist" ($404.2 million for 16 months of sales). Appx5781; Appx5721; Appx5733. Mr. Kennedy thus opined that Samsung would have ignored its costs and agreed to ***a guaranteed loss*** by giving ***all*** incremental revenue to Netlist—a nonsensical proposition.

V. Procedural History

A. *Daubert* Motions

Samsung moved to exclude Mr. Kennedy's damages theory because it sought 100% of Samsung's incremental revenue, without considering Samsung's incremental costs and risks, and did not apportion out all unpatented features or the contributions made by Samsung and other leading memory companies. Appx1481-1483; Appx1487-1488; Appx264 (199:17-19). Samsung also moved to exclude Mr. Kennedy's reliance on a non-comparable license and his parroting of the opinions of an expert, Andreas Groehn, who was scheduled to testify but never took the stand. Appx1488-1490; Appx1443-1448. The district court denied those motions. Appx282-283; Appx347; Appx2111-2112.

B. Trial

Following a six-day trial, the jury found infringement of the asserted patents and awarded Netlist $303,150,000. *See* Appx2326; Appx2328-2331.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 27

C. Post-Trial Briefing

1. ’060 and ’160 Patents

Samsung moved for JMOL of noninfringement and a new trial, arguing Netlist failed to show infringement under the district court’s construction of “array die.” *See* Appx18029-18033; Appx18046-18047. The district court denied Samsung’s motion, citing three lines of testimony in which Netlist’s expert nakedly concluded that Samsung’s DRAM dies qualify as “array dies” while he repudiated the district court’s claim constructions. Appx14 (quoting Appx5519). Netlist identified no evidence that Samsung’s accused DRAM dies are “different from a DRAM circuit” under its plain meaning. Appx18101-18105; Appx18214-18215. Netlist instead relitigated the same “DRAM circuit” prosecution disclaimer issue that it lost at claim construction. Appx18102 n.1; Appx18103-18104; Appx18214-18215; Appx1291-1292.

2. ’054 Patent

Samsung moved for JMOL of noninfringement based on the “second operable state” limitation. Appx18038-18040. The district court denied Samsung’s motion, reasoning that “whether a memory module with some low-level circuitry that is operating” is in an “operable” state is a factual dispute that the jury resolved. Appx22.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 28

3. '918 Patent

Samsung moved for JMOL of noninfringement based on the "converter circuit" limitation. Appx18034-18038. The district court denied Samsung's motion. Appx16-21. For literal infringement, the court cited Dr. Mangione-Smith's cursory testimony "identif[ying] the LDO as the converter circuit" as substantial evidence supporting infringement. Appx19 (citing Appx5339). The court rejected Samsung's argument that Dr. Mangione-Smith improperly applied a purely functional definition of converter circuit, reasoning the term "requires electrical components for voltage conversion," this component set is "broad," and there is no "authority that prevents structure from being defined broadly." Appx19. For DOE, the court found that Dr. Mangione-Smith's one-sentence testimony was not conclusory because it must "be considered, not in utter isolation, but against the backdrop of the entirety of Dr. Mangione-Smith's testimony." Appx20.

4. '339 Patent

Samsung moved for JMOL of noninfringement because Netlist failed to show that Samsung's accused products satisfy the multiple byte-wise write data paths requirement. Appx18022-18028. The district court denied Samsung's motion, reasoning that "[t]he claims distinguish between 'byte-wise' and 'byte-wide,' indicating that a 'byte-wise data path[]' does not need to be eight bits wide,"

Samsung Electronics Co., Ltd.
Ex. 1077, p. 29

and "the jury heard evidence of two paths—the two nibbles," i.e., two four-bit paths. Appx8.

5. Damages

Samsung moved for JMOL of no damages or a new damages trial because, *inter alia*, substantial evidence does not support Mr. Kennedy's theory that Samsung would have agreed to give Netlist 100% of the revenue associated with the patents. Appx18058-18075. The court denied the motion. Appx41-60.

SUMMARY OF THE ARGUMENT

This Court should reverse the denial of JMOL of noninfringement of the asserted patents. In each case, Netlist substituted the district court's claim construction and/or the plain meaning of the claim language with Netlist's own impermissible and/or unsupported interpretation.

For the '060 and '160 patents, the district court held at *Markman* that Netlist disavowed all "DRAM circuits," and thus to prove infringement Netlist had to show that Samsung's accused HBM products have memory arrays that are "different from a DRAM circuit." Samsung's HBM memory arrays are indisputably made up of DRAM dies, and so Netlist's expert failed to faithfully apply that construction and refused to provide any opinion on the plain meaning of "DRAM circuit" or "different from a DRAM circuit." Instead, he impermissibly relitigated the disclaimer before the jury and presented an infringement theory

Samsung Electronics Co., Ltd.
Ex. 1077, p. 30

based entirely on Netlist's narrowing disclaimer argument that the district court rejected at *Markman*.

For the '054 patent, Netlist ignored the claimed "second operable state," which the parties construed as a "state in which the memory module is ***operated*** after transition." Rather than identify a state where the memory module is ***operable*** to do anything, Netlist accused a state where the module is turned off and ***inoperable*** (i.e., the module is not powered, all memory functionality is inoperable, and data stored on DRAM is erased). Netlist relied on "wake up" circuitry on the module that listens for a power-on signal, but this is not an operable state and thus, there can be no infringement.

For the '918 patent, Netlist ignored evidence that "converter circuit" describes a class of circuits that are readily identifiable in the art, treated "converter circuit" as a wholly functional concept encompassing anything that alters voltage, and never invoked §112, ¶6. Netlist then mapped the '918 patent's "converter circuit" onto a structure that dissipates energy, but is not a converter circuit under the plain meaning of the term.

For the '339 patent, Netlist eviscerated the claim requirements for multiple byte-wise (i.e., eight-bit) data paths. Only by doing so could Netlist accuse Samsung's products—which undisputedly contain only a single byte-wise data path—of infringement.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 31

Damages. The district court erred in denying JMOL of no damages or a new damages trial. First, although the law requires a hypothetical negotiation to examine an infringer's incremental ***profits*** from using patented technology, Netlist instead focused on Samsung's incremental ***revenues*** and ignored Samsung's incremental costs and risk. Second, Netlist applied a 100%-of-revenue rule of thumb whereby Netlist received ***all*** alleged incremental revenue in the hypothetical negotiation and bore no incremental costs. Netlist's theory eliminates any incentive to include the accused features by guaranteeing Samsung a loss on each sale. Netlist's theory also violates prohibitions on (1) profit disgorgement (let alone revenue disgorgement) for utility patent infringement, and (2) discredited "rule of thumb" approaches. Third, Netlist failed to apportion Samsung's alleged incremental revenues because it presented no evidence that those revenues were attributable solely to patented features.

The district court's denial of Samsung's *Daubert* motions also warrants a new damages trial. Netlist's damages expert should not have been permitted to (1) rely on a non-comparable, $1.1-billion license to skew the jury's damages horizon while ignoring real-world licenses to the asserted patents or (2) use another expert's flawed, out-of-court regression analysis of an old memory generation to compute damages for a modern memory generation. In subsequent litigation

Samsung Electronics Co., Ltd.
Ex. 1077, p. 32

between Samsung and Netlist, the district court excluded the same $1.1-billion license.

STANDARD OF REVIEW

This Court reviews denial of JMOL under regional circuit law. *Mirror Worlds, LLC v. Apple Inc.*, 692 F.3d 1351, 1356 (Fed. Cir. 2012). The Fifth Circuit reviews denial of JMOL *de novo*, and reverses if substantial evidence does not support the jury's finding. *Id.* at 1356-57. Substantial evidence requires "more than a mere scintilla," and the Court "must review the record as a whole, taking into consideration evidence that both justifies and detracts from the jury's decision." *Cordis Corp. v. Bos. Sci. Corp.*, 658 F.3d 1347, 1357 (Fed. Cir. 2011); *Echeverry v. Jazz Casino Co.*, 988 F.3d 221, 228 (5th Cir. 2021). The Court may grant a new trial if the verdict is against the weight of the evidence. *See Smith v. Transworld Drilling Co.*, 773 F.2d 610, 612-13 (5th Cir. 1985). The Court reviews a district court's evidentiary rulings according to Fifth Circuit law for an abuse of discretion. *LaserDynamics, Inc. v. Quanta Comput., Inc.*, 694 F.3d 51, 66 (Fed. Cir. 2012).

Samsung Electronics Co., Ltd.
Ex. 1077, p. 33

ARGUMENT

I. Substantial Evidence Does Not Support the Jury's Infringement Verdict

Because substantial evidence does not support the jury's infringement verdict, this Court should reverse the denial of JMOL of noninfringement. Alternatively, the Court should order a new trial because the verdict is against the weight of the evidence.

A. The '060 and '160 Patents: Substantial Evidence Does Not Support the Jury's Finding That Samsung's Accused DRAM Dies Satisfy the District Court's Construction of "Array Die"

1. Netlist Lost at Claim Construction, Then Acquiesced to the District Court's Ruling

The asserted claims describe memory packages comprising stacked memory array "dies," which are the individual memory chips containing memory cells.[5] Claim 1 of the '160 patent includes a representative limitation: "stacked array dies including a first group of array dies and a second group of at least one array die." Appx717 (23:62-64).

Because the prior art was replete with decades-old examples of stacked ***DRAM*** memory dies, Netlist had to disclaim DRAM circuits from its claimed "dies." The district court held that Netlist "clear[ly] and unambiguous[ly] disavow[ed]" "DRAM circuits" during prosecution by "distinguish[ing] 'stacked

[5] Semiconductor wafers containing hundreds to thousands of memory chips are "diced" into individual "dies."

Samsung Electronics Co., Ltd.
Ex. 1077, p. 34

DRAM circuits' from 'stacked array dies'" without limitation to overcome a prior-art reference ("Rajan137") (Appx16339-16366). Appx1380-1381. The distinction Netlist used to gain allowance was the fact that the Rajan137 dies contain DRAM circuits, not that the dies are interconnected in any particular way such as through the use of external wire bonding. *Id.*; Appx1229. The district court noted that "Netlist did not explain the structural difference . . . or even how they might be stacked differently." Appx1380-1381.

The *Markman* order thus construed "'array die' as 'array die that is different from a DRAM circuit.'" Appx1380; Appx1291-1292 (noting experts could address plain meaning of "DRAM circuit"). In so doing, the district court rejected Netlist's argument that it had narrowly disclaimed only DRAM circuits that were described in Rajan137. *See, e.g.*, Appx1153-1154; *ActiveVideo Networks, Inc. v. Verizon Commc'ns, Inc.*, 694 F.3d 1312, 1325 (Fed. Cir. 2012) (rejection of a proposed construction resolved claim construction dispute).

Netlist initially objected to the district court's disclaimer holding, asking the court to narrow the disavowal from "DRAM circuits" to only Rajan137's specific DRAM circuits. Appx1386. Netlist later ***affirmatively withdrew this objection***. Appx1394-1397. As a result, Netlist waived all arguments about the scope of its prosecution disclaimer and must accept that the claimed "array dies" exclude ***all*** DRAM circuits. *See Johnson v. United States*, 318 U.S. 189, 200 (1943)

Samsung Electronics Co., Ltd.
Ex. 1077, p. 35

("[P]etitioner expressly waived any objection to the prosecutor's comment by withdrawing his exception to it."); *United States v. Musquiz*, 45 F.3d 927, 931 (5th Cir. 1995) ("By withdrawing the challenge, Gatewood waived his objection.").

2. This Court Should Reject Netlist's Relitigation of the District Court's Disclaimer Holding

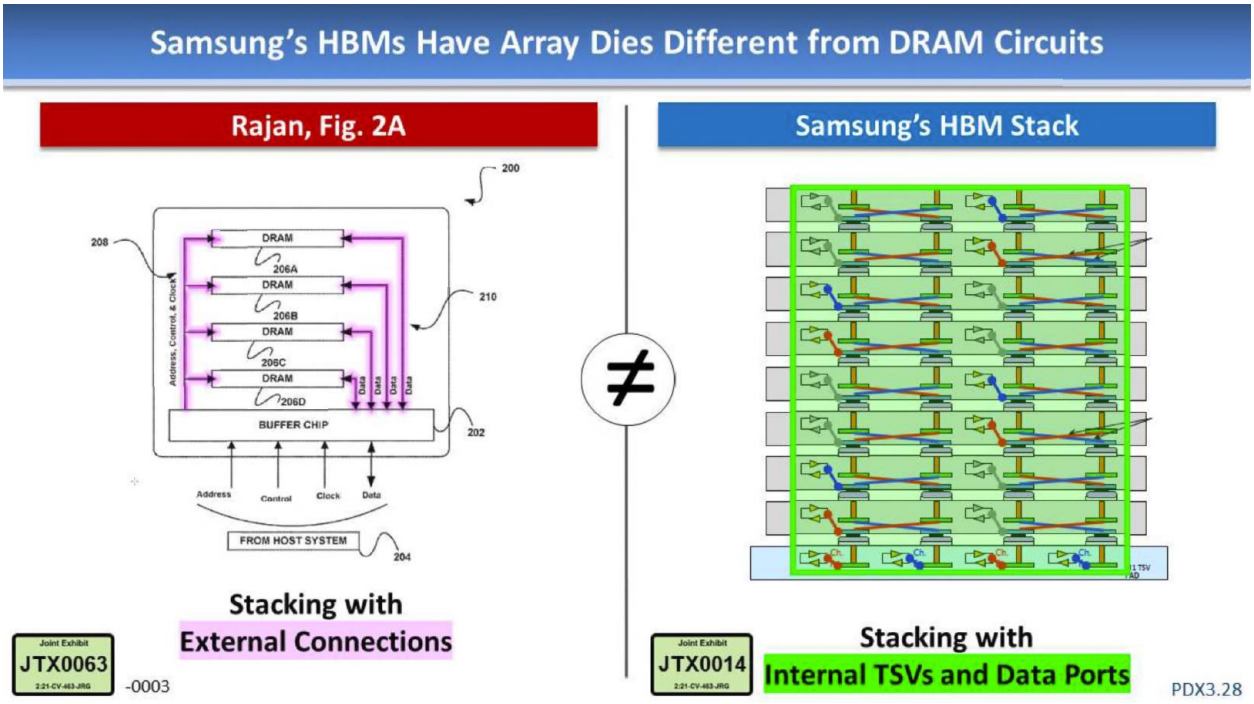
Perhaps confronted with the realization that it could not prove infringement under the district court's "different from a DRAM circuit" disclaimer, Netlist sought to relitigate its prosecution disclaimer during trial.

Specifically, through its expert, Dr. Brogioli, Netlist attempted to distract the jury from the actual disclaimer by introducing a new complexity involving the way the memory array dies were wired to each other. Netlist argued at trial that it had not disclaimed memory array dies containing DRAM circuits generally, but rather only those array dies that were wired together in the precise manner shown in Rajan137, using "wire bonding." *See, e.g.*, Appx5624-5625. Memory array dies can be connected using a variety of techniques including through-silicon-vias (TSVs), or through externally attached wire bonds. Appx708 (5:51-53); Appx709 (8:19-34); Appx6063; Appx6065. Because Rajan137 allegedly shows wire bonds, Netlist argued that only memory array dies connected with wire bonds were

Samsung Electronics Co., Ltd.
Ex. 1077, p. 36

disclaimed. *See* Appx5624-5625.[6]

Dr. Brogioli then opined that Samsung's DRAM dies meet the construction "array die that is different from a DRAM circuit" because they are ***connected differently*** from Rajan137's alleged connections. Appx5515; Appx5517-5518; *see also* Appx5509-5519 (506:2-516:14). He even used demonstratives like the following to make the entirely irrelevant comparison between the connections between Samsung's accused DRAM dies and Rajan137's DRAM dies:



Appx22117.

[6] At trial Netlist framed the infringement determination as a claim construction exercise: "The relevant question is, ***what does this term mean in the intrinsic record***." Appx5505-5506.

a. The District Court Did Not Invite Netlist To Rehash Prosecution Disclaimer Before the Jury

To justify its attempt to relitigate the disclaimer issue it lost during claim construction, Netlist cited the district court's statements during the *Markman* hearing that the meaning of DRAM circuit "would be for the experts" and that the court did "not intend to try and define what DRAM circuits that are disclaimed." Appx1291-1292 (33:23-34:8).

However, in stating that experts could discuss the plain meaning of "DRAM circuit," the district court simply acknowledged established law: after claim construction, the plaintiff must apply the plain meaning of "DRAM circuit" to the accused products. *Broadcom Corp. v. Qualcomm Inc.*, 543 F.3d 683, 696 (Fed. Cir. 2008) (explaining when a term is not construed, the plain meaning controls). However, the district court in no way invited the parties to relitigate Netlist's prosecution disclaimer before the jury.

Nevertheless, Netlist plainly relitigated its lost disclaimer argument, contradicting the court's later-issued claim construction order, which identified a clear and unambiguous disavowal of "DRAM circuits" with no basis for any narrower disclaimer. Appx1380-1381; Appx6327 (1300:13-25) (jury instructions). Nothing about the *Markman* hearing authorized Netlist to reargue claim construction issues under the guise of providing the plain meaning of "DRAM circuit." *See Cordis Corp. v. Bos. Sci. Corp.*, 561 F.3d 1319, 1337 (Fed. Cir. 2009)

Samsung Electronics Co., Ltd.
Ex. 1077, p. 38

("[I]t is improper to argue claim construction to the jury because the 'risk of confusing the jury is high when experts opine on claim construction.'" (citation omitted)); *TEK Global, S.R.L. v. Sealant Sys. Int'l, Inc.*, 920 F.3d 777, 787 (Fed. Cir. 2019) (rejecting attempt to read the plain meaning to include narrowing requirements rejected at claim construction). Critically, in a pique of studied ignorance, Netlist's expert expressly refused to offer an opinion as to the plain meaning of "DRAM circuits." *See, e.g.,* Appx5550. Netlist thus ***refused*** to provide the jury with any interpretation of the plain meaning of this phrase, which was the only issue remaining after claim construction.

b. Netlist's Untimely Prosecution Disclaimer Argument Is Contrary to the Prosecution History, the '060 and '160 Patents, and Netlist's Own Admissions

Even excusing the impropriety of contesting claim construction at trial, Netlist's disclaimer re-argument should fail. First, Netlist unambiguously disclaimed DRAM circuits from the scope of "array dies" during prosecution. In response to the examiner's rejection over the Rajan137 reference, Netlist argued unequivocally that DRAM circuits are "different from array dies" recited in claim 1. Appx14479; Appx1380-1381. Netlist never mentioned wire bonding, or any other connection techniques. *See* Appx14470-14484. The method of interconnecting array dies simply was not relevant; instead, Netlist's arguments limited its claims to "array dies" that did not employ DRAM circuits, full stop.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 39

Notably, Netlist made none of its "wire bonding" arguments during claim construction, nor did it raise these issues in its later-withdrawn objections to the claim construction order.

Second, Rajan137 nowhere restricts DRAM circuits to externally wire-bonded circuits. Instead, Rajan137 broadly teaches that its DRAM circuits include all types of DRAM circuits, regardless of electrical connection configuration. Appx16360 (stating "[t]he ***electrical connections*** between the buffer chip 202 and the stack of DRAM circuits 206A-D may be configured in ***any desired manner***," ¶50, and disclosing "***various configurations*** of a buffered stack of DRAM circuits 206A-D," ¶43); Appx6072-6074 (1052:24-1054:10). Notably, Rajan137's two references to "wires" are exemplary rather than limiting; indeed, one of those references explicitly teaches "simulating" wires rather than using physical wires. *See* Appx16358 ("[I]n the context of electrical simulation, a number of pins, wires, signals, etc. may be simulated," ¶26); Appx16360 ("may be wired," ¶50). Rajan137 simply cannot be read as limiting "DRAM circuits" to externally wire-bonded structures.

Third, the '060 and '160 patents do not support limiting "DRAM circuits" to any particular interconnection structure disclosed in Rajan137. The asserted claims say nothing about the way the memory array dies are interconnected. The specification states that the inventions encompass all types of memory (including

Samsung Electronics Co., Ltd.
Ex. 1077, p. 40

DRAM, flash, and SRAM "to name a few") and all types of array die interconnections (including TSVs and wire bonds). Appx708 (5:1-5; 5:51-54); Appx709 (8:22-27); Appx716 (22:2-4). Netlist further conceded that Rajan137's DRAM circuits "may take any form of" DRAM, including SDRAM, DDR DRAM or "any other type of DRAM." Appx1154 (quoting Appx16357 (¶18)). Presumably, the "other than DRAM circuits" disclaimer meant to focus the claims on other species of memory such as flash, which have different circuit structures from DRAM. *See* Appx5376-5377 (373:25-374:2); Appx5431 (428:9-14); Appx5433 (430:1-8); Appx5869 (863:12-19); Appx5892 (872:2-16).

Fourth, Netlist conceded before the PTAB that DRAM circuits are ***not limited*** to external wire-bonding and instead can be connected using through-silicon vias (TSVs) like the accused products' DRAM dies are. *Contra* Appx22117 (Brogioli demonstrative comparing Samsung's accused DRAM dies and Rajan137's DRAM dies). Specifically, Netlist argued that a prior art reference called Riho has "DRAM circuits" even though its disclosed memory array dies are connected using the same TSVs as the accused products. Appx22054 (54:11-22); Appx22007-22009 (7:23-8:2, 9:3-8). In IPR, Netlist never advanced the "wire bonding" distraction before the PTAB or suggested that the connection method was relevant to the instant claims. *See* Appx22053-22054 (53:12-54:2) (explaining the district court's construction did not say "different from Rajan[137]'s DRAM

Samsung Electronics Co., Ltd.
Ex. 1077, p. 41

circuit"). Critically, Netlist's lead counsel correctly told the PTAB that, to show infringement in the district court trial, Netlist "actually had to show that what was in the HBM products was not a DRAM circuit." *Id.*

Therefore, Netlist disclaimed all DRAM circuits such that the recited "array die" covers only array dies that are different from DRAM circuits, as the district court concluded.

3. Netlist Did Not Present Substantial Evidence That Samsung's Accused DRAM Dies Are "Different from a DRAM Circuit" Under the District Court's Construction

After the district court issued its prosecution disclaimer construction, Netlist could have stipulated to noninfringement and proceeded to appeal. Alternatively, Netlist could have attempted to apply the plain meaning of "array die that is different from a DRAM circuit." Netlist did neither. Appx1651 (129:7-10); Appx5550; Appx5624-5625; Appx5505-5506. Rather, as set forth above, Netlist re-urged a version of its rejected interpretation of "DRAM circuit," with a new twist—the disclaimer should be narrowed to only ***DRAM circuits using wire bonds like those in Rajan137***. Appx5498-5499; Appx5515-5517; Appx5518; Appx5582; Appx5624-5625.

None of the evidence Netlist adduced at trial is substantial evidence of infringement under the district court's correct construction. First, Netlist's evidence was directed to the rejected disclaimer argument discussed above, and

Samsung Electronics Co., Ltd.
Ex. 1077, p. 42

that evidence legally cannot support the infringement verdict under the correct scope of the disclaimer. *See Exergen Corp. v. Wal-Mart Stores, Inc.*, 575 F.3d 1312, 1321 (Fed. Cir. 2009) ("No party may contradict the court's construction to a jury."). Netlist's evidence purported to show that Samsung's DRAM circuits connected through TSV's are different from DRAM circuits that are connected through external wire-bonds—consistent with its new narrowed disclaimer theory. *E.g.*, Appx5515; Appx5517; Appx5518. But the district court's construction excludes any and all types of DRAM circuits and says nothing about interconnections. Appx1380-1381. Netlist's evidence therefore does not prove infringement under the correct construction. *See, e.g.*, *Liquid Dynamics Corp. v. Vaughan Co.*, 449 F.3d 1209, 1224 n.2 (Fed. Cir. 2006) (holding "expert opinion evidence" was "irrelevant because it was based on an impermissible claim construction").

Second, Dr. Brogioli never offered an opinion on the plain meaning of "DRAM circuit" or "different from a DRAM circuit," despite repeated direct questioning. *See, e.g.*, Appx5550; Appx5623-5624; Appx5554-5560 (551:17-557:9); Appx5570-5573 (567:13-570:3).[7] To the contrary, he admitted that, outside of the context of the prosecution history, "DRAM circuit" carries a broad

[7] Instead, Dr. Brogioli testified that Netlist had disclaimed only Rajan137's particular circuits. *See* Appx5624-5625.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 43

meaning in the field, not limited to a particular physical stacking arrangement or type.

For example, Dr. Brogioli admitted that he personally drafted a book chapter, two papers, and a Ph.D. thesis in which he discussed DRAM without limiting that concept to a particular type of wiring. Appx5562-5567 (559:19-564:6). He also agreed that Hyun Lee (named inventor of the '060 and '160 patents) testified that DRAM circuits and DRAM dies are both widely referred to as "DRAM," refuting any suggestion that "DRAM circuits" are somehow distinct from "DRAM dies." Appx5552-5553; Appx5861. Dr. Brogioli further conceded that Netlist, in its 10-K disclosure to the U.S. Securities and Exchange Commission, used the concepts of DRAM circuits and DRAM "interchangeably." Appx5561-5562. Critically, Netlist's expert expressly refused to offer an opinion as to the plain meaning of "DRAM circuits" as claimed, feigning complete ignorance of the meaning of the term "circuit." *See, e.g.,* Appx5550.

This evidence demonstrates that a skilled artisan would understand "DRAM circuit" as a broad term referring to the circuits that make up DRAM memory, not limited to only an embodiment from Rajan137 where the array dies are connected using wire bonding. Reinforcing that evidence, Samsung's expert, Dr. Robins, offered an unrebutted opinion that the plain meaning of "DRAM circuit" to a skilled artisan encompasses the individual DRAM cells on Samsung's DRAM dies,

Samsung Electronics Co., Ltd.
Ex. 1077, p. 44

as well as the "entire array" of those cells on each die without reference as to how those dies were connected. Appx6064-6065; Appx6071; Appx6073; Appx6075; Appx6101.

Dr. Brogioli admitted that each DRAM die has millions of DRAM cells on it, Appx5554; Appx5555, and that each DRAM cell is made of circuit components—a transistor and a capacitor—thus meeting the "circuit" aspect of "DRAM circuit." Appx5570-5573 (567:20-570:8). Because the plain meaning of "DRAM circuits" includes DRAM cells, Dr. Brogioli's admissions constitute a concession that Samsung's accused products are DRAM circuits and do not infringe. As noted above, Netlist's lead counsel candidly admitted that, to prove infringement, Netlist "actually had to show that what was in the HBM products was not a DRAM circuit." *See* Appx22053-22054 (53:12-54:2) (explaining the district court's disclaimer in arguments before the PTAB). Netlist could not show that "what was in the HBM products was not a DRAM circuit," *id.*, and Samsung does not infringe the '060 and '160 patents.

Therefore, substantial evidence does not support the jury's infringement verdict, under the district court's construction, and this Court should enter judgment of noninfringement.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 45

B. The ’054 and ’918 Patents

1. Substantial Evidence Does Not Support the Jury’s Finding That Samsung’s Accused Products Satisfy the ’054 Patent Claims’ “Second Operable State” Requirement

Netlist stretched the claim language “second ***operable*** state” and its construction (“state in which ***the memory module is operated*** after transition,” Appx1241) well beyond plain meaning. In particular, Netlist alleged that Samsung’s products are in the claimed ***operable*** state after transition, even though Samsung’s DRAMs are entirely ***inoperable***, the stored data ***is deleted***, and everything on the module is ***powered off*** except low-level circuitry awaiting a “wake up” signal. Appx5985 (965:3-5). Netlist’s interpretation is foreign to the ’054 patent and contradicts its main thrust—avoiding data loss by transferring data from volatile DRAM to permanent flash if operating power is declining. *See, e.g.*, Appx648 (title); Appx654-655 (Figs. 4A-4B); Appx676 (3:66-4:27) (summarizing disclosure as relating to modules having both non-volatile and volatile memory subsystems). This “second operable state” is a functioning operating state where this data transfer takes place to avoid data loss. *See* Appx686-687 (24:60-25:7).

The construction of “second operable state” is a “state in which ***the memory module is operated*** after transition.” The key to this limitation is actual operation of the memory module. The patent consistently discloses “second operable states” in which the DRAMs transmit, receive, and maintain data. Appx676 (3:59-62,

Samsung Electronics Co., Ltd.
Ex. 1077, p. 46

3:39-40), Appx686-687 (24:60-25:7). And, the critical "second operable state" transfers data from volatile DRAM to non-volatile flash memory to prevent data loss in response to a voltage change. Appx687 (25:8-31); Appx686-687 (24:60-25:7); Appx687-688 (26:66-27:1). Each disclosed "second operable state" example includes actual memory module operations, as the claim language requires. Appx687 (25:58-66); *see also* Appx686-689 (24:35-29:64). The patent never suggests that a memory module may be considered "operable" when its DRAMs are turned off, thus wiping any stored data, and rendering it incapable of transmitting, receiving, or maintaining data.

Netlist's infringement theory stretches the claim language to encompass functionality "far beyond anything described in the specification." *Kinetic Concepts, Inc. v. Blue Sky Med. Grp., Inc.*, 554 F.3d 1010, 1019 (Fed. Cir. 2009). Unlike the patent, Netlist expert Dr. Mangione-Smith theorized that the "state in which ***the memory module is operated*** after transition" includes the "off" state, in which (as both experts agree) the accused memory modules' DRAMs are inoperable and do not read or write data, and stored data is lost. Appx5381-5383 (378:16-380:6) (conceding "after an overvoltage condition . . . the DRAMs are turned off"); Appx5434-5435 (agreeing data "gets wiped out"); Appx5439-5440; Appx5914-5917 (894:5-897:7); Appx5987; Appx5988; Appx5990-5991.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 47

Netlist's entire theory rests on Dr. Mangione-Smith's claim that, when the module is off, certain non-DRAM, low-level components receive power merely to listen for a "wake up" signal. *See* Appx5381-5382; Appx5341; Appx5346; Appx5333. These listening components have nothing to do with memory, transferring data, or other memory module operations. Appx5435; Appx5439-5440; Appx5425. Netlist never explained—apart from Dr. Mangione-Smith's *ipse dixit*—its assumption that a skilled artisan would understand a "memory module" to be "operated" whenever ***any*** component receives power even if no memory is operable, no transfers may take place, and no data moves. Appx5381-5383 (378:25-380:6). With that foundational assumption unsubstantiated, Dr. Mangione-Smith's theory lacks substantial support. *Alexsam, Inc. v. IDT Corp.*, 715 F.3d 1336, 1342 (Fed. Cir. 2013) (reversing infringement JMOL denial for "cursory" expert testimony); *Becton, Dickinson & Co. v. Tyco Healthcare Grp., LP*, 616 F.3d 1249, 1260 (Fed. Cir. 2010) ("[I]nferences wholly unsupported by the record cannot" support a verdict.).

Netlist and Dr. Mangione-Smith eviscerated the requirement for the "memory module" to be "operated" in a second state. No rational factfinder could have concluded that a state in which the memory module is inoperable is somehow operable. This Court should reverse the district court's denial of JMOL of noninfringement. *See, e.g., Commil USA, LLC v. Cisco Sys., Inc.*, 813 F.3d 994,

Samsung Electronics Co., Ltd.
Ex. 1077, p. 48

997 (Fed. Cir. 2015) (reversing JMOL denial where verdict was based on conclusory expert testimony and evidence did not correspond to claims); *Ericsson, Inc. v. D-Link Sys., Inc.*, 773 F.3d 1201, 1223-24 (Fed. Cir. 2014).

2. Substantial Evidence Does Not Support the Jury's Finding That Samsung's Accused Products Satisfy the '918 Patent's "Converter Circuit" Limitation

a. Netlist's Literal Infringement Theory Applied a Purely Functional Definition of "Converter Circuit"

Because the Court did not construe "converter circuit," Netlist bore the burden of proving that the accused LDO regulators are "converter circuits" under its plain meaning. *See Broadcom*, 543 F.3d at 696; *ActiveVideo*, 694 F.3d at 1326. Yet Netlist and Dr. Mangione-Smith provided only a purely functional theory that "converter circuit" covers anything that alters a voltage—even a simple wire—and imposes "no structural limitation." Appx5419-5420; Appx6258. The district court acknowledged that Netlist applied "electrical components for voltage conversion," Appx19—a purely functional formulation.

Netlist's theory violates the longstanding prohibition against "broaden[ing] . . . product claims by describing the product in terms of function." *See Gen. Elec. Co. v. Wabash Appliance Corp.*, 304 U.S. 364, 371 & n.8 (1938); *Holland Furniture Co. v. Perkins Glue Co.*, 277 U.S. 245, 257-58 (1928). The narrow exception allowing functional claim language ***in exchange*** for limiting a patent's reach to structures equivalent to those in the specification is inapplicable here. *See*

Samsung Electronics Co., Ltd.
Ex. 1077, p. 49

35 U.S.C. §112, ¶6;[8] *Warner-Jenkinson Co. v. Hilton Davis Chem. Co.*, 520 U.S. 17, 28 (1997). Both sides agreed that this term is not a means-plus-function term. Appx1238-1241 (showing no means-plus-function contentions in joint claim construction chart). Thus, Netlist cannot treat "converter circuit" as purely functional without accepting the strictures of §112, ¶6—yet that is precisely what the district court permitted.

Specifically, Dr. Mangione-Smith opined that the accused products' LDO regulator is the recited "converter circuit" simply because it performs the function of "converting" a voltage level—impermissibly transforming the structural "converter circuit" limitation into a purely functional one. *See, e.g.*, Appx5339; Appx5646-5650 (640:24-644:23); Appx5678; Appx5674. Without analysis, he used this purely functional "converter circuit" definition to assert that the term encompasses LDO regulators. Appx5339. Such conclusory expert testimony is not substantial evidence. *See MobileMedia Ideas LLC v. Apple Inc.*, 780 F.3d 1159, 1172 (Fed. Cir. 2015).

The district court erroneously shifted the burden to Samsung, reasoning that requiring Netlist to apply the plain meaning of "converter circuit" amounts to "late claim construction." Appx20. Yet, applying a claim's plain meaning is the correct

[8] Because the respective applications "were filed before the effective date of the AIA," pre-AIA §112, ¶6 applies. *Cox Commc'ns, Inc. v. Sprint Commc'n Co. LP*, 838 F.3d 1224, 1228 (Fed. Cir. 2016); Appx600-601; Appx648-649.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 50

analysis for infringement. *E.g.*, *Broadcom*, 543 F.3d at 696. Netlist bore the burden of proving that the accused products satisfy the plain meaning of "converter circuit," but instead morphed the limitation into a functional, structure-less requirement. Samsung's argument did not request claim construction at all but simply asked that the district court apply "plain meaning" law faithfully.

Finally, although the district court purported to identify structure in Netlist's application of "converter circuit"—holding that it requires "electrical components for voltage conversion," Appx19—the phrase "electrical components" connotes no specific structure for converting a voltage, which it must do to be correct here. *See MTD Prods. Inc. v. Iancu*, 933 F. 3d 1336 (Fed. Cir. 2019). "Electrical components for voltage conversion" is a meaningless phrase because any electrical component will change voltage, even a bare wire, as Dr. Mangione-Smith recognized. Appx5419-5420. Thus, the district court's post-hoc construction is far broader than "converter circuits" and the court erred by permitting this amorphous understanding of the claim.

b. The Record Lacks Substantial Evidence To Support Netlist's Assertion of Literal Infringement Under the Plain Meaning of "Converter Circuit"

Only Samsung's expert, Mr. McAlexander, offered an evidence-based interpretation of what "converter circuit" means to skilled artisans. He explained that a "converter circuit" must combine a switch and an inductor and operates by

Samsung Electronics Co., Ltd.
Ex. 1077, p. 51

turning the switch "on and off to provide charge to an inductor." *See* Appx5909-5912 (889:23-892:19).

Netlist pointed to only the LDO regulators in the accused devices as satisfying the "converter circuit" limitation. Mr. McAlexander contrasted converter circuits from LDO regulators, which are distinct circuits that dissipate energy thermally, like a resistor, without induction. Appx5909-5912. Mr. McAlexander cited, *inter alia*, the testimony of two Samsung engineers. *See id.* (A converter is "comprised of an inductor and switch," and a regulator is "something without induction."); Appx5678; Appx5662. This evidence-based interpretation of "converter circuit" stood unrebutted. *CommScope Techs. LLC v. Dali Wireless Inc.*, 10 F.4th 1289, 1297 (Fed. Cir. 2021) ("Not only was there a lack of evidence to show that the accused product met the proper construction of the claims, there is unrebutted evidence showing the opposite.").

Applying this plain meaning, Mr. McAlexander explained that the accused LDO regulator is not the recited "converter circuit." Appx5905-5912 (885:23-892:19). He explained that, unlike a converter circuit, LDO regulators "throw[] away the extra energy" by "expending that energy through a resistor—a set of resistors which creates [] thermal radiation." Appx5909. Even Dr. Mangione-Smith admitted that an LDO regulator has "fundamental" differences from the

converter circuits that the '918 patent describes—i.e., buck, boost, and buck-boost converters. Appx641; Appx5418; Appx5424.

Struggling to find substantial evidence support under the proper interpretation of "converter circuit," the district court cited testimony that Samsung's accused products include switches. Appx20. Yet, no witness testified that a "converter circuit" is a mere switch alone. Rather, a switch must turn "on and off to provide charge to an inductor." *See* Appx5909-5912.

The record supports only one conclusion—that the accused products do not include the recited "converter circuit" under the term's plain meaning and, therefore, do not literally infringe.

c. Substantial Evidence Does Not Support Netlist's Alternative, Conclusory DOE Argument

Substantial evidence does not support a finding of infringement under the doctrine of equivalents. Netlist's DOE theory consists of a single, conclusory assertion by Dr. Mangione-Smith: "[The LDO is] equivalent because it does the same function which is to convert; it does it the same way, which is to reduce the input voltage; and it achieves the exact same result, which is a regulated output voltage." Appx5340.

Such conclusory testimony does not provide substantial evidence of infringement under DOE. *See, e.g., Hewlett-Packard Co. v. Mustek Sys., Inc.*, 340 F.3d 1314, 1322-23 (Fed. Cir. 2003) (holding insufficient expert testifying that the

Samsung Electronics Co., Ltd.
Ex. 1077, p. 53

accused product performed same function, way, result, without further explanation). Dr. Mangione-Smith's testimony is especially insufficient given the uncontroverted evidence that "converter circuits" operate in a fundamentally different "way" than LDO regulators (i.e., induction and switching versus no induction and heat dissipation). Appx5418; Appx5424; Appx5909-5912. Indeed, Dr. Mangione-Smith admitted that LDO regulators and buck converters are not interchangeable. Appx5424; *Warner-Jenkinson*, 520 U.S. at 36 (known interchangeability is factor relevant to DOE).

Netlist's post-trial argument (Appx18115-18116) that the "way" is simply to convert the voltage ignores the real-world differences between converter circuits and LDO regulators and requires accepting Netlist's erroneous wholly-functional view of "converter circuits." *See VirnetX, Inc. v. Cisco Sys., Inc.*, 767 F.3d 1308, 1323 (Fed. Cir. 2014) (reversing JMOL denial where expert failed to adequately opine on the "way" prong). It also incorrectly applies the function-way-result test: By treating converting a voltage as both the "function" and the "way," Netlist improperly conflates those prongs without legal support. *Applied Med. Res. Corp. v. U.S. Surgical Corp.*, 312 F. App'x 326, 332 (Fed. Cir. 2009) (prohibiting conflation of "the 'way' and 'function' prongs of the function-way-result test").

Infringement under DOE is a separate analysis from literal infringement and must be proven with "particularized testimony and linking argument," and Netlist

Samsung Electronics Co., Ltd.
Ex. 1077, p. 54

provided none here. *Tex. Instruments Inc. v. Cypress Semiconductor Corp.*, 90 F.3d 1558, 1567 (Fed. Cir. 1996). Even assuming, *arguendo*, that the evidence that Dr. Mangione-Smith used for literal infringement is part of Netlist's DOE theory, such showing was insufficient and never linked to any DOE test. *See* Appx20; Appx18114-18115. Critically, Netlist never linked any DOE argument to the "converter circuit" limitation or its plain meaning as Dr. Mangione-Smith's analysis assumed a purely functional definition of "converter circuit." *See* Section I.B.2.a, *supra*.

C. The '339 Patent: Substantial Evidence Does Not Support the Jury's Finding That Samsung's Accused Products Satisfy the Multiple "Byte-Wise Data Paths" Limitation

The claims recite a "byte-wise buffer" having a "first side" coupled to the computer and a "second side" coupled to multiple "ranks" of DDR DRAM devices. Appx595 (19:40-61). When data is written to the memory module, the byte-wise buffer is required to "drive a respective byte-wise section" of the write data from the first to the second side of the buffer. *Id.*; Appx1356. Because the "drive" term means "enabl[e] only one of the data paths while the other possible paths are disabled," Appx1359, the district court's construction requires multiple byte-wise write data paths.

The accused Samsung products do not have the required multiple byte-wise data paths. Instead, both experts agree that they have ***only one eight-bit data path***.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 55

Appx5395; Appx5451-5452; Appx5454; Appx5944. Under the plain meaning of this limitation, substantial evidence does not support the jury's infringement verdict.

1. Netlist's Faulty Interpretation of "Byte-Wise" Does Not Provide a Basis To Sustain the Infringement Verdict

In an attempt to avoid the fact that the accused devices have a single pathway, Netlist recasts the single, byte-wise eight-bit path in the accused devices as two four-bit pathways. Yet, Dr. Mangione-Smith agreed that the claims recite two byte-wise data paths—not a single byte-wise data path that can be thought of as divided into smaller-size data paths (e.g., a four-bit "nibble" path). *See* Appx5476-5477 (admitting the claim "does not recite actively driving a nibble-wise section of the write data along a nibble-wise data path"); Appx5452 (admitting the claims require at least "[t]wo write paths" from the first side of the buffer to the second side of the buffer). Dr. Mangione-Smith also admitted that he was resorting to a mathematical sleight of hand—the structures he identified are two four-bit data paths, ***not*** two eight-bit data paths. Appx5458.

To explain the obvious deviation from the claimed "byte-wise" requirement, Dr. Mangione-Smith made a naked assumption: "byte-wise" simply means that "the width is measured in terms of bytes." Appx5456-5457. He reasoned that a four-bit nibble counts as "byte-wise" because it can be measured in terms of bytes (i.e., half a byte). Appx5395.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 56

Dr. Mangione-Smith's assumption is conclusory, as he cited nothing to support that a skilled artisan would understand "byte-wise" as simply referring to how a data path is "measured," or even what it means to "measure" a data path. *See, e.g.*, *MobileMedia*, 780 F.3d at 1172; *Yoon Ja Kim v. ConAgra Foods, Inc.*, 465 F.3d 1312, 1320 (Fed. Cir. 2006). Furthermore, his theory eviscerates the requirement altogether. Under Netlist's theory, ***any*** size data path would qualify as a "byte-wise data path" because any number of bits can be represented as some fraction or multiple of eight bits—e.g., one-eighth of a byte, one quarter of a byte, half of a byte, etc. Such an interpretation impermissibly strips the claim limitation of all meaning. *Intel Corp. v. Qualcomm Inc.*, 21 F.4th 801, 810 (Fed. Cir. 2021) ("It is highly disfavored to construe terms in a way that renders them void, meaningless, or superfluous." (citation and quotation marks omitted)).

Perhaps most problematic for Netlist's relabeling exercise is that skilled artisans have a term for a four-bit pathway: "nibble." The record makes plain that four bits are a nibble, and two nibbles make a byte. Appx5395; Appx5446; Appx595 (20:6-14) (reciting "a ***first nibble*** of the respective ***byte-wise*** section of the N-bit wide write data" and "a ***second nibble*** of the respective ***byte-wise*** section of the N-bit wide write data"); Appx594 (describing "nibble-wide data path"); Appx583. The repeated use of "nibble" in the claims of the '339 patent demonstrates that the inventors well-understood the difference between an eight bit

Samsung Electronics Co., Ltd.
Ex. 1077, p. 57

"byte-wise" data path and a four bit "nibble" data path. Appx595-599. Consequently, Dr. Mangione-Smith's opinion that two nibble paths satisfy the requirement for multiple byte-wise data paths in claim 1 cannot be correct, as ***each*** byte-wise data path must be at least two nibbles.

2. The District Court Erroneously Eviscerated the "Byte-Wise" Limitation

The district court likewise erred in using a supposed distinction between "byte-wise" and "byte-wide" to eviscerate the "byte-wise" limitation. *See* Appx8. Contrary to the district court's assertion, the claims do not "recite both 'byte-wise' and 'byte-wide.'" Appx6. The claims nowhere mention "byte-wide," and the phrases "byte-wide" and "byte wide" are each used once in the specification, which never suggests any difference in meaning between "byte-wise," "byte-wide," and "byte wide." Appx592; Appx595.

Although claim 1's preamble recites that "each set of data signal lines is ***a byte wide***" (with no hyphen), that language does not imply that "byte-wise" has any different meaning. Claim 1 uses both phrases for grammatical, not substantive, reasons. "Byte-wise" appears much later in claim 1 as a compound adjective modifying the nouns "buffer(s)," "data path," and "section," Appx595 (19:40-67); in contrast, the preamble does not use "byte wide" as a compound adjective, such that one could draw a meaningful substantive distinction between the two phrases, Appx595 (19:14-15). Further, the preamble's recitation of the

Samsung Electronics Co., Ltd.
Ex. 1077, p. 58

language, "each set of data signal lines is ***a byte-wise***," would result in a nonsensical sentence fragment, and thus the patentee chose "a byte wide." Claims may use multiple terms to refer to the same concept, *Baran v. Med. Device Techs., Inc.*, 616 F.3d 1309, 1316 (Fed. Cir. 2010), and here the specification and claims of the '339 patent support only the conclusion that "a byte wide" and "byte-wise" both simply refer to eight bits.

Nor does expert testimony support substantively distinguishing between "byte-wise" and "byte-wide." Dr. Mangione-Smith testified merely that he is "more familiar" with the non-recited "byte-wide" term than the recited "byte-wise" claim term on which he was supposed to opine, not that "byte-wide" somehow has a different meaning than "byte-wise." Appx5457. On the other hand, Mr. McAlexander was clear that "byte-wise" is well-known to be eight bits. Appx5940.

The district court's understanding of "byte-wise" is also belied by the specification, which uses "byte-wise" to describe eight-bit-wide data paths. *See, e.g.*, Appx592 (13:43-53, 13:57-14:8, 14:15-18). The specification also distinguishes between four-bit data paths and eight-bit data paths, and never describes a four-bit data path as a "byte-wise" data path. Appx594 (17:30-32); Appx595 (20:6-8, 20:13-14).

Samsung Electronics Co., Ltd.
Ex. 1077, p. 59

Therefore, Netlist's faulty interpretation of "byte-wise" does not provide a basis to sustain the verdict. Because there is no evidence apart from Dr. Mangione-Smith's *ipse dixit* and Netlist's attorney argument that the accused products have multiple byte-wise data paths, this Court should enter judgment of noninfringement.

II. Substantial Evidence Does Not Support the Damages Verdict

Netlist took the surprising position below that it was entitled to ***100%*** of the ***revenues*** associated with the patented inventions. Netlist ignored the basic *Georgia-Pacific* framework that sets a reasonable royalty as one a rational licensee would agree to, assigning ***0%*** of revenue and ***100%*** of cost to Samsung as licensee, leaving no incentive to implement the accused features at all. What rational licensee would agree to incur the costs and risk to implement a patented feature if the associated royalty payment guaranteed a substantial loss?

Netlist's damages theory was entirely result-oriented, and Netlist neglected its duty to present evidence that the damages it sought reflect the value of the patented technology and nothing more. Because Netlist pursued an invalid damages theory and failed to present an alternative case for damages, the Court should instruct the district court to enter JMOL of zero damages. Alternatively, the Court should remand for the district court to determine whether Netlist forfeited alternative damages theories and whether to order a new damages trial.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 60

A. Mr. Kennedy Did Not Calculate Samsung's Incremental Profits and Gave Netlist 100% of the Alleged Incremental Revenue

1. Mr. Kennedy Calculated Revenues, Not Profits

"Any reasonable royalty must seek to measure the value of the patented technology." *VLSI Tech. LLC v. Intel Corp.*, 87 F.4th 1332, 1345 (Fed. Cir. 2023) (internal quotation marks omitted) (collecting cases). "[A]nticipated incremental profits under the hypothesized conditions are conceptually central to constraining the royalty negotiation." *Aqua Shield v. Inter Pool Cover Team*, 774 F.3d 766, 772 (Fed. Cir. 2014) (citation omitted).

An accused infringer's incremental profits from a patented technology are its incremental revenues (the additional revenues earned by using the technology) less its incremental costs (the additional costs incurred to implement the technology). *See VLSI*, 87 F.4th at 1347 (noting damages expert assessed "incremental costs" and "incremental revenues," "[r]eflecting the role of profits (not revenues per se) in the inquiry"). Indeed, "[a]ccording to basic accounting principles," profit equals revenues minus expenses. *U.S. Steel Grp. v. United States*, 225 F.3d 1284, 1291 (Fed. Cir. 2000).

Here, Mr. Kennedy admitted he calculated only "the ***revenue*** difference between [his] noninfringing alternative and the [accused] product"—and did not consider costs. Appx5781; *see also* Appx5704-5708. Thus, despite asserting he determined what Samsung "made off" the patents, Appx5691, at best Mr. Kennedy

Samsung Electronics Co., Ltd.
Ex. 1077, p. 61

determined Samsung's incremental revenues, rendering his reasonable-royalty analysis incomplete, unreliable, and impermissible.

Mr. Kennedy's omission of Samsung's incremental costs is inexcusable. He could have analyzed Samsung's costs to implement the alleged patented features and offset them from his calculation. For example, he could have calculated the costs of implementing accused DDR5 products' power-management integrated circuit (PMIC), which he cited as the technical benefit of the '054 and '918 patents. *See* Appx5774 (Kennedy testifying that the noninfringing alternative required "pull[ing] the PMIC off the board"); Appx5402-5403 (Netlist's technical expert testifying that a noninfringing alternative to the '054 and '918 patents "move[s] the power management off the module"). Likewise, Mr. Kennedy could have calculated Samsung's costs to implement the feature of the accused HBM products that Mr. Kennedy credited as the technical benefit of the '060 and '160 patents ("eight high or higher DIMMs"). *See* Appx5707.

Mr. Kennedy's failure to account for Samsung's incremental costs renders his 100%-of-revenue theory legally insufficient to support the damages verdict. *See ResQNet.com, Inc. v. Lansa, Inc.*, 594 F.3d 860, 872 (Fed. Cir. 2010) (explaining patentee's "burden . . . to persuade the court with legally sufficient evidence regarding an appropriate reasonable royalty" (citation omitted));

Samsung Electronics Co., Ltd.
Ex. 1077, p. 62

Transclean Corp. v. Bridgewood Servs., Inc., 290 F.3d 1364, 1376 (Fed. Cir. 2002).

That the jury awarded less than Mr. Kennedy's demand is immaterial. Netlist claims the jury "awarded exactly 75% of [Mr. Kennedy's] proposed sum for each accused product family," Appx18147, suggesting that the jury used Mr. Kennedy's proposal as a starting point and adjusted it. But "[b]eginning from a fundamentally flawed premise and adjusting it . . . nevertheless results in a fundamentally flawed conclusion." *Uniloc USA, Inc. v. Microsoft Corp.*, 632 F.3d 1292, 1317 (Fed. Cir. 2011). Thus, the jury's using Mr. Kennedy's inflated proposal "merely as a starting point," rather than the end point, does not erase the error. *Id.*

Because Mr. Kennedy relied solely on Samsung's alleged incremental revenues and failed to consider its incremental costs, his opinions are fatally flawed. The "core economic question is what the infringer in the hypothetical negotiation would have anticipated the ***profit***-making potential of use of the patented technology to be," not its ***revenue***-making potential. *VLSI*, 87 F.4th at 1346 (internal quotation marks and citation omitted). Without evidence of that "profit-making potential," substantial evidence does not support Netlist's proffered royalty. Netlist put all of its eggs into this 100%-of-revenue basket, and offered

Samsung Electronics Co., Ltd.
Ex. 1077, p. 63

the jury no alternative theory or evidence to carry its burden on damages. The Court should therefore reverse the denial of JMOL. *See* Section II.C, *infra*.

2. Mr. Kennedy Applied a 100%-of-Revenue Rule of Thumb

Compounding his failure to calculate profit at all, Mr. Kennedy opined that Samsung, in a hypothetical negotiation, would have agreed to give Netlist ***100%*** of the ***revenue*** allegedly corresponding to the patented technology. Appx5799. He did not explain why Samsung, as a rational company, would have agreed to implement the patented technology ***at a guaranteed loss*** by surrendering all incremental revenues while bearing additional costs to implement the technology.

This 100%-of-revenues theory violates precedent. The notion that a rational company "would agree to pay a royalty in excess of what it expected to make in profit" is "absurd." *Lindemann Maschinenfabrik GmbH v. Am. Hoist & Derrick Co.*, 895 F.2d 1403, 1408 (Fed. Cir. 1990) (citation omitted). "[A] reasonable royalty must be fixed so as to leave the infringer a reasonable profit." *Hughes Tool Co. v. Dresser Indus., Inc.*, 816 F.2d 1549, 1558 (Fed. Cir. 1987) (citing *Georgia-Pacific Corp. v. U.S. Plywood Corp.*, 318 F. Supp. 1116, 1122 (S.D.N.Y. 1970)); *see also Georgia-Pacific Corp. v. U.S. Plywood-Champion Papers, Inc.*, 446 F.2d 295, 299 (2d Cir. 1971) (reversing royalty that "gobble[d] up all of GP's expected profit" and "did not allow GP a reasonable profit"). For this reason alone, Mr. Kennedy's proposed royalty is untenable.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 64

Unsurprisingly, courts have rejected similar theories as untethered to economic reality. *See Contour IP Holding, LLC v. GoPro, Inc.*, 2020 WL 5106845, at *14 (N.D. Cal. Aug. 31, 2020) ("[I]t would be unreasonable and unreliable for [the patentee's damages expert] to conclude that 100% of profits associated with the infringing technology would go to Contour."); *Looksmart Grp., Inc. v. Microsoft Corp.*, 2019 WL 4009263, at *3 (N.D. Cal. Aug. 5, 2019); *Labyrinth Optical Techs. LLC v. Alcatel-Lucent USA, Inc.*, 2015 WL 12696081, at *4 (C.D. Cal. Sept. 2, 2015); *Nordock Inc. v. Sys. Inc.*, 2013 WL 989864, at *8 (E.D. Wis. Mar. 13, 2013).

Mr. Kennedy's 100%-of-revenues theory is indistinguishable from previous "rules of thumb" this Court has rejected. *Uniloc* held "the 25 percent rule of thumb is a ***fundamentally flawed*** tool for determining a baseline royalty rate in a hypothetical negotiation" because it "fails to tie a reasonable royalty base to the facts of the case." 632 F.3d at 1315. Likewise, *VirnetX* held that assuming a 50% profit split as a starting point was improper. 767 F.3d at 1333-34. Mr. Kennedy's 100%-of-revenue rule of thumb is even less supportable than *Uniloc*'s and *VirnetX*'s because at least those cases split profit, not revenue.

Neither Mr. Kennedy nor Netlist presented evidence that anyone has ever agreed to pay 100% of the revenue associated with patented technology for any license, much less under similar circumstances. Mr. Kennedy's opinion is thus

Samsung Electronics Co., Ltd.
Ex. 1077, p. 65

"untethered from the patented technology at issue and the many licenses thereto and, as such, [is] arbitrary and speculative." *LaserDynamics*, 694 F.3d at 81.

Mr. Kennedy's damages theory exceeds a remedy Congress long ago eliminated: disgorgement of profits. *See Aro Mfg. Co. v. Convertible Top Replacement Co.*, 377 U.S. 476, 505 (1964) (holding § 284 "eliminate[d] the recovery of [the infringer's] profits . . . and allow[s] recovery of damages only"). By assigning 100% of revenue to Netlist, Mr. Kennedy disgorged an amount ***greater than*** Samsung's incremental profits. *See* Appx5767-5768. Because § 284 prohibits disgorgement of profits, it must also prohibit Mr. Kennedy's 100%-of-revenues theory.

B. The Record Lacks Evidence of Proper Apportionment

Netlist failed adequately to apportion the damages in this case. Netlist did not adduce evidence as to what portion of Samsung's alleged incremental revenues was attributable to the asserted patents, rather than to the contributions made by Samsung and other leading memory companies who developed the memory standards with which the accused products comply. *See, e.g.*, Appx2249-2252. Because well-settled law requires apportionment, Netlist's failure is fatal to its damages case.

A patentee "must in every case give evidence tending to separate or apportion the defendant's profits . . . between the patented feature and the

Samsung Electronics Co., Ltd.
Ex. 1077, p. 66

unpatented features." *Garretson v. Clark*, 111 U.S. 120, 121 (1884); *VLSI*, 87 F.4th at 1345 ("Any reasonable royalty must seek to measure the value of the patented technology—it must be 'apportioned' to that value—by separating out and excluding other value in economic products or practices." (collecting cases)).

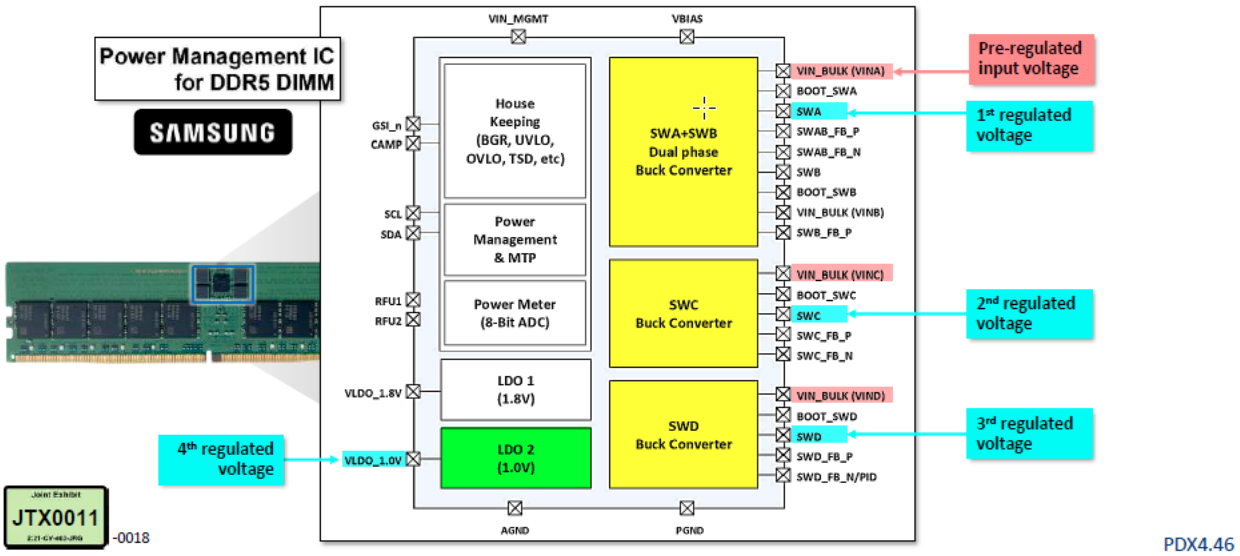
Here, Netlist failed to adduce or apply the evidence *Garretson* requires. For each accused product, Mr. Kennedy relied on technical experts to identify capabilities that the patents supposedly enabled and then tried to estimate the price of those capabilities. *See, e.g.*, Appx5706-5707. In each instance, Netlist failed to present evidence as to what portion of those capabilities was attributable to patented features. Indeed, even if Netlist intended to pursue all revenue (which it ultimately did), the law requires that Netlist prove that 100% of the incremental revenue was attributed to the patented features. Netlist did none of this, and so failed the apportionment test, providing an independent reason why the Court should reverse the damages award. *See* Section II.C, *infra*.

1. The Alleged Technical Benefit of the '054 and '918 Patents Includes Unpatented and/or Prior-Art Features

For the '054 and '918 patents, Mr. Kennedy cited Dr. Mangione-Smith's testimony that the patents enabled the accused DDR5 products' improved PMIC, where the PMIC is on a memory module (instead of off module) and which Dr. Mangione-Smith testified results in 30% higher power efficiency. Appx5699-5700; Appx5332-5333. Even assuming the unlikely premise that the 30%

Samsung Electronics Co., Ltd.
Ex. 1077, p. 67

efficiency gain was attributable to the PMIC,[9] Dr. Mangione-Smith's own demonstrative shows that the PMIC includes ***unpatented features***, like a "Power Management" component, "Power Meter," and "LDO 1":



Appx22226 (citing Appx8068) (alterations in original); *see also* Appx5337-5338. Netlist failed to present any evidence apportioning the purported 30% efficiency gain "between the patented feature and the unpatented features." *Garretson*, 111 U.S. at 121.

Dr. Mangione-Smith admitted, moreover, that prior art discloses the concept of "putting the power management block . . . on the memory module," but he did not consider what was "new or unique" about the '054 and '918 patents. Appx5413-5414; Appx5411-5412; *contra Omega Pats., LLC v. CalAmp Corp.*, 13

[9] Netlist derived this number from marketing material comparing DDR5- and DDR4-standard compliant products, thus capturing many advances from the newer DDR5 standard. Appx22203 (citing Appx12653).

F.4th 1361, 1377-78 (Fed. Cir. 2021) (holding "the incremental value that [the] patented improvement added" must be "apportioned from the value of any conventional features"). Accordingly, substantial evidence does not support the alleged incremental revenues for the '054 and '918 patents.

2. The Alleged Technical Benefit of the '339 Patent Is Untethered to the Patent and Was in the Prior Art

For the '339 patent, Mr. Kennedy based his analysis on the capability of the accused DDR4 products to use "two DIMMs per channel" instead of one, but Netlist presented no evidence tying that capability to any claimed advance of the '339 patent, much less ***solely*** to that claimed advance. Appx5706-5707. Netlist instead tried to shift the burden onto Samsung to prove the two-DIMMs-per-channel capability was ***not*** due to the '339 patent. Appx18149 ("Samsung did not identify 'any alternative technology that it could have used in place of the '339 inventions.'" (quoting Appx5403)). But Netlist did not attempt to carry its burden to prove the capability was due solely to the '339 patent. *See ResQNet.com*, 594 F.3d at 872. Without that showing, Netlist was required to apportion. *Garretson*, 111 U.S. at 121.

The district court's order denying JMOL confirms Netlist's failure of proof on this point. The only cited evidence allegedly showing that the '339 patent enabled two DIMMs per channel was Dr. Mangione-Smith's testimony that "the capacity with the '339 invention is doubled." Appx50-51 (quoting Appx5403).

Samsung Electronics Co., Ltd.
Ex. 1077, p. 69

This testimony, however, is circular because Dr. Mangione-Smith based it on the assertion he was asked to prove—that the '339 patent enables two DIMMs per channel. Appx5403. That the district court had to cite the conclusion to support its premise confirms that the premise lacks substantial evidence.

Netlist suggested the two-DIMMS-per-channel capability is somehow tied to the '339 patent's distributed buffers, *see* Appx5388, but Dr. Brogioli admitted that such buffers were known in the prior art, Appx5459. Indeed, Samsung's technical expert testified and presented unrebutted evidence that distributed buffers were "well-known in the art by 2009." Appx5934-5935. Substantial evidence therefore does not support the alleged incremental revenues for the '339 patent.

3. The Alleged Technical Benefit of the '060 and '160 Patents Is Untethered to the Patents

Mr. Kennedy computed damages for the '060 and '160 patents assuming they "allow[] Samsung to sell eight[-]high or higher DIMMs," but the record does not support that assumption. Appx5707. Although Dr. Brogioli testified, *ipse dixit*, that the '060 and '160 patents "allow Samsung to go beyond a four[-]high stack," he did not supply any ***evidence*** for that conclusory assertion. Appx5540-5541. On the contrary, he admitted that he failed to consider prior art in assessing the technical benefits of the '060 and '160 patents. Appx5614-5615; *contra Omega*, 13 F.4th at 1377.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 70

Dr. Brogioli cited a figure in the patents depicting an eight-die-high stack that merely shows the inventor contemplated using his invention with eight-high stacks, but does not show that his technology enabled eight-high stacks. Appx5541; Appx22178. Dr. Brogioli also cited an article suggesting eight-high stacks were "probably" unattainable using a prior-art solution ("the multi-drop configuration"). Appx5541-5544 (538:8-541:3); Appx22179. He did not, however, present evidence showing this "multi-drop configuration" was the next-best available and acceptable noninfringing alternative, probably because the article discloses a better alternative—***one that achieves eight-high stacks***. Appx1583-1585. Again, "the incremental value that [the] patented improvement add[s]" must be "apportioned from the value of any conventional features," like eight-high stacks, which Mr. Kennedy did not do. *Omega*, 13 F.4th at 1377-78. Thus, substantial evidence does not support the alleged incremental revenues for the '060 and '160 patents.

4. Netlist Failed To Present Evidence That Mr. Kennedy's Noninfringing Alternatives Were Available and Acceptable

Another missing link between Samsung's alleged incremental revenues and the asserted patents is the lack of evidence that Mr. Kennedy's noninfringing alternatives were available and commercially acceptable by the time of the hypothetical negotiation.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 71

To value patented technology by comparing it to a noninfringing alternative, the party advancing the alternative must show it would have been acceptable and available. *See Grain Processing Corp. v. Am. Maize-Prods. Co.*, 185 F.3d 1341, 1351 (Fed. Cir. 1999) ("[A] rational would-be infringer is likely to offer an acceptable noninfringing alternative, if available, to compete with the patent owner rather than leave the market altogether."); *In re Ecast, Inc.*, 96 F. App'x 710, 711 n.* (Fed. Cir. 2004) (non-precedential) ("[U]nder Federal Circuit precedent evidence of the availability of an acceptable noninfringing alternative is relevant to . . . a reasonable royalty analysis." (citing *Zygo Corp. v. Wyko Corp.*, 79 F.3d 1563, 1571-72 (Fed. Cir. 1996))).

Here, Netlist failed to adduce any evidence that the noninfringing alternatives Mr. Kennedy used for his incremental-revenues calculation were available and commercially acceptable by the hypothetical negotiation. Regarding the '054 and '918 patents, for example, Mr. Kennedy's alleged noninfringing alternative was a theoretical DDR5 module with the PMIC removed. Appx5774; Appx5402-5403. Netlist did not, however, present evidence that this theoretical product was available or acceptable to customers. Indeed, Mr. Kennedy admitted having no such evidence. Appx5776-5777. Likewise, for the '339 patent, Dr. Mangione-Smith admitted customers would ***not*** have accepted Mr. Kennedy's alleged noninfringing alternative (a theoretical DDR4 module "not capable of

Samsung Electronics Co., Ltd.
Ex. 1077, p. 72

running at DDR4 speeds with two or more modules per channel"). Appx5403-5404; Appx5706-5707.

Because there is no evidence that Mr. Kennedy's alleged noninfringing alternatives were available and commercially acceptable, substantial evidence does not support the damages award. *See Grain Processing*, 185 F.3d at 1350-51; *Ecast*, 96 F. App'x at 711 n.*.

C. The Court Should Not Give Netlist a "Do-Over"

To remedy Netlist's unsupported damages theory, this Court should instruct the district court to enter JMOL of no damages. In *Promega Corp. v. Life Technologies Corp.*, the Court affirmed JMOL of no damages because Promega, like Netlist, "abandon[ed] valid theories of recovery in a singular pursuit of an ultimately invalid damages theory." 875 F.3d 651, 666 (Fed. Cir. 2017). That "deliberate strategy to adhere to a single damages theory had the effect of winnowing out from the case any argument about damages" under another theory. *Id.* at 662 (citation omitted). Netlist followed the same approach as Promega, and JMOL of no damages is appropriate.

Granting a new damages trial under these circumstances would incentivize patentees like Netlist to present unsupported damages theories, knowing a new damages trial is the only risk they face. Historically, damages-only trials heavily favor the patentee where the initial premise is, "the defendant has violated the law,

Samsung Electronics Co., Ltd.
Ex. 1077, p. 73

and we are here to determine what they owe." Giving a patentee, who fails to carry its burden of proof on damages by presenting an invalid damages theory, "a second bite at the apple . . . simply encourages overreaching on the first bite." *Network Prot. Scis., LLC v. Fortinet, Inc.*, 2013 WL 5402089, at *8 (N.D. Cal. Sept. 26, 2013). In such cases, "there is a positive need to deny a second bite in order to encourage candor in the first place." *Id.*; *see also Unwired Planet, LLC v. Apple Inc.*, 2017 WL 589195, at *1 (N.D. Cal. Feb. 14, 2017).

Consider how this Court treats infringement. Parties routinely set forth what they believe to be winning theories to prove infringement. This Court sometimes finds that a party has failed to carry its burden and renders judgment without remand. The result of overreach is not a "do-over." Instead, a party that fails to faithfully apply a claim construction or posits an incorrect legal theory is left empty-handed. Damages issues deserve the same treatment.

The statute entitles a successful patent plaintiff to not less than a reasonable royalty, but patentee still bears the burden of proving its damages through competent evidence, *Lucent Techs., Inc. v. Gateway, Inc.*, 580 F.3d 1301, 1324 (Fed. Cir. 2009), and "[t]he [patentee] must show his damages by evidence," *Philp v. Nock*, 84 U.S. 460, 462 (1873) (Damages "must not be left to conjecture by the jury. They must be proved, and not guessed at."). A patentee may, for example, "waive[] the right to damages based on alternate theories." *Id.* (citing *Promega*,

Samsung Electronics Co., Ltd.
Ex. 1077, p. 74

875 F.3d at 660). That is what happened here, as Netlist "failed to present an alternative case for damages" "in a singular pursuit of an ultimately invalid damages theory." *Promega*, 875 F.3d at 666. Bearing the burden of proving damages, the patentee is free to choose whatever theories and evidence suits its claims. Failing to follow the law or presenting unsubstantiated theories simply means the patentee failed to carry its burden and thus judgment in favor of the defendant follows. This Court should not grant Netlist "multiple chances to correct deficiencies in its arguments or the record." *Promega*, 875 F.3d at 666.

Accordingly, the Court should reverse the denial of JMOL of no damages. Alternatively, the Court should "remand to the district court to determine whether [Netlist] has waived the right to establish reasonable royalty damages under a new theory and whether to order a new trial on damages." *Finjan, Inc. v. Blue Coat Sys., Inc.*, 879 F.3d 1299, 1312 (Fed. Cir. 2018). Justice demands that a plaintiff, like Netlist, that chooses a "swing-for-the-fences" damages theory must be held to its burden of proof when that theory fails.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 75

III. Allowing Mr. Kennedy To Rely on a Non-Comparable License and a Flawed Regression Model Was an Abuse of Discretion

A. The District Court Should Have Excluded Mr. Kennedy's Comparable License Analysis

1. Netlist's Technical Expert Did Not Analyze 99.4% of the Licensed Rambus Patents

To inflame the jury and make his $404.2-million proposed royalty appear reasonable, Mr. Kennedy cited a $1.1-billion portfolio license between Samsung and third-party Rambus as evidence "that Samsung has paid significant amounts" before. Appx5697; Appx5747. Netlist did not even attempt to show comparability for 99.4% of the licensed patents.

When a patentee relies on a license between an accused infringer and a third party, the licensed patents must be "comparable to the patent in suit." *Georgia-Pacific*, 318 F. Supp. at 1120. "[A]lleging a loose or vague comparability between different technologies or licenses does not suffice." *LaserDynamics*, 694 F.3d at 79. Nor does it suffice to proffer a license but supply "no explanation to the jury about the subject matter or patents covered." *Lucent Techs.*, 580 F.3d at 1328. Otherwise, "a prevailing plaintiff would be free to inflate the reasonable royalty analysis with conveniently selected licenses without an economic or other link to the technology in question." *ResQNet.com*, 594 F.3d at 872.

Here, Netlist offered no evidence of comparability for 99.4% of the licensed Rambus patents. The Rambus license includes ***2,175*** patents and "none of the

Samsung Electronics Co., Ltd.
Ex. 1077, p. 76

patents-in-suit." Appx1490 n.11; Appx5739; Appx5747. Mr. Kennedy admitted that, although he asked Dr. Mangione-Smith to assess whether the licensed Rambus patents are comparable to Netlist's, Dr. Mangione-Smith merely "looked at a few of the patents." Appx5697; Appx5753. Indeed, Dr. Mangione-Smith analyzed only ***12 (<0.6%)*** of the Rambus patents but ultimately disavowed that Netlist's patents in this case "are technically comparable to the ***entire*** Rambus patent portfolio." Appx5462-5463; Appx5464. Yet the district court denied Samsung's *Daubert* motion on this point. Appx283; *see also* Appx1489-1490; Appx1686. That was an abuse of discretion requiring a new damages trial. That the district court reversed course and excluded the ***same*** Rambus license from a follow-on trial involving some of the same technology (e.g., DDR4 LRDIMM) between the parties confirms this point. *Netlist, Inc. v. Samsung Elecs. Co.*, No. 2:22-CV-00293-JRG, Dkt. 818 at 4 (E.D. Tex. Nov. 6, 2024).

This Court has rejected reliance on a license under similar circumstances. In *MLC Intellectual Property, LLC v. Micron Technology, Inc.*, the Court affirmed the exclusion of a license because the proffering expert "conducted no assessment of the licensed technology versus the accused technology" for 40 of the 41 licensed patents (97.6%). 10 F.4th 1358, 1374-75 (Fed. Cir. 2021). In *Trell v. Marlee Electronics Corp.*, the Court vacated because the district court relied on a license but "fail[ed] to consider the fact that [it] . . . encompassed the right to other

Samsung Electronics Co., Ltd.
Ex. 1077, p. 77

inventions." 912 F.2d 1443, 1447 (Fed. Cir. 1990). Here, too, Mr. Kennedy relied on a portfolio license but failed to consider that it encompasses 2,175 patents, and Dr. Mangione-Smith analyzed only 12 for comparability.

That Mr. Kennedy told the jury the Rambus license "should not be used to calculate the specific royalty" does not diminish its prejudicial effect. Appx5748. In *Uniloc*, this Court affirmed the grant of a new damages trial because the patentee introduced evidence of the accused infringer's total revenues, and that "$19 billion cat" could not be "put back into the bag." 632 F.3d at 1320. "The disclosure that a company has made $19 billion dollars in revenue from an infringing product cannot help but skew the damages horizon for the jury." *Id.* Indeed, such disclosures "only serve to make a patentee's proffered damages amount appear modest by comparison, and to artificially inflate the jury's damages calculation." *LaserDynamics*, 694 F.3d at 68. Netlist cannot conceal its true purpose here: Mr. Kennedy's reference to the $1.1-billion Rambus license was only to skew the jury's damages horizon. Not even "a curative jury instruction"—much less Mr. Kennedy's comment that the license "should not be used to calculate the specific royalty"—"could have offset the resulting unfair prejudice." *Id.*; Appx5748.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 78

2. Mr. Kennedy's Alleged Reason for Ignoring Licenses to the Asserted Patents Applies Equally to the Rambus License

Mr. Kennedy compounded his error by shunning two far more relevant licenses—***Netlist's licenses for the asserted patents***. Appx5726; Appx5698.

The first factor in any hypothetical negotiation is "[t]he royalties received by the patentee for the licensing of the patent in suit." *Georgia-Pacific*, 318 F. Supp. at 1120. Here, Mr. Kennedy acknowledged there are "two real-world actual licenses" for the asserted patents. Appx5745-5746. Yet he told the jury not to use those licenses, claiming "they're not comparable" simply because they include non-cash consideration. Appx5746. Mr. Kennedy ignored that same reasoning for the Rambus license, which also includes significant non-cash consideration (e.g., a cross-license). Appx5790. Mr. Kennedy cannot have it both ways.

The true reason Mr. Kennedy eschewed Netlist's prior licenses is much simpler: Netlist licensed its entire portfolio for $8 million and $40 million. Appx5726; Appx5735-5736. Thus, his "exclusion of the [two] licenses expressly for the [asserted patents] served no purpose other than to [] increase the reasonable royalty rate above rates more clearly linked to the economic demand for the claimed technology." *LaserDynamics*, 694 F.3d at 80 (quotation marks and citation omitted). For this additional reason, the Court should vacate and remand.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 79

B. Netlist's Reliance on a Non-Analogous, Unreliable Regression Analysis Requires a New Damages Trial for the '054 and '918 Patents

1. Mr. Kennedy Was Merely a Mouthpiece for Dr. Groehn, Who Did Not Testify

To determine Samsung's incremental revenues from allegedly using the '054 and '918 patents in the accused DDR5 products, Mr. Kennedy turned to Dr. Groehn, who performed a "regression analysis" purporting to determine "empirically the relationship between [s]peed and price of [Samsung's] DDR4 products." Appx5703; Appx5777; Appx18517 (¶16).[10] But Dr. Groehn ***never testified at trial***, so direct testimony regarding his regression analysis is not in evidence. Netlist instead had Mr. Kennedy parrot Dr. Groehn's analysis over Samsung's pre-trial objection, and then rely on it.[11]

Because Dr. Groehn's regression analysis was an expert opinion, Netlist could not introduce it into evidence through Mr. Kennedy. "An expert witness may not simply summarize the out-of-court statements of others." *Robroy Indus.-*

[10] Mr. Kennedy assumed the alleged inventions improved power efficiency by 30%, *see* Section II.B.1, *supra*, without which a corresponding 20% speed decrease would occur. Appx5702 (696:18-22). He then relied on Dr. Groehn to price this alleged speed decrease. Appx5777-5778 (771:3-772:16). Applying Dr. Groehn's figure, Mr. Kennedy concluded that Samsung would have lost $196.3 million in revenue selling Netlist's proposed, slower noninfringing alternative. Appx5704-5705 (698:17-699:5).

[11] Nothing prevented Netlist from calling Dr. Groehn to the witness stand as "Dr. Groehn was present in the courtroom every day at trial." Appx18162.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 80

Tex., LLC v. Thomas & Betts Corp., 2017 WL 1319553, at *9 (E.D. Tex. Apr. 10, 2017) (Bryson, J., by designation) (quoting *United States v. Brownlee*, 741 F.3d 479, 482 (7th Cir. 2014)). If he does, he "is not giving expert testimony; ***he is a ventriloquist's dummy***." *Id*.

Although Federal Rule of Evidence 703 allows experts to rely on inadmissible evidence in some circumstances, it "was not intended to abolish the hearsay rule and to allow a witness, under the guise of giving expert testimony, to in effect become the mouthpiece of the witnesses on whose statements or opinions the expert purports to base his opinion." *Factory Mut. Ins. Co. v. Alon USA L.P.*, 705 F.3d 518, 524 (5th Cir. 2013) (quotation marks and citation omitted); *see also Mike's Train House, Inc. v. Lionel, L.L.C.*, 472 F.3d 398, 409 (6th Cir. 2006); *United States v. Grey Bear*, 883 F.2d 1382, 1392-93 (8th Cir. 1989). Yet that is precisely what the district court allowed Mr. Kennedy to do.

The district court wrongly faulted Samsung for not objecting. *Contra* Appx57-58. But Samsung could not have objected to Mr. Kennedy's reliance on Dr. Groehn's regression analysis because, when Mr. Kennedy testified, Dr. Groehn was still scheduled to testify. It was only after Mr. Kennedy's testimony that Netlist announced its decision to not call Dr. Groehn. Thus, Netlist alone knew its intention to deprive the record of Dr. Groehn's testimony. Samsung, therefore,

Samsung Electronics Co., Ltd.
Ex. 1077, p. 81

properly relied on its pre-trial objection to Mr. Kennedy's reliance on Dr. Groehn's opinion and did not object again. *See* Appx1488-1490; Appx1443-1448.

The district court, therefore, abused its discretion in allowing Mr. Kennedy to base his incremental-revenues calculation for the '054 and '918 patents on the opinion of an expert witness who did not testify at trial.

2. The Regression Analysis Says Nothing About DDR5 Prices

The district court further abused its discretion in denying Samsung's *Daubert* motion to exclude Mr. Kennedy's reliance on Dr. Groehn's analysis. Dr. Groehn's entire analysis focused on the speed/price relationship for Samsung's DDR***4*** products (the older 2014 standard), but Mr. Kennedy then applied that relationship to calculate damages for Samsung's DDR***5*** products (the current 2020 standard). Appx5777-5778; Appx6524; Appx6558.

His only explanation for relying on a DDR4 regression analysis to compute damages for DDR5 products was his conclusory assertion that "DDR5 is just the next generation." Appx5704; *see also* Appx18300 (¶218) ("DDR4 and DDR5 are sequential generations of DIMM and are used in the same kinds of applications."). Just because DDR4 and DDR5 are sequential generations, however, does not mean they share the same speed/price relationship, and Netlist presented no evidence on that point. The actual record evidence holds the contrary; Dr. Mangione-Smith

Samsung Electronics Co., Ltd.
Ex. 1077, p. 82

testified that DDR5 was a "significant improvement" over DDR4. Appx5384-5385.

The district court did not explain its denial of Samsung's motion to exclude Mr. Kennedy's arbitrary reliance a DDR4 regression analysis to compute damages for DDR5 products. Appx283. This Court should hold that the district court abused its discretion and remand for a new damages trial.

3. The Regression Analysis's Variables Are Incomplete and Collinear

The district court also abused its discretion in denying Samsung's *Daubert* motion to exclude Mr. Kennedy's reliance on Dr. Groehn's analysis because Dr. Groehn's regression model was fatally flawed, even for irrelevant DDR4 products.

"[A] hedonic regression model looks at the price of a good and tries to determine how much of that price is due to each of the attributes that enter into the market pricing." *Banks v. United States*, 721 F. App'x 928, 941 (Fed. Cir. 2017) (non-precedential) (quotation marks and citation omitted). Thus, as with all statistical analysis approaches, the choice of inputs usually dictates the results in regression analysis. For example, if one is trying to attribute the price of a car to its various parts, there are many choices of input variables. If one chooses as input variables the choice of engine, and then a collection of minor contributing factors like tire pressure, seat color, and the number of cupholders, the price difference

Samsung Electronics Co., Ltd.
Ex. 1077, p. 83

will largely be attributed to the engine choice. As a result, a regression model is flawed if it fails to (1) account for all independent variables that affect price or (2) avoid multicollinearity (failing to recognize that some variables are linked, like density and weight). *See Sheehan v. Purolator, Inc.*, 839 F.2d 99, 103 (2d Cir. 1988) (affirming finding that regression analysis was "flawed" because it "did not take into account various relevant . . . factors"); Appx1466 (handbook cited by Groehn explaining that "[m]issing or omitted or mismeasured variables . . . can cause serious problems for hedonic functions"); *see also Realcomp II, Ltd. v. F.T.C.*, 635 F.3d 815, 833-34 (6th Cir. 2011) (variables were properly excluded because "they were highly collinear with other explanatory variables and would render unreliable the results of the analysis").

Here, Dr. Groehn made two serious errors that collectively render his analysis unreliable. ***First***, he omitted relevant independent variables. The variables he analyzed were time, capacity, and speed. Appx18523 (¶29). Those variables, however, are not the only relevant variables of the products Dr. Groehn analyzed. Other independent variables include "Rank x Organization," which corresponds to the number of rows of memory chips, and "Component Composition," which relates to the number of dies. Appx18896-18897 (¶¶32-33); Appx5387. Dr. Groehn did not explain how he selected variables for his regression analysis, much less explain why he ignored the key Rank x

Samsung Electronics Co., Ltd.
Ex. 1077, p. 84

Organization and Component Composition variables. *See* Appx18511-18537. His regression analysis was therefore "flawed." *See Sheehan*, 839 F.2d at 103.

Second, the variables Dr. Groehn did use suffer from multicollinearity (interdependence). He admitted this fact at his deposition. Appx1677 ("[T]here is a correlation between the three variables, time, speed, and capacity, and . . . we know that speed and capacity both increase over time [and] that in general, on average, the devices become cheaper over time, and so that's a correlation between these three variables."); Appx1679. This interdependence amplifies the error introduced by Dr. Groehn's failure to include relevant independent variables. Accordingly, Dr. Groehn's analysis is unreliable. *See Realcomp II*, 635 F.3d at 834.

The district court did not explain why the flaws in Dr. Groehn's analysis "don't rise to the level of warranting outright exclusions." Appx347. That unexplained decision is part of a disturbing trend permitting black-box damages analysis that no judge or juror can parse. *See* William F. Lee & Mark A. Lemley, *The Broken Balance: How "Built-In Apportionment" and the Failure to Apply* Daubert *Have Distorted Patent Infringement Damages*, 37 Harv. J.L. & Tech. 255, 264 (2024) ("[I]n recent years, courts have repeatedly permitted parties to introduce deeply flawed regressions and conjoint survey analyses and have left it to the jury to decide their probative value and whether the expert did any real

Samsung Electronics Co., Ltd.
Ex. 1077, p. 85

apportionment.”). Damages law suffers under this flawed analysis, and this Court should intervene.

Because Dr. Groehn used incomplete and collinear variables, this Court should hold that the district court abused its discretion in permitting Dr. Groehn’s analysis and remand for a new damages trial as to the ’054 and ’918 patents.

CONCLUSION

On noninfringement, this Court should reverse the district court’s denial of Samsung’s motion for JMOL or a new trial. Should it be necessary to address damages, the Court should reverse the denial of JMOL of zero damages because Netlist failed to carry its burden of proof under a viable damages theory. Alternatively, the Court should vacate and remand for the district court to determine whether Netlist forfeited alternative damages theories and whether to order a new damages trial. Finally, the district court’s *Daubert* errors independently require a new damages trial.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 86

Dated: December 16, 2024

Respectfully submitted,

/s/ Ruffin B. Cordell

Ruffin B. Cordell
Lauren A. Degnan
Michael J. McKeon
Michael J. Ballanco
Benjamin J. Christoff
FISH & RICHARDSON P.C.
1000 Maine Avenue, S.W., Suite 1000
Washington, DC 20024
Telephone: (202) 783-5070

Francis J. Albert
FISH & RICHARDSON P.C.
12860 El Camino Real, Suite 400
San Diego, CA 92130
Telephone: (858) 678-5070

Alexander Pechette
FISH & RICHARDSON P.C.
One Marina Park Drive, Suite 1700
Boston, MA 02210
Telephone: (617) 542-5070

Attorneys for Defendants-Appellants

Samsung Electronics Co., Ltd.
Ex. 1077, p. 87

ADDENDUM

Samsung Electronics Co., Ltd.
Ex. 1077, p. 88

IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION

NETLIST, INC., *Plaintiff*, v. SAMSUNG ELECTRONICS CO., LTD., SAMSUNG ELECTRONICS AMERICA, INC., and SAMSUNG SEMICONDUCTOR, INC., *Defendants*.	§	CIVIL ACTION NO. 2:21-CV-00463-JRG [redacted]

MEMORANDUM OPINION AND ORDER

I. INTRODUCTION

Before the Court are two motions. The first, filed by Defendants Samsung Electronics Co., Ltd., Samsung Electronics America, Inc., and Samsung Semiconductor, Inc. (collectively, "Samsung") is Defendants' Combined Rule 50(b) Motion for Judgment as a Matter of Law and Rule 59 Motion for a New Trial (the "Motion"). (Dkt. No. 561.) In the Motion, Samsung moves under Rules 50(b) and 59, respectively, for judgment as a matter of law on multiple grounds or, in the alternative, for a new trial. (*Id.* at 1–4.) Plaintiff Netlist, Inc. ("Netlist") opposes the Motion. (Dkt. No. 573.) For the following reasons, the Court finds that this Motion should be **DENIED**.

The second motion before the Court is Defendants' Motion for Leave to File a Supplemental Brief in Support of Their Combined Rule 50(b) Motion for Judgment as a Matter of Law and Rule 59 Motion for a New Trial (the "Motion to File a Supplemental Brief") filed by Samsung. (Dkt. No. 591.) In the Motion to File a Supplemental Brief, Samsung requests leave of Court to file a brief in support of the Motion. (*Id.* at 1.) Netlist opposes the Motion. (Dkt. No. 594.) For the following reasons, the Court finds that this Motion should be **DENIED**.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 89

II. THE MOTION TO FILE A SUPPLEMENTAL BRIEF

The Court will briefly address the Motion to File a Supplemental Brief. Samsung argues that the Court should grant it leave to file a supplemental brief to address allegedly inconsistent positions that Netlist took in proceedings before the Patent Trial and Appeal Board ("PTAB") after the briefing on the Motion concluded. (Dkt. No. 591 at 1.) In response, Netlist disputes that it took any inconsistent positions and argues that what occurred after a jury trial ended could not have affected the outcome of the jury trial. (Dkt. No. 594 at 1–4.) In reply, Samsung argues that "Netlist's statements to the PTAB [] demonstrate that the trial evidence cannot support the jury's verdict." (Dkt. No. 595 at 1.) Samsung then contends that Netlist did indeed take inconsistent positions. (*Id.* at 1–2.) Netlist did not file a sur-reply

The Court finds that leave should not be granted to Samsung to file a supplemental brief. Netlist is correct that any positions it took in front of the PTAB after the jury trial concluded has no bearing on the record evidence presented to and considered by the jury. Samsung is not without remedy. While Samsung might argue to the PTAB that Netlist took inconsistent positions, such has no bearing on this district court litigation. Samsung effectively seeks to supplement the trial record after the fact. Such a request is **DENIED**.

III. LEGAL STANDARDS

A. Judgment as a Matter of Law

"Judgment as a matter of law is proper when 'a reasonable jury would not have a legally sufficient evidentiary basis to find for the party on that issue.'" *Abraham v. Alpha Chi Omega*, 708 F.3d 614, 620 (5th Cir. 2013) (quoting Fed. R. Civ. P. 50(a)). The non-moving party must identify "substantial evidence" to support its positions. *TGIP, Inc. v. AT&T Corp.*, 527 F. Supp. 2d 561, 569 (E.D. Tex. 2007). "Substantial evidence is more than a mere scintilla. It means such relevant

evidence as a reasonable mind might accept as adequate to support a conclusion." *Eli Lilly & Co. v. Aradigm Corp.*, 376 F.3d 1352, 1363 (Fed. Cir. 2004).

"The Fifth Circuit views all evidence in a light most favorable to the verdict and will reverse a jury's verdict only if the evidence points so overwhelmingly in favor of one party that reasonable jurors could not arrive at any contrary conclusion." *Core Wireless Licensing S.A.R.L. v. LG Elecs., Inc.*, 880 F.3d 1356, 1361 (Fed. Cir. 2018) (citing *Bagby Elevator Co. v. Schindler Elevator Corp.*, 609 F.3d 768, 773 (5th Cir. 2010)). A court must "resolve all conflicting evidence in favor of [the verdict] and refrain from weighing the evidence or making credibility determinations." *Gomez v. St. Jude Med. Daig. Div. Inc.*, 442 F.3d 919, 937–38 (5th Cir. 2006).

B. New Trial

Rule 59 provides that a new trial may be granted on all or part of the issues on which there has been a trial by jury for "any reason for which a new trial has heretofore been granted in an action at law in federal court." Fed. R. Civ. P. 59(a). Notwithstanding the broad sweep of Rule 59, "courts do not grant new trials unless it is reasonably clear that prejudicial error has crept into the record or that substantial justice has not been done, and the burden of showing harmful error rests on the party seeking the new trial." *Metaswitch Networks Ltd. v. Genband US LLC*, 2017 WL 3704760, at *2 (E.D. Tex. Aug. 28, 2017); *Erfindergemeinschaft UroPep GbR v. Eli Lilly & Co.*, 276 F. Supp. 3d 629, 643 (E.D. Tex. 2017). "A new trial may be granted, for example, if the district court finds the verdict is against the weight of the evidence, the damages awarded are excessive, the trial was unfair, or prejudicial error was committed in its course." *Smith v. Transworld Drilling Co.*, 773 F.2d 610, 612–13 (5th Cir. 1985); *see also Laxton v. Gap Inc.*, 333 F.3d 572, 586 (5th Cir. 2003) ("A new trial is warranted if the evidence is against the great, and not merely the greater, weight of the evidence."). Furthermore "[u]nless justice requires otherwise, no error in admitting or excluding evidence—or any other error by the court or a party—is ground for granting a new

trial . . . the court must disregard all errors and defects that do not affect any party's substantial rights." FED. R. CIV. P. 61.

IV. BACKGROUND

At trial, Netlist asserted U.S. Patent Nos. 10,949,339 (the "'339 Patent"), 8,787,060 (the "'060 Patent"), 9,318,160 (the "'160 Patent"), 11,016,918 (the "'918 Patent"), and 11,232,054 (the "'054 Patent") (collectively, the "Asserted Patents"). (*See* Dkt. No. 561 at 1.) Samsung asserted that all asserted claims of the '339 Patent, the '918 Patent, and the '054 Patent were invalid as failing to comply with the written description requirement of 35 U.S.C. § 112 ¶ 1. (*See* Dkt. No. 561 at 37.) The jury found that at least one asserted claim of the '339 Patent was infringed, at least one asserted claim of the '918 and '054 Patents was infringed, and at least one asserted claim of the '060 and '160 Patents was infringed; that Samsung's infringement was willful; and that Samsung had not shown any claim to be invalid. (*See* Dkt. No. 480 at 4–6.) The jury awarded damages in the amount of $33,150,000.00 for the infringement of the '339 Patent, $147,225,000.00 for the infringement of the '918 and '054 Patents, and $122,775,000.00 for the infringement of the '060 and '160 Patents. (*Id.* at 7.)

V. ANALYSIS

Samsung moves for judgment as a matter of law ("JMOL") of non-infringement, no willfulness, invalidity, and no damages. (Dkt. No. 561 at 11–71.) Samsung also moves for a new trial on the same grounds it seeks JMOL. In addition, Samsung alleges that Netlist engaged in improper conduct at trial and that the Court erroneously excluded certain evidence. (*Id.* at 11–75.) The Court will consider each in turn.

A. Non-Infringement

Samsung moves for JMOL of non-infringement on all Asserted Patents or, in the alternative, a new trial on non-infringement on all Asserted Patents. (*Id.* at 11–37.) The Court will address each ground in turn.

i. The '339 Patent

At trial, Netlist asserted that Samsung infringed Claims 1, 8, and 9 of the '339 Patent. (Dkt. No. 480 at 2.) Claims 8 and 9 depend from Claim 1. Claim 1 of the '339 Patent recites a memory module comprising, among other things, "a plurality of byte-wise buffers." (*Id.* at 12.) Each buffer must have a "first side" coupled to the computer and a "second side" coupled to multiple "ranks" of DDR DRAM devices. (*Id.*) When data is written to the memory module, the buffer must "actively drive a respective byte-wise section" of the data from the first side of the buffer to the second side of the buffer. (*Id.*) The Court construed the phrase "to drive" to mean "enabling only one of the data paths while the other possible paths are disabled." (*Id.* at 12 (citing (Dkt. No. 114 at 10).) Netlist accused certain Samsung DDR4 LRDIMMs of infringing these claims of the '339 Patent. (Dkt. No. 561 at 9.)

In the Motion, Samsung argues that "the accused Samsung LRDIMM products do not satisfy the 'drive' limitation because the buffers used in Samsung's products do not have the required 'fork-in-the-road.'" (*Id.* at 13–14.) According to Samsung, at trial Netlist relied on a path that was 8-bits (*i.e.*, one byte) wide, contending that the path was actually two paths, or "nibbles," each 4-bits wide. (*Id.* at 14.) Samsung argues that the nibbles cannot satisfy the "byte-wise" limitation because the "byte-wise" limitation requires the paths to be 8-bits wide under the specification of the '339 Patent. (*Id.* at 14–16.)

Samsung next argues that Netlist has not shown that the LRDIMM products "enable[e] only one of the data paths while the other possible paths are disabled," as the claims require. (*Id.*

at 17.) According to Samsung, Netlist pointed to data paths that were not within the buffer. (*Id.*) Samsung contends that paths not within the buffer cannot satisfy the claim language. (*Id.* (quoting Dkt. No. 114 at 10 ("[T]he claims also require logic configurable to enable a data path to 'actively drive' write data from one side of the buffer to the other side.")).) Samsung urges that Netlist never contended that the nibbles referenced above meet the selective enabling requirement. (*Id.* at 18.)

In response, Netlist first contends that the "enablement" limitation is satisfied. (Dkt. No. 573 at 1.) To this point, Netlist first argues that the Court's construction only requires selective enablement ***if*** there are other paths. (*Id.* at 2 (quoting Dkt. No. 114 at 10 ("A skilled artisan would understand 'driving' data from one side of the buffer to the other means, when there are multiple paths in a buffer through which data can be driven, enabling only one of the data paths while the other possible paths are disabled.")).) Second, Netlist argues that if the construction requires multiple paths, it presented evidence of them—the two nibbles. (*Id.* at 2.) Further, Netlist contends that its expert, Dr. Mangione-Smith, testified that the data paths can operate independently and can be programmed so that one path is enabled while the other path is disabled. (*Id.* at 2–3.) Finally, Netlist argues that Samsung's expert, Mr. McAlexander, presented evidence that a read path is disabled when a write path is enabled, thus satisfying the claim construction. (*Id.* at 3–4.)

Netlist next contends that the "byte-wise data paths" limitation is satisfied by the two nibble-wide paths that operate independently. (*Id.* at 4–5.) Netlist urges that "byte-wise" was given its plain and ordinary meaning, and that Samsung is attempting to equate "byte-wise" with "byte-wide." (*Id.* at 5.) The claims recite both "byte-wise" and "byte-wide" and different terms are generally deemed to have different meanings. (*Id.* (citing *Applied Med. Res. Corp. v. U.S. Surgical Corp.,* 448 F.3d 1324, 1333 n. 3 (Fed. Cir. 2006)).) Next, Netlist argues that the specification does not support Samsung's interpretation. (*Id.* at 6.)

CONFIDENTIAL INFORMATION HAS BEEN REMOVED FROM THIS PAGE

In reply, Samsung asserts that Netlist's arguments regarding the drive limitation fail. (Dkt. No. 577 at 1–2.) First, Samsung contends that the Court already rejected "Netlist's single-path argument" when it adopted the "fork-in-the road" construction. (*Id.* at 2.) Second, Samsung argues that the ability to set different latency delays for the nibbles does not mean that only one is enabled for a write operation. (*Id.*) Indeed, Samsung urges, it is undisputed that a [redacted] (*Id.*) Third, Samsung argues that this Court should reject Netlist's argument that the Court's construction is satisfied by enabling a write path while a read path is disabled because Magistrate Judge Payne rejected this same argument in other litigation. (*Id.* (citing *Netlist, Inc. v. Micron Tech., Inc.*, No. 2:22-cv-00203 (E.D. Tex.), Dkt. No. 249 at 25–26).)

Samsung then argues that Netlist's arguments concerning the "byte-wise" limitation are insufficient. (*Id.* at 2–3.) Samsung contends that this Court should reject Netlist's argument because Magistrate Judge Payne rejected this same argument in similar litigation. (*Id.* at 2–3 (citing *Netlist, Inc. v. Micron Tech., Inc.*, No. 2:22-cv-00203 (E.D. Tex.), Dkt. No. 249 at 24, 26).) Next, Samsung contends that Netlist's interpretation is untenable because it renders the claim term meaningless. (*Id.* at 3 (citing *Intel Corp. v. Qualcomm Inc.*, 21 F.4th 801, 810 (Fed. Cir. 2021)).)

In sur-reply, Netlist first addresses the "enabl[ement]" limitation. (Dkt. No. 588 at 1–3.) Netlist again argues that the Court's construction does not require multiple paths. (*Id.* at 1.) Netlist also contends, in a footnote, that decisions in other cases are unavailing because the only question before the Court is whether Netlist met its burden to prove infringement under the constructions provided. (*Id.* at n. 1.) Next, Netlist again urges that the two nibbles operate independently. (*Id.* at 1–2.)

Samsung Electronics Co., Ltd.
Ex. 1077, p. 95

Regarding the "byte-wise" limitation, Netlist argues that Samsung is making a claim construction argument, which it has waived. (*Id.* at 2–3 (citing *ePlus, Inc. v. Lawson Software, Inc.*, 700 F.3d 509, 520 (Fed. Cir. 2012)).)

As a preliminary matter, the Court agrees with Netlist that the only question before the Court is whether the jury's verdict is supported by sufficient evidence. In other words, it is irrelevant how Magistrate Judge Payne ruled on any similar issues in another case which is at a different stage.

First, the Court finds that there is substantial evidence that the LRDIMMs satisfied the "byte-wise data paths" limitation. The claims distinguish between "byte-wise" and "byte-wide," indicating that a "byte-wise data path[]" does not need to be eight bits wide. *Applied Med. Res. Corp,* 448 F.3d at 1333 n. 3 ("[T]he use of two terms in a claim requires that they connote different meanings … ."). Further, the jury heard evidence of two paths—the two nibbles. (*See, e.g.*, Dkt. No. 488 at 392:1–392:10 (Mangione-Smith) ("Q. Can you explain that for us? A. Sure. If we zoom in on this figure a little bit more, hopefully it becomes apparent that in the bottom right-hand edge, it refers to a lower nibble and an upper nibble. A nibble is just half of a byte. And if you look on the righthand side of that blow-up, hopefully it's apparent that there's two sets of these circuits stacked on top of each other. So there's actually one circuit for the upper half byte and one for the lower half byte.").) Samsung's remaining challenges regarding this term are claim construction challenges, which have long been waived. *See ePlus, Inc.*, 700 F.3d at 520 ("[I]f [defendant] desired such a narrow definition, it could (and should) have sought a construction to that effect.").

Second, the Court finds that there is substantial evidence that the LRDIMMs satisfied the "drive" limitation. Notably, the Court's Claim Construction Order required multiple paths when it

adopted the "so-called fork-in-the-road approach." (Dkt. No. 114 at 10.) There cannot be a fork within a single path. However, Netlist presented substantial evidence of two sets of multiple paths.

Netlist presented substantial evidence that the two nibbles are set up such that one is enabled while the other is disabled. (Dkt. No. 488 at 395:24–396:18 (Mangione-Smith) ("Q. And you're showing -- it's a little faded in the back. You're showing the back path is red and the front path is green. Do these two paths get enabled at the same time? A. No, they actually operate independently. They have their own latency delays that are programmed in."); *id.* at 396:4–18 (Mangione-Smith) ("Q. And is that how the system works? Is the first path enabled while all the other possible paths are disabled? A. That's exactly how it works.").)

The jury also heard evidence that a read path is disabled while the write path is enabled. (Dkt. No. 492 at 999:24–1000:15 (McAlexander) ("Q. And there is two data paths on the buffer, a read data path and a write data path. Is that correct? A. There is a read path and a write path, yes. Q. And they are both data paths. Is that correct? A. Sure. Q. And the -- on each of those read and write data paths, there's tri-state buffers. Correct? A. Yes. Q. And when the tri-state buffers are off, you can't transfer data on the path. Correct? A. Generally speaking, yes, that's correct. Q. And when the TX and RX are enabled, you will be able to transfer data on the path. Is that correct? A. When one of the other is enabled, yes. Q. And the read and write data paths can be turned off and on separately. Correct? A. One would hope so."); *id.* at 1000:23–25 (McAlexander) ("Q. And either the read path or the write path will be on. They will not be both on simultaneously. Correct? A. That is correct.").)

Nothing in the Court's Claim Construction Order required that the paths be from the same from one side of the buffer to the same other side of the buffer; in other words, nothing in the

Court's Claim Construction Order mandated that the enabled path and the disabled path be paths that flow in the same direction.

ii. The '060 and '160 Patents

Netlist accused certain Samsung HBM products of infringing claims 1, 5, and 7 of the '060 Patent and claim 5 of the '160 Patent. (Dkt. No. 561 at 19.) Claims 5 and 7 of the '060 Patent are dependent upon claim 1.

Samsung contends that there is not substantial evidence for the jury to find that the accused HBM products practiced two key limitations. First, all asserted claims of the '060 and '160 Patents require "array dies." The Court previously construed "array die" to mean an "array die that is different from a DRAM circuit " because Netlist disclaimed the use of "DRAM circuits" during prosecution. (Dkt. No. 114 at 32.) Netlist had argued in its *Markman* brief that it disclaimed only the use of the specific DRAM circuits in prior art reference Rajan, to the extent it disclaimed anything at all during prosecution. (*See* Dkt. No. 114 at 31–32.) However, the Court rejected that position. (*Id.* at 32.) Further, it is undisputed that the accused HBM products contain "DRAM cells" in "DRAM cores." (Dkt. No. 488 at 582:4–6 (Brogioli).)

Second, the asserted claims of the '060 and '160 Patents also require "die interconnect[s]" that are "not in electrical communication with" certain "array dies." (Dkt. No. 561 at 23.) It is undisputed that "the signal path in the accused products travels across each and every core die." (*See id.*)

Turning now to the parties' arguments, Samsung first argues that there is insufficient evidence to support the jury's finding that the accused Samsung HBM products contain an "array die that is different from a DRAM circuit." (*Id.* at 19–23.) Samsung contends that Dr. Brogioli, one of Netlist's technical experts, improperly distinguished between a "DRAM circuit" and a "DRAM circuit as that term is used in the Court's claim construction," because his distinction is

really a distinction between a DRAM circuit and the DRAM circuits used in Rajan. This is erroneous because the Court did not limit the disclaimer to the DRAM circuits in Rajan. (*Id.* at 20–22.) Next, Samsung contends that Dr. Brogioli erroneously distinguished between the claimed "array dies" and the disclaimed "DRAM circuits" based on the type of die interconnect used. (*Id.* at 22–23.) According to Samsung, Netlist's contention that "array dies" must use through-silicon vias ("TSVs") and that "DRAM circuits" must use "wire bonding," is contrary to the teaching of the '060 and '160 Patents. (*Id.* at 22.) Samsung urges that the '060 and '160 Patents list wire bonding as an acceptable method of connecting array dies. (*Id.*) Further, Samsung argues that the fact that TSVs are specified as an additional limitation of a dependent claim reinforces the point that "array dies" in claim 1 do not require the use of TSVs. (*Id.*) Samsung also argues that Dr. Brogioli offered no evidence to support his contention that all "DRAM circuits" use wire bonding, other than as a reference to Rajan. (*Id.* at 23.)

Regarding the "die interconnect" limitation, Samsung argues that Netlist failed to prove that the accused products have "die interconnects" that are "not in electrical communication with" certain "array dies." (*Id.* at 23–24.) According to Samsung, Netlist's argument that the die interconnects are not in communication with certain array dies because the path does not connect to a "data port" or "receiver[]" is legally insufficient because it does not comport with the specification. (*Id.* at 23–24.) Samsung urges that the specification refers to die interconnects being in electrical communication with various parts of a memory product, and also refers to electrical communication with the array dies themselves. (*Id.* at 24.) Thus, Samsung argues, when the claims recite electrical communication with the "array dies" themselves, they mean to claim electrical communication with any portion of the array die. (*See id.*) Samsung contends that since it is

CONFIDENTIAL INFORMATION HAS BEEN REMOVED FROM THIS PAGE

undisputed that the signal path electrically connects every core die, Netlist cannot show that the accused HBM products practice this limitation. (*Id.*)

In response, Netlist first addresses the "array dies" limitation. (Dkt. No. 573 at 6–10.) Netlist first notes that Samsung does not dispute that its HBM products contain memory dies having array dies. (*Id.* at 7–8.) Next, Netlist contends that Dr. Brogioli explained to the jury that the core dies in Samsung's HBM products are not a "DRAM circuit." (*Id.* at 8.) Netlist notes that Samsung raised the argument that Dr. Brogioli did not apply the plain and ordinary meaning at the *Daubert* stage, and it was rejected. (*Id.* (citing Dkt. No. 202 at 7–9, Dkt. No. 432 at 16).) Netlist then contends that Dr. Brogioli did apply the plain and ordinary meaning of "DRAM circuit"—the only place "DRAM circuit" appears in the intrinsic record evidence is in connection with Rajan. Samsung does not dispute that the DRAM circuits in its HBM products are different than those described in the intrinsic evidence. (*Id.* at 8–9.) Further, Netlist contends that [redacted] (*Id.*at 9.) Netlist argues that "the Court's construction of 'array die' does not require that the 'array dies' cannot contain 'DRAM circuits.'" (*Id.* at 10.) Finally, Netlist argues that since [redacted] the jury could have found that Samsung's accused products have array dies that contain DRAM circuits, but which are different from DRAM circuits. (*Id.*)

Regarding the "not in electrical communication" limitation, Netlist argues that Samsung is making a claim construction argument by construing "electrical communication" to mean "electrical connection," and has waived this construction. (*Id.* at 10–11(citing *Music Choice v. Stingray Digital Grp., Inc.*, No. 2:16-cv-00586, 2019 WL 8110069, at *3 (E.D. Tex. Nov. 19, 2019)).) Netlist then contends that the jury could accept Dr. Brogioli's testimony that mere

CONFIDENTIAL INFORMATION HAS BEEN REMOVED FROM THIS PAGE

electrical connection does not constitute electrical communication. (*Id.* at 11.) Dr. Brogioli explained [redacted]. (*Id.* at 11–13.)

In reply, Samsung first addresses the "array dies" limitation. (Dkt. No. 577 at 3–5.) Samsung first argues that Netlist is equating "DRAM circuit" to the DRAM circuit described in Rajan, which the Court has already rejected. (*Id.* at 3–4.) Samsung also argues that Dr. Brogioli did not apply the plain and ordinary meaning of "DRAM circuit," but the definition laid out in Rajan. (*Id.* at 4.) Samsung then contends that Netlist has not responded to its argument that numerous documents use the terms "DRAM" and "DRAM circuit" without requiring such circuits to be wire bonded as Netlist urges. (*Id.* at 4–5.) Samsung then urges that it is irrelevant that the Court denied Samsung's *Daubert* motion concerning Dr. Brogioli because the Court "must disregard" expert testimony that is "based on an incorrect understanding of the claim construction," as is Dr. Brogioli's. (*Id.* at 5 (quoting *Cordis Corp. v. Boston Sci. Corp.*, 658 F.3d 1347, 1357 (Fed. Cir. 2011)).) Finally, Samsung argues that Netlist's alternative argument—that the jury may have found that the accused HBM products contain DRAM circuits but that the array die as a whole is different than a DRAM circuit—was not presented to the jury by Netlist's expert. (*Id.*)

Regarding the "not in electrical communication with" limitation, Samsung argues that Netlist is improperly attempting to rewrite the plain language of the claims by contending that the interconnects do not connect to the data ports or receivers on certain dies. (*Id.* at 5–6.)

In sur-reply, Netlist first addresses the "array dies" limitation. (Dkt. No. 588 at 3–4.) Netlist argues that the jury heard ample evidence to support a finding that the accused products have array

CONFIDENTIAL INFORMATION HAS BEEN REMOVED FROM THIS PAGE

dies containing DRAM circuits. (*Id.* at 3.) Netlist then contends that it did not limit its definition of "DRAM circuits" to the DRAM circuits described in Rajan; [redacted]. (*Id.* at 3–4.) Netlist then argues that the Court has already rejected these arguments at the *Daubert* stage and should reject them here as well. (*Id.* at 4.)

Regarding the "not in electrical communication" limitation, Netlist largely re-urges the same arguments it made in its response. (*See id.* at 4–5.)

The Court finds that Samsung's arguments, as to both disputed limitations, are without merit. First, there is legally sufficient evidence to support the jury's finding of infringement with respect to the "array dies" limitation. (*See* Dkt. No. 489 at 516:10–12 (Brogioli) [redacted].) Moreover, Samsung raised these same arguments at the *Daubert* stage and lost—Samsung does not now get a second bite at that same apple. (*See* Dkt. No. 202 at 7–9.) The fact that the Court ruled against Samsung at the *Daubert* stage shows that the Court already is persuaded that Dr. Brogioli's testimony is not inconsistent with the claim language. (*See* Dkt. No. 202 at 7–9 (Samsung arguing that Dr. Brogioli's opinion regarding the "array dies" limitation is inconsistent with the Court's construction); Dkt. No. 432 at 8 (denying Samsung's Motion to Strike Expert Testimony of Dr. Michael Brogioli).) Samsung's arguments that depend on the insufficiency of Dr. Brogioli's testimony are rejected on this ground.

Regarding Netlist's alternative theory—that the jury could have found that Samsung's core dies have components and functionality other than the basic DRAM cell and that the jury could have found that Samsung's accused products have array dies that contain DRAM circuit's but that

CONFIDENTIAL INFORMATION HAS BEEN REMOVED FROM THIS PAGE

are themselves different from DRAM circuit's—the Court finds that it was not squarely presented to the jury, and so it is insufficient on its own.

The Court finds that substantial evidence was presented to the jury concerning the "not in electrical communication with" limitation. Netlist presented evidence [redacted] (Dkt. No. 489 at 517:20–518:10 (Brogioli) [redacted]; *id.* at 492:16–24 (Brogioli) [redacted]; *id.* at 519:21–520:9 (Brogioli) [redacted]

Samsung Electronics Co., Ltd.
Ex. 1077, p. 103

CONFIDENTIAL INFORMATION HAS BEEN REMOVED FROM THIS PAGE

[redacted] Additionally, the Court rejects Samsung's attempt to equate "electrical communication" with "electrical connection." Simply because a die interconnect is electrically connected to an array die, does not mean it is in electrical communication with the array die.

iii. The '918 Patent

Netlist asserted that Samsung's SODIMM, UDIMM, and RDIMM products infringe claims 16, 18, and 19 of the '918 Patent and that Samsung's RDIMM products also infringe claims 1, 5, and 13 of the same. (Dkt. No. 561 at 24.) All asserted claims require a "converter circuit" configured to "provide a fourth regulated voltage." (*Id.*) The Court did not construe the term "converter circuit," so the plain and ordinary meaning applies. (*See id.* at 25.) Netlist contended that a low drop out ("LDO") regulator in the accused products was the required converter circuit. (*Id.*)

Samsung argues that Netlist failed to present any evidence that the plain and ordinary meaning of "converter circuit" encompasses an LDO regulator. (*Id.*) Instead, Samsung contends, Netlist treated the limitation as purely functional, such that anything that converts a voltage satisfies the limitation. (*Id.*) Netlist asserted that "[t]he claims impose no structural limitation on the claimed converter circuit," and Samsung contends that this assertion is wrong as a matter of law. (*Id.* at 26.) Samsung urges that purely functional limitations fail to comply with the Patent Act's requirement that the inventor "particularly point out and distinctively claim the part, improvement, or combination which he claims as his invention or discovery," unless they are means-plus-functions claims, and that Netlist failed to seek such a construction. (*Id.* at 26–27 (quoting *Halliburton Oil Well Cementing Co. v. Walker*, 329 U.S. 1, 9–10 (1946)).) Thus, Samsung argues, Netlist has relied on an impermissible functional construction. (*Id.* at 27.)

In a footnote, Samsung contends that a "converter circuit" must include a "switch" in which the switch is continually switching on and off the voltage to provide charge to an inductor. (*Id.* at 26 (quoting Dkt. No. 493 at 889:23–892:19 (McAlexander)).)

Samsung also argues that Netlist has failed to prove infringement under the doctrine of equivalents because the only testimony on the subject was a conclusory sentence by Dr. Mangione-Smith: "[an LDO is] equivalent because it does the same function which is to convert; it does it the same way, which is to reduce the input voltage; and it achieves the exact same result, which is a regulated output voltage." (*Id.* at 27 (quoting Dkt. No. 489 at 337:15–18 (Mangione-Smith)).) Samsung also argues that Dr. Mangione-Smith did not properly address the "way" prong of the function-way-result test because his testimony in this regard was conclusory. (*Id.*) Finally, Samsung argues that Dr. Mangione-Smith's own testimony contradicts any claim that an LDO regulator is equivalent to a converter circuit. (*Id.*) The only circuits the specification describes as converter circuits are buck converters, boost converters, and buck-boost converters. (*Id.*) However, Samsung notes, Dr. Mangione-Smith admitted that "buck converters and LDOs have fundamentally different characteristics as well as operational benefits and limitations," and that LDO regulators and buck converters are not interchangeable. (*Id.* at 27–28 (quoting Dkt. No. 489 at 415:5–8, 421:10–13).) According to Samsung, these admissions negate his doctrine of equivalents opinion. (*Id.* at 28.)

In response, Netlist argues that it presented substantial evidence that the LDO regulator is the claimed "converter circuit." (Dkt. No. 573 at 14–15.) Netlist then argues that the term "converter circuit" denotes structure, a circuit, and that Samsung is improperly attempting to import additional structural limitations to the term. (*Id.* at 15.) Netlist contends that this is a claim construction argument, which, at this stage, has been waived. (*Id.* at 15–16.) Even if the Court

CONFIDENTIAL INFORMATION HAS BEEN REMOVED FROM THIS PAGE

were to entertain claim construction at this stage, Netlist contends that importation of a "switch" limitation into the term would contravene the basic principles of claim construction. (*Id.* at 16.) In any event, Netlist argues, its expert, Dr. Mangione-Smith, testified that the [redacted]." (*Id.* at 19 (quoting Dkt. No. 489 at 349:9–350:10 (Mangione-Smith)).)

Netlist also argues that it did not rely on an impermissible functional construction because the term "converter circuit" requires structure, namely "electrical components for voltage conversion." (*Id.* at 17–19.) Further, Netlist contends it put forth substantial evidence that Samsung infringed under the doctrine of equivalents—[a]ll of the evidence of record, including that discussed above, regarding the function and structure of the LDO supports not only literal infringement by infringement by equivalents." (*Id.* at 19–20.) Moreover, Netlist urges, Dr. Mangione-Smith's testimony that buck and boost converters are different from LDO regulators is irrelevant to the doctrine of equivalents analysis, which requires a comparison of the accused products to the claims, not the specification. (*Id.* at 20.) Finally, Netlist argues that Dr. Mangione-Smith did properly address the "way" prong of the function-way-result test by testifying that the converter circuit accomplishes the same result of a converter circuit by reducing the input voltage. (*Id.* at 20–21 (citing Dkt. No. 489 at 337:11–18 (Mangione-Smith)).)

In reply, Samsung re-urges that the term does not have structure, and argues that "electrical components" is a "meaningless phrase in the context of a memory module, which is nothing but electrical components." (Dkt. No. 577 at 6.) Samsung also contends that Netlist's definition is at odds with Magistrate Judge Payne's claim construction order in a similar case, where he rejected Netlist's position that the plain meaning of "converter circuit" was "a circuit for voltage conversion." (*Id.* at 6–7 (citing *Netlist, Inc. v. Micron Tech., Inc.*, No. 2:22-cv-00203 (E.D. Tex.), Dkt. No. 249 at 30, 32).) Samsung then argues that there is no testimony supporting Netlist's

contention that an LDO regulator is a voltage converter. (*Id.* at 7–8.) Finally, Samsung urges that Netlist fails to explain how the evidence in the record provides support for its doctrine of equivalents theory. (*Id.* at 8.)

In sur-reply, Netlist argues that Dr. Mangione-Smith testified that the LDO regulator in the accused products satisfied the "converter circuit" limitation. (Dkt. No. 588 at 5.) Netlist argues that Magistrate Judge Payne's *Markman* order confirms Netlist's interpretation. (*Id.* at 6–7 (citing *Netlist, Inc. v. Micron Tech., Inc.*, No. 2:22-cv-00203 (E.D. Tex.), Dkt. No. 249 at 31–32).) Finally, Netlist argues that Samsung's reply ignores Netlist's arguments and evidence that it presented in its opposition brief. (*Id.* at 7.)

The Court finds that there is substantial evidence that the LDO regulator in the accused products satisfies that "converter circuit" limitation. (*See, e.g.*, Dkt. No. 489 at 336:1–2 (Mangione-Smith) ("Q. You identified the LDO as the converter circuit? A. Yes, sir.") Samsung's contention that the "converter circuit" limitation does not contain structure is unpersuasive. The Court agrees with Netlist that the language of the claim requires electrical components for voltage conversion. Samsung contends that the phrase "electrical components for voltage conversion" is "meaningless," but it is just broad, not without meaning. (*See* Dkt. No. 577 at 6.) Samsung has not pointed to any authority that prevents structure from being defined broadly.

With regard to the doctrine of equivalents, the Court understands that the testimony elicited by Netlist might be viewed as conclusory and deficient when analyzed in complete isolation. Indeed, Netlist only elicited one sentence from Dr. Mangione-Smith that, standing alone could be considered as conclusory. (Dkt. No. 489 at 337:15–18 (Mangione-Smith) ("[an LDO is] equivalent because it does the same function which is to convert; it does it the same way, which is to reduce the input voltage; and it achieves the exact same result, which is a regulated output voltage.").)

However, this is to be considered, not in utter isolation, but against the backdrop of the entirety of Dr. Mangione-Smith's testimony. *Paice LLC v. Toyota Motor Corp.*, 504 F.3d 1293, 1305 (Fed. Cir. 2007) ("Our 'particularized testimony' standard does not require [an expert] to re-start his testimony at square one when transitioning to a doctrine of equivalents analysis. Indeed, we think it desirable for a witness to incorporate earlier testimony in order to avoid duplication.") In this context, Samsung objection as conclusory fails.

Samsung's argument that Dr. Mangione-Smith did not adequately address the "way" prong of the way-function-result test is also without merit because Dr. Mangione-Smith testified that the "way" the converter circuit accomplishes the result is by reducing the input voltage. (Dkt. No. 489 at 337:11–18 (Mangione-Smith).)

The fact that Dr. Mangione-Smith admitted that "buck converters and LDOs have fundamentally different characteristics as well as operational benefits and limitations" is wholly irrelevant. (*Id.* at 415:5–8.) Samsung appears to be making a claim construction argument: since the specification only describes buck converters, boost converters, and buck-boost converters as converter circuits, the "converter circuit" term must be equivalent to those structures under the doctrine of equivalents. Not so. Samsung cites no law indicating as much.

The Court will not entertain Samsung's late claim construction request to import a limitation from the specification—a "switch"—into the claim language. Even if the Court were to grant such a request, it would be satisfied. Netlist put on substantial evidence that the LDO regulators have a switch. (*See, e.g.*, Dkt. No. 389 at 349:9–350:10, 372:19–373:12 (Mangione-Smith).) This gets Samsung nowhere.

Finally, as discussed above, the question before the Court is whether the jury's verdict is supported by substantial evidence and so proceedings in other related cases are largely irrelevant, including Magistrate Judge Payne's *Markman* Order.

iv. The '054 Patent

Netlist asserted claims 16 and 17 of the '054 Patent against Samsung's DDR5 SODIMM, UDIMM, and RDIMM memory modules. (Dkt. No. 561 at 28.) Claim 17 depends from claim 16, which recites a "memory module comprising," among other things:

> a voltage monitor circuit coupled to an input voltage received from the host system via the interface, the voltage monitor circuit configured to detect an amplitude change in the input voltage, wherein, in response to the voltage monitor detecting an amplitude change in the input voltage, the memory module transitions from a first operable state to a second operable state.

(*Id.* at 28–29.) The parties agreed to construe "operable state" as a "state in which the memory module is operated." The parties further agreed to construe "first operable state" as "state in which the memory module is operated before transition," and "second operable state" as "state in which the memory module is operated after transition." (Dkt. No. 91-1 at 4.) Both sides' experts agreed that in the event of an amplitude change in the input voltage, the DRAM on the module is disabled, though some low-level circuitry receives power. (Dkt. No. 561 at 29.)

Samsung argues that Netlist has not shown sufficient evidence, as a matter of law, that the memory module has entered a "second operable state." (Dkt. No. 561 at 29.) According to Netlist's expert, Dr. Mangione-Smith, the memory module is in an operable state as long as any component receives power. (*Id.*) Samsung urges that Dr. Mangione-Smith's testimony is conclusory, as he presented no explanation as to why "powered" means "operated." (*Id.* at 30.) Further, Samsung argues that "[n]o reasonable jury could have found that the memory module is in an 'operable state' when the memory is turned off, thereby erasing the contents and rendering the memory unusable." (*Id.*)

Netlist argues in response that the jury is entitled to credit Netlist's evidence that the memory module is still in an operable state when module components are still operating. (Dkt. No. 573 at 21–22.) Netlist contends that many of Samsung's arguments in this regard are premised on an erroneous conclusion that the specification requires that the second state be for the purpose of transferring data from volatile memory to non-volatile memory on the module. (*Id.* at 22.) Netlist contends that this is improper *post-hoc* claim construction that Samsung has waived. (*Id.* at 22–23.) Netlist then describes various evidence presented to the jury. (*See id.* at 23–24.)

In reply, Samsung argues that the evidence shows that in the state identified by Netlist, the memory products are unable to perform memory operations and thus the memory module is not "operable" in the state identified by Netlist. (Dkt. No. 577 at 8–9.) Samsung then contends that it is not advancing a new claim construction, but demonstrating how Netlist's evidence fails to support the verdict under the plain and ordinary meaning of the agreed-upon construction, which is consistent with the specification. (*Id.* at 9.)

In sur-reply, Netlist notes that the claim recites "operable state" of a memory module, not memory devices. (Dkt. No. 588 at 7.) Netlist then describes various evidence presented to the jury. (*See id.* at 7–9.)

The Court finds that Samsung's request should be denied. Fundamentally, this dispute boils down to whether a memory module with some low-level circuitry that is operating, but nothing else (including the memory devices themselves), is in an "operable" state. Samsung said no, Netlist said yes, and the jury agreed with Netlist. This is a classic factual dispute—both sides presented their evidence, and the jury was entitled to credit Netlist's over Samsung's. Samsung has pointed to no reason why Netlist's evidence is deficient. Rather, Samsung merely re-argues the position it took at trial.

CONFIDENTIAL INFORMATION HAS BEEN REMOVED FROM THIS PAGE

To be clear, there is substantial evidence that the memory module has a "second operable state." (*See, e.g.*, Dkt. No. 389 at 378:22–379:13 (Mangione-Smith) ("In the second operable state, after an overvoltage condition, for example, the DRAMs are turned off but other elements in the system continue to operate, such as the SPD that we talked about and temperature sensors and the RCD."); Dkt. No. 493 at 967:7–10 (McAlexander) [redacted]

v. Infringement of DDR5 Products

Netlist's infringement theories for the '918 and '054 Patents focused primarily on the Power Management Integrated Circuits ("PMICs") used in the accused DDR5 products. (Dkt. No. 561 at 31.) Across its entire line of DDR5 memories, consisting of 143 different product models, Samsung uses numerous different PMICs, supplied by at least four different companies: Samsung, Texas Instruments, Renesas, and Monolithic Power Solutions. (*Id.*)

Samsung argues that Netlist only introduced evidence as to three PMICs—the Samsung S2FPD01, Samsung S2FPC01, and Renesas P8911—and no others. (*Id.* at 31–32 (quoting *Medtronic Vascular, Inc. v. Boston Sci. Corp.*, 2008 WL 2744909, at *3 (E.D. Tex. July 11, 2008) ("A patentee . . . cannot simply assume that all of the accused products are like the one plaintiff's expert tested and thereby shift to the accused infringer the burden to show that is not the case." (cleaned up))).) Samsung also contends that, as to the three PMICs discussed at trial, Netlist's expert improperly relied on different PMICs to satisfy different limitations. (*Id.* at 32.) Specifically, Samsung urges that Dr. Mangione-Smith mixed-and-matched limitations from two Samsung PMICs, the S2FPD01 and the S2FPC01, and failed to opine that either PMIC both had "the required 'buck converters *and* was programmed to receive a pre-regulated input voltage.'" (*Id.* (emphasis in original).) Additionally, Samsung argues, Netlist did not present evidence that the Renesas P8911 PMIC had "the required buck converters or that it is 'configured to receive a pre-

CONFIDENTIAL INFORMATION HAS BEEN REMOVED FROM THIS PAGE

regulated input voltage.'" (*Id.* at 32–33) Finally, Samsung argues that Netlist's failure infects all asserted claims of the '918 and '054 Patents. (*Id.* at 33.)

In response, Netlist first argues that Dr. Mangione-Smith testified that there was no material difference between the accused products. (Dkt. No. 573 at 25–26 (citing *TiVo, Inc. v. EchoStar Commc'ns Corp.*, 516 F.3d 1290, 1308 (Fed. Cir. 2008)).) Next, Netlist argues that Dr. Mangione-Smith testified with regards to the S2FPD01 that "each of those blocks [the buck converters and LDO converter circuit] ***receives the pre-regulated input voltage*** and produces a regulated output voltage." (*Id.* at 26 (quoting Dkt. No. 489 at 334:25–335:18) (alterations in original).) Netlist also contends that Dr. Mangione-Smith presented the jury with the PMIC diagram from the [redacted] (Dkt. No. 573 at 26–27.) Netlist then argues that Samsung's arguments have no merit, which could infect the claims. (*Id.* at 27.)

In reply, Samsung argues that the context in which Dr. Mangione-Smith testified that there is "no material difference between these products" makes clear that "these products" was referring to the three types of DDR5 products—RDIMMs, UDIMMs, and SODIMMs—not all 143 products. (Dkt. No. 577 at 9–10 (quoting Dkt. No. 489 at 332:7–8).) Samsung then argues that *TiVo* is inapposite. (*Id.* at 10 (citing 516 F.3d 1290).)

In sur-reply, Netlist argues that Dr. Mangione-Smith testified that he analyzed all the products and found that there were no material differences between them, and that he frequently reminded the jury that his analysis applied to all accused products. (Dkt. No. 588 at 9–10.) Netlist also notes that Samsung does not cite any evidence that there were any material differences among the DDR5 products. (*Id.* at 10 (citing *TiVo*, 516 F.3d at 1308).) Netlist then argues that Samsung's

expert, Mr. McAlexander, made admissions from which the jury could have reasonably found infringement of all accused DDR5 devices. (*Id.* at 10–11.)

The Court is not persuaded by Samsung's arguments. As Netlist correctly notes, "there is nothing improper about an expert testifying in detail about a particular device and then stating that the same analysis applies to other allegedly infringing devices that operate similarly, without discussing each type of device in detail." *TiVo*, 516 F.3d at 1308. Samsung has not identified a material difference in the products, and Dr. Mangione-Smith testified that he "analyzed all of these products" and that found there was "no material difference between these products in terms of infringement." (Dkt. No. 489 at 332:1–12.) In light of these circumstances, the Court is not persuaded by Samsung's arguments.

The Court also finds that Netlist presented substantial evidence that the S2FPD01 and the S2FPC01 PMICs contained buck converters and were programmed to receive a pre-regulated input voltage. Samsung does not dispute that Dr. Mangione-Smith testified with regards to the S2FPD01 that "each of those blocks [the buck converters and LDO converter circuit] receives the pre-regulated input voltage and produces a regulated output voltage." (Dkt. No. 489 at 334:25–335:18.)

Similarly, the Court finds that Netlist presented substantial evidence that the Renesas P8911 PMIC contained buck converters and was programmed to receive a pre-regulated input voltage. Samsung does not contest that Netlist presented the jury with a Renesas data sheet, JTX30, showing three buck converters and a pre-regulated input voltage of 5 Volts. (*See also* Dkt. No. 489 at 337:7–10 (Mangione-Smith) ("And I put a figure for JTX 30 on the right-hand side which is the data sheet for Renesas's PMIC.").)

vi. Infringement of the DDR4 Products

Samsung argues that Netlist took an improper "mix-and-match" approach to the '339 Patent. (Dkt. No. 561 at 33.) Netlist accused at least 90 different DDR4 products under the '339

Patent but presented evidence as to only two models: the M386A8K40BM1 and M386A8K40BM2. (*Id.* at 34.) Thus, Samsung argues it is entitled to JMOL that no other model infringes. (*Id.* (citing *Medtronic Vascular*, 2008 WL 2744909, at *3).)

Samsung argues that Netlist's evidence also fails as to the M386A8K40BM1 and M386A8K40BM2 models because Netlist failed to present any evidence that the products have the claimed "data buffers." (*Id.*) According to Samsung, Netlist's expert, Dr. Mangione-Smith relied on an integrated circuit manufactured by third-party Renesas, the 4DB0232KC2, but did not present any evidence that the M386A8K40BM1 or M386A8K40BM2—or any other accused product—contained a Renesas 4DB0232KC2 data buffer. (*Id.*) Samsung then contends that this failure is apparent from Netlist's demonstratives. (*Id.* at 34–36.)

In response, Netlist contends that Dr. Mangione-Smith explained that his analysis applied to Samsung's "DDR4 LRDIMM products," and used two representative models, the M386A8K40BM1 ("–BM1") and M386A8K40BM2 ("–BM2"). (Dkt. No. 573 at 28 (citing *TiVo, Inc.*, 516 F.3d at 1308).) Netlist then contends that Dr. Mangione-Smith presented evidence that the –BM1 and –BM2 devices contained a Renesas 4DB0232KC2 ("–KC2") data buffer. (*Id.*) Netlist contends that its slide (as corrected) referenced JTX28—which is a datasheet covering two representative DDR4 LRDIMM products, –BM1 and –BM2—and contained blue boxes showing the location of data buffers. (Dkt. No. 573 at 29.) Further, Netlist urges that Dr. Mangione-Smith made clear that he was discussing JTX-28 when discussing this slide and the data buffers in the accused products. (Dkt. No. 573 at 30.) Finally, Netlist argues that Dr. Mangione-Smith explained how all the same functionality that he discussed concerning the Renesas data buffer operates in the –BM1 and –BM2 products, as reflected in the datasheets. (*Id.* at 31–32.)

In reply, Samsung argues that Dr. Mangione-Smith never (a) testified that he was treating the –BM1 and –BM2 products as representative, (b) claimed to have examined any of the other 88 products, or (c) discussed how the other 88 products related to the two he addressed. (Dkt. No. 577 at 10–11.) Even as to the –BM1 and –BM2 products, Samsung contends that Netlist presented no evidence that the –KC2 buffer is used in either product. (*Id.* at 11.) Samsung also argues that Netlist does not show how the same functionality in the –KC2 buffer is present in –BM1 and –BM2. (*Id.*) Further, Samsung argues, Netlist does not explain how the products satisfy the "drive" limitation. (*Id.*)

In sur-reply, Netlist notes that Dr. Mangione-Smith began his analysis by testifying that the "DDR4 LRDIMM products" are the accused Samsung products that use the inventions of the '339 Patent. (Dkt. No. 588 at 11.) Netlist argues that Dr. Mangione-Smith then relied on JTX28 and "testified in detail" how they infringe the '339 Patent. (Dkt. No. 588 at 11 (quoting *TiVo*, 516 F.3d at 1308).) Next, Netlist argues that Dr. Mangione-Smith did analyze the functionality of the –KC2 buffer and shows how that same functionality is present in the –BM1 and –BM2 products. (*Id.*)

The Court is not persuaded by Samsung's arguments. Dr. Mangione-Smith clearly treated the accused DDR4 products as materially the same for the purposes of infringement throughout the trial, testifying that the "DDR4 LRDIMM products" are the accused Samsung products under the '339 Patent. (Dkt. No. 489 at 387:15–16.) Further, Samsung has not identified any material differences between any two models. Dr. Mangione-Smith then analyzed the –BM1 and –BM2 models in depth, indicating that these models were representative. In light of all these circumstances, the jury was likewise entitled to treat the –BM1 and –BM2 models as representative. *See TiVo*, 516 F.3d at 1308.

CONFIDENTIAL INFORMATION HAS BEEN REMOVED FROM THIS PAGE

Additionally, there was substantial evidence that the accused products contain data buffers. Dr. Mangione-Smith discussed JTX28 when discussing data buffers in the accused products. (Dkt. No. 395:2–5 ("This image is from JTX 28, and it sort of shows that there's a relationship, but if we look -- relationship in terms of position. It shows that the signal lines go into those data buffers.").) Dr. Mangione-Smith also showed a slide to the jury containing a portion of JTX28 with boxes over the data buffers, highlighting their location. (*See* Dkt. No. 573 at 28–29.)

The Court finds that there was substantial evidence for the jury to find that the –BM1 and –BM2 products contained the –KC2 data buffer. Dr. Mangione-Smith provided testimony concerning the functionality of Renesas –KC2 data buffer and how that same functionality is present in –BM1 and –BM2 as reflected in the datasheets, JTX28. (Dkt. No. 489 at 391:16–393:22.) [redacted]. (Dkt. No. 493 at 921:5-14 (McAlexander) ("[redacted]).) Taking this evidence together, there is substantial evidence that the –BM1 and –BM2 products have the claimed "data buffers."

vii. New Trial on Infringement

Samsung argues that if the Court does not grant JMOL of non-infringement, the Court should grant a new trial because the verdict was against the great weight of the evidence and was based on Netlist's improper and highly prejudicial arguments. (Dkt. No. 561 at 36.) Samsung then repeats some of the arguments it made in connection with its grounds for JMOL detailed above.

Samsung Electronics Co., Ltd.
Ex. 1077, p. 116

(*Id.* at 37.) In response, Netlist argues no new trial is warranted because Samsung's arguments are simply rehashes of its arguments for JMOL, which should be denied. (Dkt. No. 573 at 32.) In reply, Samsung argues that Netlist failed to offer any independent argument as to why the verdict is not against the clear weight of the evidence. (Dkt. No. 577 at 11.) Samsung then contends that it is clear a new trial is warranted if JMOL is not granted. (*Id.*) Netlist does not address this topic in sur-reply. (*See* Dkt. No. 588.)

The Court finds that a new trial is not warranted. The verdict on infringement is not against the great weight of the evidence, and the Court is not persuaded that Samsung has been unduly or unfairly prejudiced by Netlist's arguments.

B. Invalidity

Samsung asserted that all asserted claims of the '339 Patent, the '918 Patent, and the '054 Patent were invalid as failing to comply with the written description requirement of 35 U.S.C. § 112 ¶ 1, but the jury disagreed. (*See* Dkt. No. 561 at 37.)

A patent satisfies the written description requirement when "the disclosure of the application relied upon reasonably conveys to those skilled in the art that the inventor had possession of the claimed subject matter as of the filing date." *Mentor Graphics Corp. v. EVE-USA, Inc.*, 851 F.3d 1275, 1276 (Fed. Cir. 2017). Written description may be supplied by any "words, structures, figures, diagrams, formulas, etc." in the patent that "set forth the claimed invention." *Blue Calypso, LLC v. Groupon, Inc.*, 815 F.3d 1331, 1346 (Fed. Cir. 2016). "Whether a patent claim is adequately supported by the written description under 35 U.S.C. § 112 is a question of fact that [is] review[ed] for substantial evidence following a jury trial." *Intellectual Ventures I LLC v. Motorola Mobility LLC*, 870 F.3d 1320, 1324 (Fed. Cir. 2017).

i. The '918 and '054 Patents

In 2008, Netlist filed a patent application for a hybrid memory module that transfers data between volatile DRAM memory and non-volatile flash memory on the same module. (Dkt. No. 561 at 38.) In December 2020 and May 2021, Netlist filed the sixth and seventh continuations of that application in an attempt to claim memory modules without any non-volatile memory at all. (*Id.*) Those continuations issued as the '918 and '054 Patents, respectively. (*Id.*)

Samsung argues that the shared specification of the '918 and '054 Patents dictates that the alleged invention is a hybrid memory module, as its expert, Mr. McAlexander, testified. (*Id.* at 38–40.) According to Samsung, a hybrid memory module is one that incorporates both volatile and non-volatile memory. (*Id.* at 40.) Samsung notes that Mr. Milton, the only named inventor to testify, admitted that the "[t]he concept [he] invented was putting that all on one module." (*Id.* (quoting Dkt. No. 481 at 244:22–24).)

Next, Samsung argues that Netlist "impermissibly broadened" the claims of the '918 and '054 Patents during prosecution, such that they no longer claim the invention: the hybrid memory module described in the specification. (*Id.*) Instead, Samsung notes that the claims recite a "new alleged invention" which Netlist has described as "intelligent on-module power management," which does not require a hybrid module capable of transferring data between volatile and non-volatile memory. (*Id.* at 40–41.) Samsung asserts that "[b]y dropping the hybrid memory structure required by the written description, Netlist improperly expanded the scope of the asserted claims to cover non-hybrid devices." (*Id.* at 41–42 (citing, among others, *Atl. Rsch. Mktg. Sys. v. Troy*, 659 F.3d 1345, 1355 (Fed. Cir. 2011); *PIN/NIP, Inc. v. Platte Chem. Co.*, 304 F.3d 1235, 1247 (Fed. Cir. 2002)).) Indeed, Samsung contends, the specification teaches against prior embodiments without non-volatile memory. (*Id.* at 42–43.)

Samsung next contends that Netlist's evidence cannot support the verdict. (*Id.* at 42–46.) It claims that Netlist's argument that the invention is an "intelligent on-module power management" without a hybrid memory system system finds no support in memory system 1010 of Figure 12, which Netlist relied on at trial because the embodiment in Figure 12 incudes non-volatile memory 1040. (*Id.* at 43.) Further, Samsung contends that Netlist's argument that Figures 15A and 15B of the specification show an embodiment without non-volatile memory is deficient because the specification explains that these figures depict only the "volatile memory subsystem" portion of a complete hybrid module that includes non-volatile memory. (*Id.* at 43–44.) Also, Samsung argues that testimony that the specification discloses an embodiment in which the non-volatile memory is a hard disk separate from the memory module cannot support the verdict because its own expert, Mr. McAlexander testified that this excerpt of the specification actually discloses a hybrid memory module with both volatile and non-volatile memory. (*Id.* at 44–45.)

Samsung then argues that "[n]one of the excerpts discussed above support[] the jury verdict" because "Netlist failed to present evidence that any of these alleged non-hybrid embodiments included all claim limitation or could otherwise be combined with other disclosures in the specification to form the claimed memory modules." (*Id.* at 45–46 (quoting *Novozymes A/S v. DuPont Nutrition Biosciences APS*, 723 F.3d 1336, 1349 (Fed. Cir. 2013) (criticizing a patentee "seek[ing] to derive written description support from an amalgam of disclosures plucked selectively")).)

Lastly, Samsung argues that Netlist's infringement read of "converter circuit" to include an LDO regulator means that the claims fail for a lack of written description because the specification only describes buck converters, boost converters, and buck-boost converters. (*Id.* 46–47.)

CONFIDENTIAL INFORMATION HAS BEEN REMOVED FROM THIS PAGE

In response, Netlist begins by arguing that the invention in the shared specification is not limited to hybrid modules. (Dkt. No. 573 at 33.) Indeed, Netlist contends, the evidence it adduced at trial including testimony from Mr. McAlexander, showed that the shared specification contemplated a memory module where non-volatile memory is located elsewhere in the system. (Dkt. No. 573 at 33–35.) Netlist also notes that Mr. Milton, a named inventor, testified that "there are several inventions embodied in that application." (*Id.* at 35–36 (quoting Dkt. No. 487 at 239:23–240:2).) Also, Netlist argues that the shared specification discloses that non-volatile memory may be located in the controller. (*Id.* at 36.) "[U]ltimately," Netlist argues, "the fact that the patents discuss hybrid modules does not change the fact that they ***also*** describe embodiments with non-hybrid modules where non-volatile memory is not on the module." (*Id.* (quoting *Allergan Sales, LLC v. Sandoz, Inc.*, 717 F. App'x 991, 995 (Fed. Cir. 2017) ("Even a single representative embodiment can support written description.")).)

Netlist goes on to argue that it did not improperly expend the scope of its claims through continuation applications, as the specification supports claims without non-volatile memory. (*Id.* at 37–38.) Netlist then contends that it was Samsung's burden to prove lack of written description by clear and convincing evidence, not Netlist's burden to present evidence that there was written description support. (*Id.* at 38–39.)

Finally, Netlist addresses Samsung' argument concerning the "converter circuit" limitation. (*Id.* at 39–41.) Netlist argues that Samsung's arguments in this regard are waived because they did not present them at trial or move on them under Rule 50(a). (*Id.* at 39–40.) Netlist then contends that even if these arguments are not waived, they fail on the merits. (*Id.* at 40–41.) Specifically, Netlist argues [redacted], and

CONFIDENTIAL INFORMATION HAS BEEN REMOVED FROM THIS PAGE

that the specification states that "various converter circuits" may be used, including "other types of converters" that differ from buck, boost, or buck-boost converters. (*Id.* at 41.)

In reply, Samsung argues that Mr. McAlexander did not admit that the specification discloses an embodiment "where non-volatile storage is not part of the module." (Dkt. No. 577 at 12–13.) Samsung also argues (again) that Figures 15A and 15B show only the volatile memory subsystem of a larger hybrid module. (*Id.* at 13.) Samsung then contends (again) that "Netlist fails to cite evidence that any of the alleged non-hybrid embodiments in the specification includes a disclosure of all claim limitations or could be combined with other disclosures to form the claimed memory modules." (*Id.* at 13–14.)

Concerning the "converter circuit" limitation, Samsung summarily asserts that "Netlist does not identify any basis on which a person of skill in the art would have that understanding of the specification." (*Id.* at 14.) Samsung also asserts that it raised the issue at trial and in its Rule 50(a) motion. (*Id.*)

In sur-reply, Netlist argues that Samsung ignores that the specification expressly discloses a non-hybrid memory module where the non-volatile memory is located elsewhere in the system and ignores testimony to the same effect. (Dkt. No. 588 at 11–12.) Netlist then asserts that [redacted]. (*Id.* at 12.) Netlist then urges that Samsung's contentions that the specification does not disclose such an embodiment are wrong. (*Id.* at 12–13.) Netlist next states (again) that it was Samsung's burden to prove lack of written description, not Netlist's. (*Id.* at 13.)

Finally, Netlist addresses Samsung's arguments regarding the "converter circuit" limitation. (*Id.* at 13–14.) Netlist again asserts that Samsung has waived the argument, and again

Samsung Electronics Co., Ltd.
Ex. 1077, p. 121

urges that it has shown substantial evidence that a person of skill in the art would understand an LDO to be a converter circuit. (*Id.* at 14.)

The Court is not persuaded by Samsung's arguments. It was Samsung's burden to prove "by clear and convincing evidence" that the patents were invalid. *Core Wireless*, 880 F.3d at 1364. *See also*, 35 U.S.C. § 282 ("A patent shall be presumed valid. … The burden of establishing invalidity of a patent or any claim thereof shall rest on the party asserting such invalidity."). The jury found that Defendants did not meet their substantial burden, and the Court sees no basis to overturn their determination.

Notably, in *Core Wireless*, the Federal Circuit held that "[b]ecause the burden rests with the alleged infringer to present clear and convincing evidence supporting a finding of invalidity, granting judgment as a matter of law for the party carrying the burden of proof is generally 'reserved for extreme cases.'" 880 F.3d at 1364. An "extreme case[]" would be when something drastic occurs, "such as when the opposing party's witness makes a key admission." *Id.* No such admissions were made here.

Samsung's arguments fail on the merits. Samsung contends that the specification does not disclose an embodiment where non-volatile memory is not on the memory module, but Netlist presented substantial evidence otherwise. First, the shared specification discloses such an embodiment:

> For example, in one embodiment, the host system (e.g., a disk controller) writes data to the volatile memory subsystem **1030** which then writes the data to non-volatile storage ***which is not part of the memory system*** **1010,** such as, for example, a hard disk. The disk controller may wait for an acknowledgment signal from the memory system **1010** indicating that the data has been written to the hard disk or is otherwise secure. The memory system **1010** of certain embodiments can decrease delays in the system operation by indicating that the data has been written to the hard disk before it has actually done so. In certain embodiments, the memory system **1010** will still be able to recover the data efficiently in the event of a power outage because of the backup and restore capabilities described herein. In certain

CONFIDENTIAL INFORMATION HAS BEEN REMOVED FROM THIS PAGE

> other embodiments, the memory system **1010** may be operated as a write-through cache or as some other type of cache.

'918 Patent at 27:42–58 (emphasis supplied). Samsung contends that this portion of the specification omits context indicating that memory system 1010 is a hybrid module containing both volatile memory subsystem 1030 and non-volatile memory subsystem 1040 (*see* Dkt. No. 577 at 12 (citing '918 Patent at 27:16–40)). However, this overlooks the oft repeated prefatory phrases like "in certain embodiments" or "in one embodiment." *See* '918 Patent at 27:14–58. The passage quoted above is described as a standalone embodiment, meaning that the specification expressly contemplates a memory module that does not contain non-volatile storage.

Witnesses for both Netlist and Samsung corroborate this. Named inventor Mr. Milton testified that the specification teaches an embodiment where the "non-volatile memory or flash is not on the module." (Dkt. No. 481 at 197:10–14.) Indeed, Mr. Milton testified that "there are several inventions embodied in that application." (Dkt. No. 487 at 239:23–240:2.) Netlist's expert, Dr. Mangione-Smith testified that "in this embodiment of the patent, the non-volatile memory is not on that memory module." (Dkt. No. 489 at 469:22–13.) Samsung's expert, Mr. McAlexander, gave similar testimony. (Dkt. No. 493 at 962:24–963:4 [redacted]

Figures 15A and 15B support neither Netlist's nor Samsung's theory. The specification makes clear that these figures illustrate a "volatile memory subsystem," but nothing in the specification either requires that the subsystems depicted be part of a memory module containing non-volatile memory or forbids the subsystems depicted from being part of a memory module that does not contain non-volatile memory. *See* '918 Patent at 23:41–64. This, however, is without

Samsung Electronics Co., Ltd.
Ex. 1077, p. 123

CONFIDENTIAL INFORMATION HAS BEEN REMOVED FROM THIS PAGE

import in light of the substantial evidence described above that the specification discloses an embodiment of a memory module that does not contain non-volatile memory.

Samsung's arguments regarding the "converter circuit" limitation also fall short. As an initial matter, however, Samsung did not waive its arguments. Contrary to Netlist's suggestions, this argument was sufficiently presented at the Rule 50(a) hearing. Counsel for Samsung argued that

> Doctor Mangione-Smith also admitted that the written description did not support the full scope of the claims. For example, Doctor Mangione-Smith admitted that there was no written description support for that LDO regulator that the parties are -- are disputing that he's now claiming that is an issue for infringement.

(Dkt. No. 495 at 1251:13–18.) Though Samsung did not reference the "converter circuit" claim, such is not fatal. In fact, the "LDO regulator" was consistently brought up in connection with the "converter circuit" term such that Netlist was on fair notice. Further, this argument was presented to the jury, though not clearly so. (*See* Dkt. No. 493 at 876:24–877:25 (McAlexander) [redacted]

CONFIDENTIAL INFORMATION HAS BEEN REMOVED FROM THIS PAGE

[redacted]

Ultimately, the Court does not conclude that Netlist was not on fair notice of this argument.

Nonetheless, Samsung's arguments fail on the merits. The jury heard substantial evidence that the written description for "converter circuit" supports an LDO regulator and that an LDO regulator falls within the plain meaning of "converter circuit." [redacted] (Dkt. No. 493 at 953:2–8.) [redacted] (*See* Dkt. No. 489 at 349:24–350:10, 372:18–23.) Moreover, the specification does not limit "converter circuit" to only buck converters, boost converters, and buck-boost converters as Samsung contends. (Dkt. No. 561 at 47.) The specification contemplates that "various converter circuits" may be used and that this may include "other types of converters." '918 Patent at 29:24, 29:59–60.

Samsung's argument that "Netlist failed to present evidence that any of these alleged non-hybrid embodiments included all claim limitations or could otherwise be combined with other disclosures in the specification to form the claimed memory modules" fails. (*See* Dkt. No. 561 at 45.) First, this turns the burden of proof on its head. 35 U.S.C. § 282. Second, Samsung has only argued that Netlist has failed to show (1) that the specification contemplates a memory module without non-volatile memory and (2) that the specification supports a reading of "converter circuit" that includes an LDO, but, as described above, both these arguments fail—substantial evidence was presented on both these grounds. Samsung summarily asserts that "Netlist has offered no evidence that a person of ordinary skill in the art … would know to combine the alleged non-hybrid embodiments with other disclosures in the manner required by the asserted claims," but this

Samsung Electronics Co., Ltd.
Ex. 1077, p. 125

again ignores that it is Samsung's burden to show invalidity by clear and convincing evidence. (Dkt. No. 561 at 46.)

In short, Samsung has failed to persuade the Court why or how it might have carried its heavy burden.

ii. The '339 Patent

Samsung argues that Netlist's infringement read of the "drive" limitation of the '339 Patent stretches the claims beyond written description support. (Dkt. No. 561 at 47–48.) Dr. Mangione-Smith asserted that the claims are met by "a single byte-wise data path in the byte-wise buffer," without a "fork-in-the-road." (*Id.* at 47 (quoting Dkt. No. 489 at 448:25–449:4).) However, in the Court's Claim Construction Order, the Court recognized that "the sole embodiment describing path selection during a write operation disables one path within the buffers when the other path is enabled." (*Id.* at 47–48 (quoting Dkt. No. 114 at 10).) Samsung asserts that Netlist did not offer any support for its assert that the '339 Patent discloses embodiments without a "fork-in-the-road." (*Id.* at 48.)

Netlist argues in response that the '339 Patent teaches that a single path may be sufficient with two rank memory modules. (Dkt. No. 573 at 41.) Netlist contends that the jury chose to reject Samsung's arguments and accept Netlist's and that the Court should not overturn the jury's conclusion. (*Id.*)

In reply, Samsung argues that Netlist again ignores the Court's Claim Construction Order. (Dkt. No. 577 at 14.) Samsung also contends that the portions of the specification and the testimony cited by Netlist in its response do not support its assertions. (*Id.* at 14–15.)

In sur-reply, Netlist argues that the specification teaches embodiments where a single path is sufficient. (Dkt. No. 588 at 14.) Netlist then argues that the Court's Claim Construction Order is

not to the contrary and that the statement referenced by Samsung concerns a different embodiment. (*Id.* at 14–15.)

The Court is not persuaded by Samsung's arguments. First, Netlist's infringement read does not depend on a single path—indeed, the Court's construction requires a "fork-in-the-path." (Dkt. No. 114 at 10.) In any event, Samsung does not dispute that the specification discusses embodiments containing a single path. '339 Patent at 4:48–64, 5:4–16. Samsung complains that these provisions come from the section discussing prior art, but the background can supply written description support just like any other part of the specification. *See Storz Instrument Co. v. Alcon Labs., Inc*., 1998 WL 50947, at *4–6 (Fed. Cir. 1998) (finding statement in "Background and Summary of the Invention" supported written description).[1]

iii. New Trial

Samsung argues that the Court should grant a new trial if it does not grant Samsung JMOL on written description because the jury's verdict stands against the "clear weight of the evidence." (Dkt. No. 561 at 48.) Samsung also argues that the Court's failure to instruct the jury that a "broad claim is invalid when the specification clearly indicates that the invention is of a much narrower scope" was prejudicial and that this necessitates a new trial on written description. (*Id.* (citing *Cooper Cameron Corp. v. Kvaerner Oilfield Prods.*, 291 F.3d 1317, 1322 (Fed. Cir. 2002)).)

Netlist argues in response that there is no basis for a new trial on written description because the jury's decision was supported by substantial evidence. (Dkt. No. 573 at 42.) Netlist then contends that a new trial based on improper jury instructions is only proper if "substantial and

[1] *Storz* is unpublished. 1998 WL 50947. Federal Cirucit Rule 32.1(d) provides that "[t]he court may refer to a nonprecedential or unpublished disposition in an opinion or order and may look to a nonprecedential or unpublished disposition for guidance or persuasive reasoning but will not give one of its own nonprecedential dispositions the effect of binding precedent." Since the Federal Circuit would consider *Storz* persuasive but not binding, the Court finds it persuasive in this context.

ineradicable doubt whether the jury has been properly guided in its deliberations." (*Id.* (quoting *OPTi, Inc. v. VIA Techs., Inc.*, 65 F.Supp.3d 465, 482 (E.D. Tex. 2014)).) Netlist asserts that Samsung has identified "no such doubt." (*Id.*) Further, Netlist urges that the Court gave the jury a complete instruction on written description. (*Id.*) Netlist then argues that Samsung has not identified any error in the Court's instructions to the jury, much less prejudicial error. (*Id.*) Indeed, Netlist argues, the Court's instruction regarding written description comports with the case cited by Samsung, *Cooper Cameron*, 291 F.3d 1317. (Dkt. No. 573 at 42–43.)

In reply, Samsung argues that Netlist uses the wrong standard for setting aside a verdict based on jury instructions. (Dkt. No. 577 at 15.) Samsung also argues that the Court's instructions failed to address Samsung's core argument in this case—*i.e.*, the asserted claims omit essential elements described in the specification. (*Id.*)

In sur-reply, Netlist argues that Samsung's argument that the Court did not instruct the jury about a theory that "the asserted claims omit essential elements described in the specification" is waived since it was raised for the first time in reply. (Dkt. No. 588 at 15.) Netlist also contends that it is wrong because it is contrary to Federal Circuit law which holds that a claim need not include all features of an invention recited in the specification. (*Id.* (citing, among others, *AllVoice Computing PLC v. Nuance Commc'ns, Inc.*, 504 F.3d 1236, 1248 (Fed. Cir. 2007) ("each claim need not include every feature of an invention")).)

The Court finds that a new trial is not warranted. The verdict is supported by substantial evidence and is not against the great weight of the evidence. The Court also finds that Samsung has failed to show that the jury instruction was erroneous, regardless of the standard used. Samsung has not argued that the instruction caused any specific prejudice, or that the outcome would have been different had the Court used Samsung's preferred instructions. Further, the Court's

instruction—the "specification must describe the full scope of the claimed invention, including each element therefore -- thereof, either expressly or inherently"—captures Samsung's proposed instruction—a "broad claim is invalid when the specification clearly indicates that the invention is of a much narrower scope." (Dkt. No. 493 at 1319:2–5; Dkt. No. 561 at 48.) Samsung also argues that the failure to include its instruction meant that the Court's instructions failed to address Samsung's invalidity argument that the asserted claims omit essential elements, but neither does Samsung's proposed instruction does not address this argument. (*See* Dkt. No. 577.) Samsung's proposed instruction is just a re-encapsulation of the idea communicated by the Court's instruction—that the specification must describe the full scope of the claims—and has no bearing on whether the claims must claim an aspect of an invention described in the specification. In any event, a claim need not include all features of an invention recited in the specification. *AllVoice Computing PLC v. Nuance Commc'ns, Inc.*, 504 F.3d 1236, 1248 (Fed. Cir. 2007) ("each claim need not include every feature of an invention"); *Golight, Inc. v. Wal-Mart Stores, Inc.*, 355 F.3d 1327, 1331 (Fed. Cir. 2004) ("patentees were not required to include within each of their claims all of these advantages or features described as significant or important in the written description"). Accordingly, a new trial on invalidity is not warranted.

C. Damages

The jury awarded Netlist $303.15MM as a reasonable royalty in the form of a running royalty. (*See* Dkt. No. 480 at 7.) Samsung argues that the Court should grant JMOL of no damages, grant JMOL of damages no more than $19.3MM, grant a new trial on damages, or condition a new trial on damages on remittitur. (Dkt. No. 561 at 48–63.)

CONFIDENTIAL INFORMATION HAS BEEN REMOVED FROM THIS PAGE

i. Award of 100% of the Revenue Conferred by the Asserted Patents and Apportionment

Netlist's damages expert, Mr. Kennedy, admitted that his damages calculations award to Netlist 100% of the revenue associated with the technology at issue. (Dkt. No. 491 at 775:1–7 [redacted] ")

Samsung contends that no reasonable jury could find that Samsung would have agreed to Netlist's proposed royalty of 100% of the revenue associated with the technology at issue. (Dkt. No. 561 at 50.) Samsung argues that courts have consistently held that awarding a patentee 100% of the alleged infringer's profits is "unreasonable" and "insupportable." (*Id.* at 50–51 (citing *Contour IP Holding, LLC v. GoPro, Inc.*, 2020 WL 5106845, at *14 (N.D. Cal. Aug. 31, 2020), *Looksmart Grp., Inc. v. Microsoft Corp.*, 2019 WL 4009263, at *3 (N.D. Cal. Aug. 5, 2019), *Nordock, Inc. v. Systems, Inc.*, 2013 WL 989864, at *8 (E.D. Wis. Mar. 13, 2013), *Zegers v. Zegers, Inc.*, 458 F.2d 726, 728 n.8 (7th Cir. 1972)).) Based on this, Samsung argues that an award of 100% of the revenues cannot be justified. (*Id.* at 51–52.) Samsung also contends that Mr. Kennedy's "damages theory" is really a disgorgement theory in disguise, which Congress eliminated long ago. (*Id.* at 52 (citing *Aro Mfg. Co. v. Convertible Top Replacement Co.*, 377 U.S. 476, 505 (1964) (holding that Congress amended Section 284 "precisely to eliminate the recovery of [the infringer's] profits . . . and allow recovery of damages only")).) Last, Samsung argues that Mr. Kennedy's is a more improper version of the 25% and 50% rules of thumb that have been rejected by the Federal Circuit. (*Id.* at 52–53 (citing *See Labyrinth Optical Techs. LLC v. Alcatel-Lucent USA, Inc.*, 2015 WL 12696081, at *4 (C.D. Cal. Sept. 2, 2015)).)

In response, Netlist first argues that "[a] jury's decision with respect to an award of damages must be upheld unless the amount is grossly excessive or monstrous, clearly not supported by the evidence, or based only on speculation or guesswork." (Dkt. No. 573 at 43 (quoting *Lucent Techs., Inc. v. Gateway, Inc.*, 580 F.3d 1301, 1310 (Fed. Cir. 2009)).) Netlist then argues that the facts in evidence before the jury constitute substantial evidence that supports the jury's award. (*Id.* at 43–44.) Netlist contends that Mr. Kennedy analyzed the *Georgia-Pacific* factors, relying primarily on the value that Samsung obtained from the technology. (*Id.* at 44.) Netlist then walks through the analysis Mr. Kennedy performed for each class of products—DDR4, DDR5, and HBM. (*Id.* at 44–46.) Netlist argues that these analyses are precisely what the Federal Circuit requires: to identify the incremental benefits of the patent, and to "estimate the value of the benefit provided by the infringed features by comparing the accused product to non-infringing alternatives." (*Id.* at 46–47 (quoting *Apple Inc. v. Motorola, Inc.*, 757 F.3d 1286, 1315 (Fed. Cir. 2014), *overruled on other grounds by Williamson v. Citrix Online, LLC*, 792 F.3d 1339 (Fed. Cir. 2015)).)

Netlist also argues that it was proper for Mr. Kennedy to allocate as damages 100% of revenue associated with the technological benefit conferred upon Netlist. (*Id.* at 47–52.) Netlist also urges that Samsung's JMOL arguments are simply rehashes of its *Daubert* motion against Kennedy, which is improper at the Rule 50(b) stage. (*Id.* at 47–48 (citing *Versata Software Inc. v. SAP Am. Inc.*, 717 F.3d 1255, 1264 (Fed. Cir. 2013)).) Netlist then contends that "[t]he Federal Circuit has repeatedly held that a patentee is entitled to the entire apportioned value between the next best NIA and the infringing instrumentality." (*Id.* at 48 (citing *Prism Techs. LLC v. Sprint Spectrum L.P.*, 849 F.3d 1360, 1376 (Fed. Cir. 2017); *Powell v. Home Depot U.S.A., Inc.*, 663 F.3d 1221, 1240 (Fed. Cir. 2011)).) Netlist asserts that Mr. Kennedy properly considered the *Georgia-*

CONFIDENTIAL INFORMATION HAS BEEN REMOVED FROM THIS PAGE

Pacific factors and properly concluded [redacted] *Id.* at 48–49.) Netlist further argues that a case cited by Samsung, *Labyrinth Optical*, is inapposite to the facts at hand. (*Id.* at 49–50.) Netlist then argues that Mr. Kennedy established a quantitative connection between the asserted facts and his ultimate conclusion by apportioning to the technical benefit of the patents-in-suit. (*Id.* at 50–51.) Also, Netlist contends that "Mr. Kennedy also did not suggest that Netlist should receive 100% of the profit associated with the accused products, but rather that Samsung would not retain any of the already fully-apportioned value of the patented technology." (*Id.* at 51.) Next, Netlist urges that Samsung's emphasis on the revenue rather than the profits was addressed at the *Daubert* stage. (*Id.* at 51–52.) Finally, Netlist argues that the jury's award was only based on 75% of the apportioned value of the patents, not 100%, since the jury awarded 75% of what Netlist asked for. (*Id.* at 52.)

In reply, Samsung first argues that its arguments (here and below) are challenges to the sufficiency of the evidence, not *Daubert* challenges. (Dkt. No. 577 at 15–16.) Samsung then contends that the Federal Circuit has recently distinguished *Versata*, cited by Netlist, on the basis that the defendant in *Versata* "argue[d] that the district court should not have admitted SSC's expert testimony on damages," not "that the jury's verdict is not supported by substantial evidence because [the plaintiff's] expert testimony on damages" was flawed. (*Id.* at 15–16 (quoting *Enplas Display Device Corp. v. Seoul Semiconductor Co.*, 909 F.3d 398, 411 n.2 (Fed. Cir. 2018)).) Samsung then contends that even where a *Daubert* admissibility challenge has been waived, Rule 50(b) still requires review for sufficiency of the evidence. (*Id.* at 16 (citing *KAIST IP US LLC v. Samsung Elecs. Co.*, 439 F. Supp. 3d 860, 889 (E.D. Tex. 2020)).) Samsung asserts that there is no way for the Court to evaluate the sufficiency of the evidence at trial prior to the actual trial itself and thus denial of a *Daubert* does not require denying JMOL. (*Id.*)

Samsung argues that *Prism* and *Powell* are distinguishable, and that Netlist has not distinguished the cases cited by Samsung. (*Id.* at 16–17 (citing *Prism*, 849 F.3d 1360; *Powell*, 663 F.3d 1221).) Samsung again urges that *Labyrinth* is applicable and dictates that Netlist's damages theory fails as a matter of law. (*Id.* at 17 (citing *Labyrinth*, 2015 WL 12696081).) Samsung argues that Netlist's arguments about the "quantitative connection" for Mr. Kennedy's 100% award are insufficient because "these conclusory expert assertions are insufficient to show that Samsung, in a hypothetical negotiation, would agree to give Netlist 100% of the revenue corresponding to the patented technology." (*Id.* at 17–18.) Samsung then argues that it is irrelevant to their "rule-of-thumb" argument that the jury only awarded 75% of what Netlist requested. (*Id.* at 18.)

In sur-reply, Netlist asserts that Samsung's arguments are premised on critiques of Mr. Kennedy's methodology and are therefore *Daubert* arguments. (Dkt. No. 588 at 16.) Netlist then argues that Samsung has failed to establish that evidence presented on damages is so one-sided as to justify JMOL. (*Id.* at 16–17 (citing *Core Wireless*, 880 F.3d at 1361).)

Netlist also argues that Samsung has not engaged with the facts Netlist cited in its opposition brief. (*Id.* at 17.) Further, Netlist contends, Samsung cites no authority requiring Netlist to show that "anyone has ever agreed" to the arrangement Mr. Kennedy proposed. (*Id.* (citing Dkt. No. 577 at 17).) Netlist then urges that *Labyrinth* is not applicable. (*Id.* at 17–18 (citing 2015 WL 12696081).)

The Court is not persuaded by Samsung's arguments. There is no doubt that Mr. Kennedy opined that a proper award to Netlist would cover 100% of the revenue associated with the technology at issue. (Dkt. No. 491 at 775:1–7.) While Samsung contends that this is impermissible and that no reasonable jury could find that Samsung would agree to award Netlist 100% of the revenue, Samsung is wrong.

Georgia-Pacific factors 10 and 11 direct that the value of the patented technology should be considered. *Georgia-Pac. Corp. v. U.S. Plywood Corp.*, 318 F. Supp. 1116 (S.D.N.Y. 1970), *modified sub nom. Georgia-Pac. Corp. v. U.S. Plywood-Champion Papers, Inc.*, 446 F.2d 295 (2d Cir. 1971) ("10. The nature of the patented invention; the character of the commercial embodiment of it as owned and produced by the licensor; and the benefits to those who have used the invention. 11. The extent to which the infringer has made use of the invention; and any evidence probative of the value of that use.") No limit is placed on the amount of value a patentee may claim.

Caselaw supports this. The Federal Circuit in *Prism* awarded the full amount of cost-savings to plaintiff, which is another metric for value. 849 F.3d at 1376. In other words, the Federal Circuit has approved the patentee receiving the entirety of the value of the patent in a hypothetical negotiation. *See id.* The Federal Circuit did much the same in *Powell* by allowing the patentee to keep an award based on the savings a patent would provide. 663 F.3d at 1240.

Moreover, there is substantial evidence that Samsung would have agreed to let Netlist keep 100% of the value (which is really a better word than "revenue" in this context) of the patented invention. Mr. Kennedy opined that Netlist "never licensed its patents outside of a strategic agreement," which is significant because "in the hypothetical negotiation there is no strategic agreement. Samsung needs a license to the patents, and they are not giving anything in return for that except royalties," and further, "the negotiators would know that Netlist has never entered into a bare patent license agreement … so … those first two strategic relationships that they entered into should not be used to determine … the royalty." (Dkt. No. 491 at 686:24–687:3, 692:11–19.) Mr. Kennedy also pointed to Samsung documents showing Samsung was "behind [its] competitor's schedule already with DDR5," and thus the technical benefit provided by Netlist's

patents would be "particularly important to Samsung." (*Id.* at 694:10–15.) It would be important to Samsung to stay competitive, Mr. Kennedy opined, because it "control[s] 95 percent of the market place" with Micron and SK hynix and so "Samsung would want to be very competitive, not to lose market share []. There's two other big players out there that would willingly take those sales if [Samsung] couldn't offer that technology." (*Id.* at 706:4–14.) Mr. Kennedy also opined, relying on Samsung testimony, that Samsung had previously been cut out of the marketplace by SK hynix when it launched a product before Samsung, meaning Samsung had strong reason to want to remain competitive. (*Id.* at 706:18–22.) In sum, Netlist presented substantial evidence for the reasons why Samsung would have agreed at a hypothetical negotiation to pay Netlist 100% of the benefit of the patented invention.

Samsung cites a litany of cases, asserting they support its position, but they do not. (*See* Dkt. No. 561 at 50–51 (citing *Looksmart*, 2019 WL 4009263, at *3; *Nordock*, 2013 WL 989864, at *8; *Zegers*, 458 F.2d at 728 n.8.) First, none of the cases cited by Samsung in this regard are binding. Secondly, none of these cases hold that a patentee may not seek 100% of the value of its invention. Third, none of these cases touch on a situation where there was substantial evidence that an infringer has strong reasons to agree to give the patentee 100% of the value, as exists here.

Samsung also cites *Labyrinth*. (*See id.* at 51 (citing 2015 WL 12696081, at *4).) The District Court for the Central District of California held that:

> There is no evidence that Defendant or anyone else has ever agreed to pay 100% or even 25% royalty, nor evidence that Plaintiff has ever received such a royalty. Without such evidence, Woods' opinion is inadmissible because it is untethered from the patented technology at issue and the many licenses thereto and, as such, is arbitrary and speculative.

Id. (quotations and brackets omitted). Relying on this Samsung argues that "[b]ecause Netlist did not present evidence that anyone has ever agreed to such an arrangement, Mr. Kennedy's bare assertions do not constitute substantial evidence." (Dkt. No. 577 at 17.)

These arguments are not persuasive. As above, *Labyrinth* is not controlling. Moreover, if Samsung were to have its way, it would inject an improper limitation on the hypothetical negotiation. Neither Samsung nor *Labyrinth* cite controlling law for why a patentee must show that the infringer "agreed to such an arrangement" for that "arrangement" to be valid as a matter of law. Further, as above, *Labyrinth* does not touch on a situation where there was substantial evidence that an infringer has strong reasons to agree to give the patentee 100% of the value.

Samsung's remaining arguments are also not persuasive. Seeking 100% of the value of the patented invention is not disgorgement, and has been allowed by the Federal Circuit, as discussed above. Also, seeking 100% of the value of the patented invention is not an improper rule of thumb when, as discussed above, there is substantial evidence that the infringer would have agreed to award 100% of the value of the patented invention.

Two other points should be made. First, Netlist's argument that the jury only awarded 75% of what Netlist requested is irrelevant and therefore unavailing. (Dkt. No. 573 at 52.) Second, the Court finds that Samsung's arguments are attacks on Mr. Kennedy's methodologies because they go directly to the sufficiency of his theories and have been either waived (if Samsung failed to raise them at *Daubert*) or are improperly re-raised again at the Rule 50(b) stage. *KAIST*, 439 F. Supp. 3d at 888.

ii. Apportionment

Samsung argues that JMOL is appropriate because Mr. Kennedy failed to apportion. (Dkt. No. 573 at 53–55.) Specifically, Samsung argues that Mr. Kennedy was required to establish that the differences between the accused products and the alternatives are limited to incremental

benefits provided by the patents. It says that Netlist failed to do so. (*Id.* at 53 (citing *Finjan, Inc. v. Blue Coat Sys., Inc.*, 879 F.3d 1299, 1310 (Fed. Cir. 2018).) Mr. Kennedy claimed, based on a Samsung document, that DDR5 could offer an increase of 30% power efficiency over the prior generation DDR4 products, but, Samsung argues, he did not offer any analysis connecting this improvement to the claims of the '918 and '054 Patents. (*Id.* at 54.) Next, Samsung argues that Netlist offered no evidence or analysis connecting the improvement associated with the '339 Patent to the two-DIMMs-per-channel configuration Netlist claims the '339 Patent enables. (*Id.*) Regarding HBM products, Samsung argues that Netlist's technical expert, Dr. Brogioli, opined that the '060 and '160 Patents are necessary to achieve 8-high stacks, but Samsung argues (1) that Dr. Brogioli failed to address other prior art with 8-high stacks and (2) that Dr. Brogioli and Mr. Kennedy never provided specific evidence to demonstrate why the '060 and '160 Patents were necessary for 8-high stacks. (*Id.*) Finally, Samsung argues that Netlist failed to account for any benefits already provided by the prior art. (*Id.* at 55.)

In response, Netlist notes that this same argument was rejected at the *Daubert* stage. (Dkt. No. 573 at 53.) Further, Netlist contends that one of its technical experts, Dr. Mangione-Smith, connected the claims of the '918 and '054 Patents to the 30% efficiency gain and the claims of the '339 Patent to the two-DIMMs-per-channel configuration. (*Id.* at 53–55.) Netlist also contends that Dr. Brogioli adequately explained the drawbacks of the prior art with 8-high stacks and explained why the '060 and '160 patents are necessary to achieve 8-high stacks. (*Id.* at 55.) Netlist also argues that Samsung cites no law requiring Netlist to take account of the benefits provided by the prior art but that the jury was presented with substantial evidence on the contributions of the prior art. (*Id.* at 55–56.)

In reply, Samsung argues that Netlist has "cherry-picked" testimony from its technical experts and ignores that Mr. Kennedy failed to abide by the principle that damages must be apportioned according to the incremental benefit conferred by the patent over conventional elements. (Dkt. No. 577 at 20–21 (citing *Realtime Data, LLC v. Actian Corp.*, 2017 WL 11661896, at *3 (E.D. Tex. Mar. 24, 2017)).) Further, Samsung contends, none of the evidence cited by Netlist "establishes that the differences between the products were solely attributable to the asserted patents rather than, for example, conventional elements that existed in the prior art." (*Id.* at 21.)

In sur-reply, Netlist points out that Kennedy was permitted to rely upon the testimony of other technical experts. (Dkt. No. 588 at 20–21.) Netlist also argues that *Realtime* is distinguishable. (*Id.* at 21 (citing 2017 WL 11661896, at *3).) Netlist reasserts that the evidence it cited was sufficient to show apportionment. (*Id.* at 21–22.)

The Court is not persuaded by Samsung's arguments. For each product class (DDR4, DDR5, and HBM), Mr. Kennedy included the apportionment to his damages number. As discussed above, Mr. Kennedy's approach was to award to Netlist the difference in value between the accused product and the next-best alternative. (*See* Dkt. No. 491 at 775:1–7.) In turn, Netlist's next-best alternatives were based on what the accused product would be if the features enabled by the patents (and only the features enabled by the patents) were removed. Accordingly, Mr. Kennedy was directly measuring the value of the patented features by comparing the accused products to the next-best alternative (which are just the accused products *sans* features enabled by the patents).

For DDR4 products, Netlist asserted that the '339 Patent provides the ability to use two DIMMs per channel. Samsung argues that "Netlist offered no evidence or analysis connecting the purported improvement … to the specific type of buffer claimed in the '339 [P]atent." (Dkt. No. 561 at 54.) However, Dr. Mangione-Smith testified that "the capacity with the '339 invention is

CONFIDENTIAL INFORMATION HAS BEEN REMOVED FROM THIS PAGE

doubled" and that Samsung did not identify "any alternative technology that it could have used in place of the '339 inventions." (Dkt. No. 489 at 400:8–23.) This is sufficient and constitutes substantial evidence.

For DDR5 products, Netlist offered evidence that the '918 and '054 Patents enabled a 30% improvement in power efficiency based on Samsung documents. (*See* Dkt. No. 489 at 381:8–383:12 (Mangione-Smith).) Samsung argues that "Netlist offered no evidence or analysis connecting the purported 30% improvement solely to the claims of the '918 and '054 [P]atents." (Dkt. No. 561 at 54.) However, Dr. Mangione-Smith testified that the '918 and '054 Patents "address … challenges with power management" with "on-module power management" and that as a result Samsung was "able to achieve a 30 percent improvement in power efficiency." (Dkt. No. 489 at 329:18–330:10.) This is substantial evidence connecting the improvement to the claims of the '918 and '054 Patents.

For HBM products, Netlist offered evidence through its other technical expert, Dr. Brogioli, [redacted]. (*See id.* at 540:18–541:18.) Samsung argues that Netlist "failed to address other prior art with 8-high stacks, and neither he nor Mr. Kennedy provided specific evidence to demonstrate why the later '060 and '160 [P]atents were supposedly necessary for 8-high stacks." (Dkt. No. 561 at 54.) Samsung does not articulate why Netlist would need to address the prior art. Nonetheless Netlist did address it. [redacted] (Dkt. No. 489 at 540:18–541:3.) [redacted]

CONFIDENTIAL INFORMATION HAS BEEN REMOVED FROM THIS PAGE

[redacted] (*Id*. at 541:10–18.) Accordingly, in the Court's view Netlist adequately addressed the 8-high stacks in the prior art.

Samsung also argues that "Netlist failed to present any evidence accounting for the benefits already provided by the prior art," but it cites no caselaw requiring Netlist to put forth such evidence. (*See* Dkt. No. 561 at 55.) The Court does not find this argument persuasive.

iii. Samsung and SK hynix Licenses

Samsung argues that Mr. Kennedy failed to take account of two highly comparable Samsung and SK hynix licenses, and since comparable licenses must be considered, Mr. Kennedy's damages analysis fails as a matter of law. (Dkt. No. 561 at 55–56.) Also, Samsung argues, Mr. Kennedy relied on a noncomparable license, the Rambus license. (*Id.* at 56–57.) As a result, Samsung argues, Mr. Kennedy's analysis fails as a matter of law. (*See id.* at 57–58.)

In response, Netlist argues that these arguments have been waived because they were presented at the *Daubert* stage (where they failed), and they substantively constitute direct challenges to Mr. Kennedy's methodology, which are inappropriate at the 50(b) stage. (Dkt. No. 573 at 57–58 (*KAIST*, 439 F. Supp. 3d at 888).) Netlist contends that Mr. Kennedy adequately explained why the Samsung and SK hynix licenses are not comparable, and Mr. Meyer, Samsung's damages expert, conceded some of these same facts. (*Id.* at 58–60.) Finally, [redacted] (*Id.* at 61 (quoting Dkt. No. 491 at 742:17–22).)

In reply, Samsung again argues that the Samsung and SK hynix licenses are comparable and that the jury awarded an award many multiples higher than any rate supported by those licenses. (Dkt. No. 577 at 18–19.)

CONFIDENTIAL INFORMATION HAS BEEN REMOVED FROM THIS PAGE

In sur-reply, Netlist again argues that it presented sufficient evidence that the Samsung license and the SK hynix agreement are not comparable. (Dkt. No. 588 at 18–19.)

The Court is not persuaded by Samsung's arguments. Whether a license is properly considered to be comparable or not is a classic challenge to expert's methodology. All of Samsung's arguments are therefore waived. *KAIST*, 439 F. Supp. 3d at 888. Nonetheless, Netlist presented substantial evidence through Mr. Kennedy regarding the issue of comparability as to the Samsung, SK hynix, and Rambus agreements. (*See, e.g.*, Dkt. No. 491 at 687:6–8, 689:19–24, 690:23–25, 691:1–23, 693:2–4.) [redacted]. (*Id.* at 742:10–22.) [redacted]. (*See, e.g.*, Dkt. No. 494 at 1191:25–1192:17, 1199:11, 1204:1–6, 1204:23–25.) On balance, this constitutes substantial evidence on the issue of comparability and non-comparability.

A verdict is not unreasonable simply because it is higher than the license agreements in the record. *LaserDynamics v. Quanta Comput.*, 694 F.3d 51, 81 (Fed. Cir. 2012) ("we do not hold that [a party's] past licenses create an absolute ceiling on the amount of damages to which it may be entitled"). Further, whether or not an agreement is comparable or not is a factually intensive inquiry. Both Netlist and Samsung thoroughly presented their versions of the facts to the jury. The jury clearly reached its conclusion by weighing these competing factual assertions. As such, the Court sees no reason to overturn the jury's ultimate determination of the facts in this regard.

iv. Non-Infringing Alternatives

Netlist calculated its damages as the difference between a "non-infringing alternative and the price of the product." (*See* Dkt. No. 491 at 775:1–7 (Kennedy).)

Samsung Electronics Co., Ltd.
Ex. 1077, p. 141

Samsung argues that Netlist failed to offer evidence that the alternatives it relied upon for the DDR4 and DDR5 products were both non-infringing and available. (Dkt. No. 561 at 58.) Specifically, Samsung argues that Netlist did not show that the alternative to the DDR4 products was non-infringing, that the alternative to the DDR5 products was available or commercially acceptable. (*Id.* at 58–60.)

In response, Netlist argues that Samsung has misinterpreted the law. (Dkt. No. 573 at 62.) Specifically, Netlist contends it only has the burden to show that there was an available next-best alternative to which it compared its claimed technology, not an acceptable non-infringing alternative. (*Id.* (citing *Salazar v. HTC Corp.*, No. 2:16-CV-01096-JRG-RSP, 2018 WL 2033709, at *3 (E.D. Tex. Mar. 28, 2018)).) Nonetheless, Netlist asserts that it did provide substantial evidence of available non-infringing alternatives as to the accused products. (*Id.* at 62–63.) Netlist also argues that the jury need only find the alternatives potentially available, not immediately available. (*Id.* at 64 (citing *Mars, Inc. v. Coin Acceptors, Inc.*, 527 F.3d 1359, 1372–73 (Fed. Cir. 2008)).)

In reply, Samsung again argues that Netlist has failed to show that the alternatives Mr. Kennedy relies upon are non-infringing, available, and commercially acceptable. (Dkt. No. 577 at 19–20.)

In sur-reply, Netlist again argues that the Court has already rejected these arguments at the *Daubert* stage. (Dkt. No. 588 at 19.) Netlist also argues that it is irrelevant that Dr. Brogioli criticized the "multidrop configuration." (*Id.* at 19–20.)

The Court is not persuaded by Samsung's arguments. First, Netlist need only show that there was an available next-best alternative to which it compared its claimed technology, not an acceptable non-infringing alternative. *Salazar*, 2018 WL 2033709, at *3. In lost profits cases, "the

CONFIDENTIAL INFORMATION HAS BEEN REMOVED FROM THIS PAGE

patentee must show a 'reasonable probability that, "but for" infringement, it would have made the sales that were made by the infringer.'" *Id.* (citations omitted). This, "in turn, requires the patent owner to prove an absence of 'acceptable' non[-]infringing substitutes." *Id.* (citations omitted). "But 'acceptable non-infringing alternatives' don't play the same role in a reasonable-royalty determination. Rather, in a reasonable royalty context, courts consider the next-best available alternative, which is not necessarily an 'acceptable' alternative." *Id.* (citations omitted).

Netlist has shown that there is a next-best alternative for each product class. For DDR4 products, Mr. Kennedy explained that that "as we heard from Doctor Mangione, without Netlist technology, the product sold wouldn't be usable in this two DIMMs per channel configuration and that the next best alternative would be to try to sell a bigger DPC LRDIMM at a lower price." (Dkt. No. 491 at 700:13–17.) Samsung argues that relying a single LRDIMM is relying on an infringing article, but the alternative need only be "next best," not non-infringing. (Dkt. No. 577 at 19.)

For DDR5 products, Samsung argues that there is no non-conclusory evidence that such an alternative was available. (Dkt. No. 577 at 20.) However, Dr. Mangione-Smith explained that Samsung "[redacted]" (Dkt. No. 489 at 700:4–7.) The Court finds this adequate.

For HBM products, Samsung argues that Mr. Kennedy relied on accused products—four-high HBMs—as opposed to non-infringing alternatives. (Dkt. No. 577 at 20.) Mr. Kennedy explained that "[s]o as we heard from Doctor Brogioli, Netlist technology allows Samsung to sell eight high or higher DIMMs, and they wouldn't be able to do that without Netlist technology. They'd have to try to sell four high HBM to replace that." (Dkt. No. 489 at 701:18–25.) As mentioned above, the alternative need not be non-infringing, just next best. *Salazar*, 2018 WL

2033709, at *3. While Netlist points to the "multi-drop" configuration, it does not point to any place in the record where such evidence was presented to the jury.

v. Evidence for Damages for the DDR5 Products

Netlist accused the DDR5 products of infringing the '918 and '054 Patents. (*See* Dkt. No. 561 at 60.)

Samsung argues that Kennedy's opinions for damages for the DDR5 products relied on a fundamentally flawed hedonic regression analysis performed by Dr. Groehn, who did not testify at trial. (*Id.*) Samsung notes that it had challenged the admissibility of Dr. Groehn's opinions on multiple grounds and that the Court had denied Samsung's admissibility challenge because Samsung could question Dr. Groehn on cross examination. (*Id.* at 60–61.) However, Samsung argues, it could not cross examine Dr. Groehn because he was never called to testify, even though Netlist placed him on the witness list every day. (*Id.* at 61.) Thus, Samsung contends, Mr. Kennedy improperly parroted the opinions of another expert, having never conducted a regression analysis himself. (*Id.* at 61–62.)

In response, Netlist argues that Samsung has waived these arguments as they are the same ones raised in its *Daubert* motion against Dr. Groehn, because they constitute a late *Daubert* challenge to Mr. Kennedy's analysis, and because Samsung did not raise any such objection about calling Dr. Groehn at trial. (Dkt. No. 573 at 64–66.) Netlist notes that Dr. Groehn was present in the courtroom every day at trial, and Samsung elected not to call him adversely when it was apparent Netlist would not call him. (*Id.* at 67.) Finally, Netlist argues it was not necessary to go through Dr. Groehn' s regression analysis with Dr. Groehn before the jury because "both parties had already walked through it with Mr. Kennedy." (*Id.*)

In reply, Samsung argues that Mr. Kennedy did not offer any opinions on regression in his report and thus Samsung could not have raised a *Daubert* challenge against Mr. Kennedy "before

knowing that Netlist would decline to call Dr. Groehn at trial, and that Mr. Kennedy would adopt Dr. Groehn's opinions as his own at trial." (Dkt. No. 577 at 21–22.) Samsung also argues that it need not have objected at trial to Mr. Kennedy's discussion of the regression because the issue did not become ripe until Netlist rested its case without calling Dr. Groehn. (*Id.* at 22.) Samsung again argues that Netlist's failure to call Dr. Groehn at trial is Netlist's fault because "Netlist had the burden of proof on damages."

In sur-reply, Netlist argues that Mr. Kennedy made clear numerous times that Dr. Groehn performed the analysis. (Dkt. No. 588 at 22.) Netlist also notes that Mr. Kennedy made equally clear in his report that he relied on Dr. Groehn's analysis. (*Id.*) Then, Netlist again argues that Samsung should have called Dr. Groehn adversely if it wished to test his opinions via cross examination. (*Id.* at 22–23.) Finally, Netlist argues that it is irrelevant that Netlist carries the burden of proof on damages because nothing requires Netlist to meet that burden using expert testimony. (*Id.* at 23 (citing *Rehco LLC v. Spin Master Ltd.*, 2020 WL 7025091, *10 (N.D. Ill. Nov. 30, 2020)).)

The Court finds that Samsung's arguments are unpersuasive. Samsung could have called Dr. Groehn adversely in its case-in-chief if it wished to subject him to vigorous cross-examination. They elected not to. Netlist was not obligated to call every witness and expert that its damages expert relied upon. If Samsung believed Mr. Kennedy should not have presented any facts about the regression analysis or a damages model based on that regression analysis unless and until Dr. Groehn had testified, Samsung should have raised that objection at trial. It did not. Samsung's contention that no issue arose until evidence closed without Dr. Groehn testifying is misguided. If no issue arose until the close of evidence, then there was no issue with Mr. Kennedy testifying in the first place. If Samsung was concerned that Mr. Groehn had not testified when Netlist closed its

rebuttal case, Samsung could have moved the Court to allow it a sur-rebuttal case in which to call Mr. Groehn adversely. Samsung did none of these things. Accordingly, Samsung's arguments are waived.

Additionally, these arguments essentially challenge Mr. Kennedy's reliance on Dr. Groehn, which is another classic challenge to Mr. Kennedy's methodology which was raised and rejected at the *Daubert* stage. (*See* Dkt. No. 205 at 10–11; Dkt. No. 432 at 4.) The Court finds that Samsung's arguments in this regard have been waived.

Samsung suggests that Netlist failed to show sufficient evidence solely because it did not call Dr. Groehn live at trial. This does not mean that Mr. Kennedy's testimony was insufficient on its own. If it were, the Court would have excluded it or would have prevented Mr. Kennedy from testifying about it unless Dr. Groehn first testified. Samsung goes too far when it now argues that Mr. Kennedy's testimony was flawed because Dr. Groehn did not first testify at the time and that it could not have known about this until after evidence closed.

vi. Reducing Damages to $19.3MM

Samsung contends that if the Court does not grant JMOL of no damages, then the Court should grant JMOL that damages are no more than $19.3MM. This is the number that Samsung's expert opined was an appropriate amount of damages. (Dkt. No. 561 at 62–63 (citing *Tronzo v. Biomet, Inc.*, 236 F.3d 1342 (Fed. Cir. 2001)).) In response, Netlist argues that substantial evidence supports the jury's award so there is no basis to reduce the damages. (Dkt. No. 573 at 67.) Neither Samsung nor Netlist address this topic in reply and sur-reply. (*See* Dkt. Nos. 577, 588.)

Since the Court has found that the damages awarded by the jury are supported by substantial evidence, reducing the damages awarded would be improper. The Court declines to do so.

vii. New Trial on Damages

Samsung argues that if the Court does not grant Samsung's JMOL on damages, then it should order a new trial on damages for two reasons: (1) the damages award is against the great weight of the evidence and (2) the damages award is based on Mr. Kennedy's testimony, which should have been excluded. (Dkt. No. 561 at 63–65.) In response, Netlist argues that Samsung's arguments are essentially a rehash of its JMOL arguments, none of which support granting a new trial. (Dkt. No. 573 at 68.) In reply Samsung asserts that Netlist does not respond to its argument that the damages award is against the great weight of the evidence. (Dkt. No. 577 at 23.) In sur-reply, Netlist argues that the standard for a new trial is highly deferential. (Dkt. No. 588 at 23 (quoting *i4i Ltd. P'ship v. Microsoft Corp.*, 598 F.3d 831, 857–58 (Fed. Cir. 2010)).)

The Court finds that the damages are not against the great weight of the evidence as they are supported by substantial evidence, as described above. Additionally, the Court has found that Mr. Kennedy's testimony should not have been excluded. The Court sees no reason to order a new trial on damages.

viii. New Trial on Damages Conditioned on Remittitur

Samsung argues that the court may also order a new trial conditioned upon remittitur of the "maximum amount the trier of fact could have properly awarded," which is $19.3MM according to Samsung. (Dkt. No. 561 at 65 (quoting *Delahoussaye v. Performance Energy Servs., L.L.C.*, 734 F.3d 389, 394–95 (5th Cir. 2013)).) In response, Netlist argues that Samsung's arguments do not comply with Federal Rule of Civil Procedure 7(b)(1)(B), which requires that motions "state with particularity the grounds for seeking the order." (Dkt. No. 573 at 68.) Netlist also argues that Samsung does not explain why $19.3MM is the "highest amount a jury could have properly awarded." (*Id.* (quoting *Longoria v. Hunter Express, Ltd.*, 932 F.3d 360, 364-65 (5th Cir. 2019)).) In reply, Samsung argues it clearly explained why $19.3MM is the highest amount a jury could

have awarded: because the Samsung and SK Hynix agreements are comparable and should control the outcome of the hypothetical negotiation. (Dkt. No. 577 at 23–24.) In sur-reply, Netlist argues that Samsung is rearguing a fact question that did not turn out as it hoped. (Dkt. No. 588 at 24.)

Since the Court sees no reason to order a new trial and sees no compelling reason to lower the damages awarded, the Court sees no reason to order a new trial conditioned on remittitur. It declines Samsung's invitation to do so.

D. Willfulness

The jury found that Samsung had willfully infringed at least one of the asserted claims of the '339 Patent, at least one of the asserted claims of the '918 and '054 Patents, and at least one of the asserted claims of the '918 and '054 Patents. (Dkt. No. 480 at 6.) At trial, the Court's ruling that Samsung did not infringe prior to the termination of the JDLA license in July of 2020 was in effect. (*See* Dkt. No. 561 at 66.)

Samsung argues that Netlist relied on a 2015 presentation (PX464) discussing a potential partnership with Samsung, but that the presentation does not mention the '339 Patent which had not even been filed at the time. (Dkt. No. 561 at 67.) Samsung also argues that Netlist relied on two documents discussing Netlist and the JDLA from 2019 to show willful infringement of the '339 Patent—PX1756 and PX1663—but PX1756 shows that Samsung had "not checked" whether it used any Netlist patents because of its license. (Dkt. No. 561 at 67.)

Samsung argues that Netlist pointed to PX586, an email from 2019 discussing a request from Samsung to discuss DDR5 technology, to show willful infringement of the '918 and '054 Patents, but there was no evidence that any of Netlist's patents were discussed in connection with the email. (Dkt. No. 561 at 67–68.) Samsung also contends that Netlist pointed to PX621 in its closing argument to show willful infringement of the '918 and '054 Patents but the document does not mention the patents, which had not yet been filed. (Dkt. No. 561 at 68.)

Samsung asserts that Netlist's argument for willful infringement of the '060 and '160 Patents centered on PX446, a patent list sent to Samsung in November 2016 listing only the '060 Patent. (Dkt. No. 561 at 68.) However, Samsung contends it was licensed to the '060 Patent at the time, foreclosing any inference that the document led Samsung to believe the HBM products infringe. (*Id.*) Samsung contends that it provided unrebutted testimony that Netlist's sharing of PX446 had nothing to do with Samsung. (Dkt. No. 561 at 68.)

Samsung contends that Netlist also identified a presentation, PX1778, that Netlist alleges was given orally to Samsung in 2015, but the email thread is entirely internal to Netlist, and Netlist produced no documentary evidence that the presentation was ever shown to Samsung even though a slide mentions the '060 Patent. (Dkt. No. 561 at 69.) Samsung then contends that Netlist never showed that any Samsung employee had knowledge of the Asserted Patents or alleged infringement was involved in the continued sales of the accused products—let alone anyone with sufficient authority such that their acts and knowledge can be imputed to Samsung. (*Id.* at 69–70.)

Finally, Samsung argues that no reasonable jury could have found that Samsung was willfully blind to infringement. (Dkt. No. 561 at 70.) Samsung first contends that the Court should not hold that willful blindness to infringement can constitute willful infringement. (*Id.*) Next, Samsung argues that Netlist cannot show that Samsung was willfully blind because there "was no evidence that Netlist informed Samsung of the alleged infringement until filing suit, and none of the communications from before or during the license period was sufficient to show that anyone at Samsung believed years later that infringement was highly likely and took deliberate steps to avoid this knowledge—let alone that any such individual had responsibility for the alleged infringement." (*Id.* at 70–71.)

In response, Netlist argues that there was ample evidence presented to the jury for it to find willful infringement. (Dkt. No. 573 at 68–70.) Netlist argues that Samsung's willful infringement after the filing of the complaint supports the jury's verdict. (*Id.* at 70 (citing *Packet Intel. LLC v. NetScout Sys., Inc.*, 2019 WL 2375218, at *8 (E.D. Tex. June 5, 2019), *aff'd in relevant part*, 965 F.3d 1299 (Fed. Cir. 2020)).)

In reply, Samsung argues that Netlist does not dispute that any claim of willful infringement is limited to a time after the alleged termination of the license in July 2020 and, for four of the five patents, after the lawsuit was filed in December 2021. (Dkt. No. 577 at 24.) Samsung argues, however, that the documents cited by Netlist predate the license and the termination. (*Id.*) Samsung re-urges that the documents either do not mention the asserted patents or are entirely internal to Netlist. (*Id.*) As such, Samsung asserts that Netlist's willfulness argument comes down to nothing more than Samsung's continued selling of its products after Netlist terminated the JDLA and filed suit. (*Id.*)

In sur-reply, Netlist argues that "pre-damages conduct can support a finding of willfulness infringement." (Dkt. No. 588 at 24.) Netlist then summarily asserts that Samsung "continued to sell infringing products despite its knowledge that it copied Netlist's technology." (*Id.* at 25.)

The Court is not persuaded by Samsung's arguments. First, willful infringement is an inquiry that considers the "totality of the circumstances." The inquiry is, by its nature, factually intensive. This Court sees no basis to overturn the jury's determination in this regard. *See WCM Indus., Inc. v. IPS Corp.*, 721 F. App'x 959, 970 (Fed. Cir. 2018). There were many potentially relevant facts before the jury, which were necessarily considered by the jury in reaching this result. (*See* Dkt. No. 573 at 68–70.)

Further, Samsung offers no compelling reason why communications made while it was licensed (or before) could not contribute to Samsung's willful infringement. It is logical that Samsung may understood that it would have been infringing but for the license, and this knowledge may properly contribute to a post-license finding of infringement that is willful.

Finally, Samsung's post-complaint actions may also properly support a finding of willfulness. *Packet Intel. LLC v. NetScout Sys., Inc.*, 2019 WL 2375218, at *8 (E.D. Tex. June 5, 2019) (emphasis in original), *aff'd in relevant part*, 965 F.3d 1299 (Fed. Cir. 2020). It is undisputed that Samsung released a product accused of infringement after the filing of the complaint. (Dkt. No. 493 at 1089:25–1090:23.)

E. New Trial

As it rounds out its 50(b) arguments, Samsung says that it is entitled to a new trial. (Dkt. No. 561 at 71–75.) First, Samsung argues that Netlist's improper trial conduct unfairly prejudice Samsung. (*Id.* at 71–74.) Samsung argues that (a) Netlist improperly argued claim construction to the jury, including during closing (b) violated a motion *in limine* precluding discussion of the value of the supply provisions of the SK hynix agreement and the JDLA, and (c) made "jingoistic comments in closing arguments to inflame the jury against Samsung on the basis of national origin." (Dkt. No. 561 at 71–74.) Samsung also argues that the Court improperly excluded relevant evidence concerning the standard setting body JEDEC. (*Id.* at 74–75.)

In response, Netlist argues that the jury was capable of following the Court's instructions regarding claim construction, which stated that the jury must accept and apply the Court's constructions, and that Samsung raised no objections during Netlist's closing arguments. (Dkt. No. 573 at 70–71.) Netlist also contends that the Court only precluded Mr. Kennedy from referencing the value of the supply provisions of the SK hynix agreement and that it was Samsung's expert, Mr. Meyer, who announced the number in open court. (*Id.* at 71–73.) As to any comments Netlist's

counsel made in closing related to national origin, they were simply to rebut the statements counsel for Samsung had made when it opened that door in its closing argument. (*Id* at 73–74.) Finally, Netlist argues that the Court properly excluded evidence regarding JEDEC, given that Samsung took the position that the patents-in-suit were not essential to any JEDEC standards. (*Id.* at 74–75.)

In reply, Samsung argues that it did not waive any objections during closing because the Court had instructed that it did not "want there to be objections to the other one's closing unless there is absolutely no alternative." (Dkt. No. 577 at 24–25 (quoting Dkt. No. 495 at 1229:17–18).) Samsung also argues that it did not open the door to Netlist's "prejudicial references to Samsung's national origin." (*Id.* at 25.) Finally, Samsung urges that exclusion of JEDEC evidence was improper because it prevented Samsung from discussing "non-accused technology." (*Id.*)

In sur-reply, Netlist argues that if its comments were so prejudicial as to warrant a new trial, then Samsung had no alternative but to object and should have. (Dkt. No. 588 at 25.) However, Netlist denies that it ever disparaged anyone's national origin. (*Id.*) Finally, Netlist argues that the Court's exclusion of JEDEC evidence did not prevent Samsung from discussing "non-accused technology."

After fulsome review of the record, the Court is persuaded that a new trial is not necessary or proper in this case. If the comments made by counsel for Netlist during closing were so prejudicial that they would warrant a new trial, then Samsung should have objected. They did not. Regarding the mention of the value of the supply provision of the SK hynix agreement, Samsung does not dispute that its own expert put this number into evidence, and this opened the door—about which it cannot complain. (Dkt. No. 495 at 1201:3 (Meyer).) Further, Samsung does not contest that the jury was able to follow, and did in fact follow, the Court's instructions regarding

claim construction. The comments that counsel for Netlist and Samsung made related to national origin were neither disparaging nor so prejudicial that a new trial is warranted. Finally, the Court continues to be persuaded that evidence related to JEDEC is not relevant because the Asserted Patents are not and were not alleged to be standard essential. To have allowed evidence about a nonapplicable standard would have injected serious confusion into this trial. Further, exclusion of evidence about JEDEC did not prevent Samsung from discussing "non-accused technology."

VI. CONCLUSION

For the foregoing reasons, the Court finds that the Motion (Dkt. No. 561) and the Motion to File a Supplemental Brief (Dkt. No. 591) should be and hereby are **DENIED**.

The parties are directed to jointly prepare a redacted version of this Order for public viewing and to file the same on the Court's docket as an attachment to a Notice of Redaction within five (5) business days of this Order.

So ORDERED and SIGNED this 12th day of July, 2024.

RODNEY GILSTRAP
UNITED STATES DISTRICT JUDGE

IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION

NETLIST, INC.,	(	CAUSE NO. 2:21-CV-463-JRG
Plaintiff,	)	
vs.	)	
SAMSUNG ELECTRONICS CO., LTD., et al.,	)	MARSHALL, TEXAS
Defendants.	)	MARCH 28, 2023
	)	9:00 A.M.

VOLUME 1

PRETRIAL CONFERENCE

BEFORE THE HONORABLE RODNEY GILSTRAP
UNITED STATES CHIEF DISTRICT JUDGE

SHAWN McROBERTS, RMR, CRR
100 E. HOUSTON STREET
MARSHALL, TEXAS 75670
(903) 923-8546
shawn_mcroberts@txed.uscourts.gov

A P P E A R A N C E S

FOR THE PLAINTIFF:

IRELL & MANELLA, LLP - LOS ANGELES
1800 AVENUE OF THE STARS
SUITE 900
LOS ANGELES, CA 90067-4276
(310) 203-7096
BY: MR. JASON SHEASBY
MS. YANAN ZHAO

McKOOL SMITH, P.C. - MARSHALL
104 E. HOUSTON ST., SUITE 300
MARSHALL, TEXAS 75670
(903) 923-9000
BY: MR. SAMUEL BAXTER
MS. JENNIFER TRUELOVE

FOR THE DEFENDANTS:

FISH & RICHARDSON PC - WASHINGTON DC
1000 MAINE AVE., SW
SUITE 1000
WASHINGTON, DC 20024
(202) 783-5070
BY: MR. RUFFIN CORDELL
MS. LAUREN DEGNAN
MR. BRIAN LIVEDALEN
MR. MATTHEW MOSTELLER
MR. MICHAEL McKEON

FISH & RICHARDSON, PC - SAN DIEGO
12860 EL CAMINO REAL
SUITE 400
SAN DIEGO, CA 92130
(858) 678-5070
BY: DR. FRANCIS ALBERT

FISH & RICHARDSON, PC - DALLAS
1717 MAIN STREET, SUITE 5000
DALLAS, TEXAS 75201
(214) 747-5070
BY: MR. MATTHEW COLVIN
MR. THOMAS REGER

FISH & RICHARDSON, PC -
NEW YORK
7 TIMES SQUARE, 20TH FLOOR
NEW YORK, NEW YORK 10036
(404) 724-2764
BY: MS. KATHERINE REARDON

GILLAM & SMITH, LLP
303 SOUTH WASHINGTON AVENUE
MARSHALL, TEXAS 75670
(903) 934-8450
BY: MS. MELISSA SMITH

GILLAM & SMITH, LLP
102 N. COLLEGE, SUITE 800
TYLER, TEXAS 75702
(903) 934-8450
BY: MR. TRAVIS UNDERWOOD
MR. TOM GORHAM

OFFICIAL REPORTER: SHAWN M. McROBERTS, RMR, CRR
100 E. HOUSTON STREET
MARSHALL, TEXAS 75670
(903) 923-8546

THE COURT: Be seated, please.

This is the time set for pretrial matters before the Court in the case of Netlist, Inc., versus Samsung Electronics Company, Ltd., et al. This is Civil Case No. 2:21-CV-463.

The Court will ask for announcements at this time. What says the Plaintiff?

MS. TRUELOVE: Good morning, Your Honor. Jennifer Truelove here for Plaintiff. With me today at counsel table, we have Mr. Jason Sheasby, Ms. Yanan Zhao, and Mr. Michael Rosen.

We are ready to proceed.

THE COURT: Thank you.

What's the announcement for the Samsung Defendants?

MS. SMITH: Good morning, Your Honor. Melissa Smith on behalf of Samsung.

I have many people in the courtroom that have joined me today, but what I'd like to do is introduce Your Honor to those that will be arguing today on behalf of Samsung in the order of argument: Mr. Ruffin Cordell, Mr. Mike McKeon, Ms. Lauren Degnan, Dr. Frank Albert, Mr. Brian Livedalen, Mr. Matt Colvin, Mr. Matthew Mosteller, Mr. Tom Reger, and in the back Ms. Katherine Reardon.

Your Honor, we also have a client representative in the courtroom today, Mr. Michael Nguyen, and we're ready to proceed, Your Honor.

THE COURT: All right. Thank you.

Counsel, we have a tremendous amount of material to cover in the next two days, and the next two days are the only time I have on my calendar to cover pretrial matters before we try this case in April. We are all going to have to move at a rapid clip to get through this. Anything we don't get through in the next two days, I'll decide on the papers.

And this is the first case in almost 12 years on the bench I have thought about putting an artificial limit on the number of dispositive motions that could be filed just in light of what you've done here. Not only have I got north of 20 dispositive motions, but most of them have multiple subparts, sometimes 10, 11, and more subparts to them. It's a tremendous amount of material.

But we have two days, and we're going to use every bit of it. You're just going to have to work with the Court to move promptly through this so that we can cover all of it or at least as much of it as humanly possible.

Let me begin by giving you some background information and housekeeping data.

Currently the case is set for jury selection on Monday, April the 17th. I'm going to change that. We're going to pick the jury on Friday, April the 14th. So you need to be here for jury selection on Friday, the 14th of April, at 9:00 a.m.

After the jury's selected and seated, we'll proceed with opening statements. We may get the first witness or so from the Plaintiff's case on that day. We'll use the entirety of that day to begin the case, and then we'll continue the following Monday, et cetera.

Each side's going to have 30 minutes to examine the venire panel for jury selection purposes. As is the Court's usual practice, both sides are entitled to use up to three minutes at the beginning of their time to address the panel to give a very high-level, non-argumentative thumbnail sketch of what's at issue here.

I continue to have lawyers who lapse into being argumentative during those first three minutes, and I take appropriate action in those cases from calling them down in front of the panel to taking further steps if necessary.

So let me reiterate to you, any high-level overview of what's at issue that you choose to put forward during your first three minutes or some portion of your first three minutes needs to be non-argumentative.

And if it becomes argumentative, I will not hesitate to take appropriate steps, including but not limited to, limiting the amount of your time before the panel.

I'm going to seat eight jurors to hear the evidence in this case. Each side's going to be afforded four peremptory challenges.

I'm going to allocate 12.5 hours to each side for presentation of their evidence. That's not including opening statements, jury selection, or closing arguments.

As is the Court's usual practice, I'm not going to require objections to deposition witnesses to the designations or counterdesignations urged with regards to deposition witnesses in advance of beginning of trial, other than to tell you I will take them up on a rolling basis through the trial. But I do expect any disputes regarding deposition witnesses to be brought to the Court's attention not later than the day before the day you intend to present those deposition witnesses.

I do not want to have a break in the trial or a delay because deposition disputes are raised too late in the process to make the proper adjustments and keep the case -- the trial on track. So bring me any disputes regarding deposition witnesses not later than the day before the day. And I'll take them up rolling on that kind of a basis throughout the remainder of the trial.

I'm going to follow the similar approach that I've used in earlier cases. You're to report to the Court each evening by 10:00 as to any progress you've made with regard to overnight disputes regarding demonstratives and other things. If at your 10:00 report you still have disputes that haven't been resolved through the meet-and-confer process, then you're

to continue your efforts to resolve those through the meet-and-confer process.

And if by 7:00 the next morning, there are unresolved disputes from overnight, then you're to present a binder, a jointly-prepared three-ring binder, to chambers by 7:00 a.m., setting forth what's in dispute and each party's respective positions to what's in dispute. That means if it's a disputed demonstrative, I need to see the demonstrative slide. And I also need at least a single paragraph, not more than, but a single paragraph from each side about why they have a problem with that demonstrative slide.

You would be amazed at the times I get a binder at 7:00 a.m. that gives me two paragraphs about a slide that's not included and I can't see what they're arguing about or, vice versa, I get the slide and no commentary from either side about why they have a problem with it.

So it doesn't do me any good to get partial information that doesn't allow me to be prepared to deal with it directly and appropriately when I meet with counsel before we bring in the jury.

I'm available each morning by 7:30 during the trial for the sole purpose of helping resolve these overnight disputes, streamline the process, and maximizing your designated trial time.

If we can't do it right before the jury comes in, we'll

do it after the jury's here and the time that's consumed in doing that will be charged to your trial time. So you will benefit by helping the Court do that in a way that's effective and efficient. That means I need complete information about what may have survived the overnight meet-and-confer resolution process.

With regard to closing arguments, the Plaintiff is entitled to split their time, but they must use at least 50 percent or 20 minutes of their designated 40 minutes in their first closing argument.

I'll defer any argument presented regarding motions under Rule 50(a) of the Federal Rules of Civil Procedure until all the evidence has been presented. And after all the evidence has been presented and I have excused the jury, then I'll take up motions from both sides as to issues that either may elect to raise under Rule 50(a).

After I've dealt with motions under Rule 50(a), I'll conduct an informal charge conference to review the then most current proposed final jury instruction and verdict form.

I see many lawyers in the courtroom who've tried cases before me in the past. You're well aware of the Court's policy of prohibiting reference to individuals on the record by first name only as being inherently confusing and being inappropriate regarding the proper decorum of the United States District Court.

Let me remind you to remind your witnesses of this. If your witnesses start talking about Jim, Jane, and John, I'm going to hold you responsible for that.

I'll refer you to the Court's standing order on protecting proprietary and confidential information by sealing the courtroom which seals that appropriate part of the record, and I'll remind you that the post-trial redaction process is not appropriate for that purpose.

So if it comes out in open court during the trial in front of the jury and whoever may be in the courtroom, it's -- it's out. It's not going to be protected by some kind of post-trial redaction process. So follow the -- follow the procedures outlined in the Court's standing order on sealing the courtroom.

My only request there, Counsel, is if you have various sections of your examination that may require sealing, that you try to consolidate them as much as possible so that the process of sealing the courtroom isn't like a light switch that gets flipped on and off and on and off during the examination. That can be disruptive, and it also takes time for everybody who's not subject to the protective order to leave the courtroom.

Again, your time is your time and you are the marshals of the effective use of it. But to the extent you can consolidate your information that needs to be protected by

sealing the courtroom, that's helpful.

I'll refer you to the Deputy-In-Charge for the Marshall Division with regard to access and use of the juror questionnaires that have been submitted. Her instructions to you are to be taken as my instructions to you.

I will tell you, though, in advance of that you are not to copy, retain, scan, or keep any records of what's set forth in those juror questionnaires. We encourage citizens to be candid in answering those questions to assist you in selecting a fair and impartial jury. Part of the way we do that is we assure them that their information in those candid responses aren't going to be retained in some lawyer or jury consultant's database for future appearance at an unexpected time.

So you're not to copy them, scan them, email them, retain them in any fashion. Other than that, the use of them, including the return of the hard copies to the clerk, I'll refer you to Ms. Clendening and her instructions in that regard.

I don't know exactly how many deposition witnesses we're going to have. To the extent your deposition witnesses are presented by video recording and you have the captioning, it's helpful that you include the captions in the video presentation if it's available.

I will tell you that the Court was surprised when I

learned that you had submitted proposed disputes regarding the pre-admission of exhibits in this case that paid absolutely no heed to the Court's standing order on the submission of exhibits for pre-admission.

I obviously by order struck your previous submissions and gave you a short period of time to resubmit in compliance with the Court's standing order, and that's been done, but I don't know why in the world you would have ignored my standing order on pre-admission of exhibits.

We're going to get to pre-admitted exhibits in this pretrial process, but nonetheless, going forward, you're on notice that the Court's standing order with regard to the use and number of pre-admitted exhibits and their submission in the pretrial process in patent infringement cases such as this is not optional.

I've also looked at what may be floating around out there with regard to an early version of proposed jury instructions and verdict form. I don't think it's going to be helpful to the Court at all. I'm going to direct that you jointly prepare and submit a more current version of a proposed final jury instruction and verdict form not later than Monday, April the 10th, at 5:00 p.m.

Also I'm going to direct that you jointly work together to prepare 12 juror notebooks to be used during the trial, and those juror notebooks should be delivered to chambers not

later than 12:00 noon on Monday, April the 10th. Those juror notebooks, as in the past, should be three-ring binders.

They should include a single-sided copy of each of the patents-in-suit. They should include a side-by-side ledger or chart on the Court's claim construction reflecting on the left side the claim language that's been construed and, correspondingly, on the right side the adopted construction by the Court.

They should include a section for witness pages where each witness who's to be called or may be called as identified on a single page with their head-and-shoulders photograph superimposed at the top of that page, their complete name underneath without any characterization, and the remainder of those pages as ruled lines for potential note-taking.

Those witness pages should also be tabbed with the witness' name on the tabs to facilitate the members of the jury quickly finding those as the witnesses are called. Behind the tabbed witness pages, there should be a new legal pad inserted with three holes in the notebooks for additional note-taking, and you should include a pen that doesn't click or make noise in the front pocket of each notebook.

And those 12 iterations of the juror notebooks should be delivered, again, to chambers not later than noon on Monday, April the 10th.

It's clear that there are going to be a large number of

expert witnesses testifying in this case, and I don't think there is any doubt or should be any doubt in counsel's mind that those experts are going to be confined to the four corners of their reports.

Much of this pretrial process is going to be presentation and argument regarding efforts to limit or strike some or all of those reports. Whatever reports in whatever shape survive the pretrial fashion comprise the four corners and the parameters of what each expert witness is permitted to testify to.

Objections raised before the jury during the trial that an expert is attempting to testify beyond the scope of their report are often necessary, but they are inherently disruptive and time-consuming. Make them where you believe they're merited. I simply would caution you that if the Court comes to the conclusion that objections of that type are being raised for strategic purposes of any type, the Court will consider that a sanctionable offense and will take whatever action is appropriate.

There had been an earlier dispute about designating an interpreter for use during the trial. My latest understanding is that has been worked out. Are there any remaining disputes about the designation of interpreters for non-English-speaking witnesses that either side's aware of?

MR. SHEASBY: Nothing from Plaintiff, Your Honor.

MR. CORDELL: Nothing from Defendant, Your Honor.

THE COURT: All right. I'll consider that that's been resolved.

All right. Let me ask at this point if there are any questions from either side as to these instructions and housekeeping matters that I've just gone over with you.

Anything from Plaintiff?

MR. SHEASBY: Nothing from Plaintiff, Your Honor.

MR. CORDELL: Nothing from Defendants, Your Honor. Thank you.

THE COURT: All right. All right. Well, without further ado, let's get into these many disputed dispositive motions and see how quickly we can work our way through them.

I'm going to begin with Docket No. 201, Plaintiff's motion for partial summary judgment on Samsung's license affirmative defenses. Let me hear from the moving Plaintiff first, please.

MR. SHEASBY: May it please the Court, Jason Sheasby for Plaintiff.

The Central District of California has ruled dispositively on four issues: One, that there was a supply obligation; two, that the clause is material; three, that there was a material breach; and, four, that there was a termination of the relationship between the parties.

And that relationship went far beyond just a license to

patents or supply. There was joint development relationships, there was right of first refusals for technology, and that all went away.

The relevant issue for the Court is in their opposition to what should have been a simple motion for summary judgment, there was an argument that -- of essentially a cross motion for summary judgment that they should be entitled to take advantage of whatever license existed for the period of time before the agreement was terminated.

Can I have my slide deck 5, please, Mr. Huynh?

THE COURT: Let me ask you this, Mr. Sheasby. Does the Plaintiff contend that there was an effective termination of the JDLA before July 15th, 2020?

MR. SHEASBY: Not effective termination. Under New York law, the standard is that there was a substantial failure to perform. And upon a substantial failure to perform under New York law, the rights to -- for executory contract cease. So if you look at Mega Group, this is the mid -- sort of interim level appellate division: When promises in a contract are mutually dependent upon a failure to act by one party, the other is excused of performance, and that's a question of fact for the jury.

You can also see this in Merrill Lynch, which is a Second Circuit opinion. That would be the -- sort of the most authoritative statement on New York law but for a decision

from the court of appeals from New York: A party's performance under contract is excused when the other party has substantially failed to perform its side of the bargain. That's a question of law for the jury.

And, of course, I don't think there's any dispute that a license agreement, which is -- a non-exclusive license agreement which is nothing more than a commitment not to sue or to avoid having liability for infringement, is an executory contract.

So the short answer is that under New York law there's two questions. There was a complex relationship between the parties that was terminated. It was terminated on a date certain. What Samsung is arguing is, even though it was terminated on a date certain, before that termination we were entitled to take advantage of our -- of the commitment to forego patent liability.

And New York law has two tiers to that. New York law says that from the time period of when there's substantial non-performance, the obligation for us to perform has ceased as well. So this is a contract issue as to whether there's substantial non-performance.

Does that answer your question, Your Honor? It's not an effective or putative termination. The termination occurred on a date certain. Under New York law before the termination, if there was substantial non-performance, the agreement -- we

don't have an obligation under the agreement to sort of give them a right, a license.

THE COURT: Let me ask you this. Is there a difference between being excused from performing under the JDLA and revoking the licenses that are granted under the JDLA?

MR. SHEASBY: Right. So that's what I think the Ninth Circuit has addressed in *In re CFLC*, which is, a license under the JDLA or non-exclusive license is nothing more than a commitment not to sue. And so the -- sort of the right to be free of a lawsuit, the right to not have patent infringement liability asserted against them, we believe ceased upon substantial non-performance of an executory contract.

And so I -- the Ninth Circuit's opinion, which is the only controlling authority we've found on the subject, has recognized, and I think this is consistent with the Federal Circuit, a non-exclusive license is just a commitment not to sue. They lost their commitment to be free of patent infringement liability from the point of substantial non-performance.

THE COURT: As far as you know, there's no dispute here that after July 15th, 2020, there's not a license? This is only -- this argument really only relates to prior to July 15th, 2020.

MR. SHEASBY: That's correct. And it's -- under New

York law, it's a question of fact for the jury as to whether -- so the -- under New York law, if it's a mutually dependent promise, if it's an executory contract, their right to be free of patent infringement liability ceased upon substantial non-performance.

They do make an argument that there was an issue preclusion or claim preclusion associated with this.

I'll go to slide 28, Mr. Huynh.

So the Court in the Central District of California determined the date on which the license agreement was terminated under a declaratory judgment clause. There was never a situation in which Samsung said, oh, but we get patent rights before the termination. That was not an issue that was joined in the Central District of California case.

THE COURT: Now, Judge Scarsi said that the contract, the JDLA, had been terminated. Did Judge Scarsi specify an exact date of termination?

MR. SHEASBY: He -- he -- he said that on this date the contract was terminated. He did not say --

THE COURT: Termination is effective as of this date.

MR. SHEASBY: I don't believe -- in other words, I -- I want to be precise. There was -- the issue in which Samsung says, well, maybe we've breached and maybe it's terminated, but we get our rights for all the period of time

before that termination. That issue was never joined. It was never decided by Judge Scarsi. It was not a subject of any issue that Judge Scarsi had to decide.

And, of course, it was a declaratory judgment claim, and under Fifth Circuit law, the Fifth Circuit is very parsimonious as to when there will be a claim preclusion associated with declaratory judgment actions.

And to be honest, this is not sort of a gotcha or -- or us being clever. There was no patent infringement fight in front of Judge Scarsi. Because there was no patent infringement tried in front of Judge Scarsi, the question was not joined.

Samsung has raised a license defense in this case for the first time. They say, even though the agreement is terminated, we still have a license for the period before termination, and that issue needs to be joined.

And that issue is three questions, none of which were actually before Judge Scarsi. One is the license and executory obligation. We believe it is. That's a question of fact. Two is their mutual dependence of the executory obligation. And three was their substantial non-performance. So the issues that must be decided by the jury before that period before the formal termination need to be -- are different.

So I want to be very clear. Judge Scarsi pointed to the

date on which we sent our termination notice. Judge Scarsi acknowledged that that termination was effective. What Judge Scarsi never did was say, well, what happens for the period of time before termination, because that wasn't relevant to any issue that was before the Court.

So I have to suggest I think there's an undercurrent of sharp practice here. I don't think there was sharp practice. I don't think the parties in any way, shape, or form engaged what was the consequences of the material breach before the formal termination.

And just from a practical standpoint, if you'd go to slide 8, Mr. Huynh.

So this is January of 2018 when they basically say there's zero performance. They're stopping performing under the contract. Under New York law, there is a non-trivial question as to whether they are free of patent infringement liability from the point in time when they announced a total non-performance.

That's not a -- a gotcha. That's not sort of a trick for a jury to get confused. There's a serious material issue as to whether they have the right to take advantage of this commitment not to sue before that date.

And that's Exhibit 16 to our opposition on their motion for summary judgment at 9 to 10, if Your Honor has questions about that document. I'll stand down unless Your Honor has

any further questions.

THE COURT: Let me hear from Samsung, please.

MR. CORDELL: May I approach, Your Honor?

THE COURT: You may.

MR. CORDELL: May it please the Court. Ruffin Cordell for Samsung.

THE COURT: Please proceed.

MR. CORDELL: Thank you, Your Honor.

What we're finding out now, very late in the case, that none of this was disclosed earlier. I'll show you interrogatory responses from early in the case, and in fact the complaint in this case tells Your Honor in this case, tells this Court, a representation, a pleading filed followed by a sworn interrogatory, that the JDLA terminated on July 15, 2020.

There was no mention of repudiation. There was none of this business about there being an earlier termination or somehow loss of rights earlier in the case, all the way through this until the Hong deposition where their CEO, Mr. Hong, finally hatches this theory out of whole cloth.

What's really going on here, the Paul Harvey rest of the story here, is that what Netlist wants to do, you are going to hear this all day today and apparently tomorrow as we face motion after motion, it's been revealed that Netlist's true strategy here is to relitigate the JDLA breach of contract

action in this court. That's what they want to do. We're going to hear it time and again.

Mr. Sheasby, last time we were here, won the right to depose J.B. Lee, our president of the semiconductor business, and he asked about a dozen questions to what Your Honor had authorized him to do, the meetings that had taken place. He spent the vast majority of his time asking President Lee about the JDLA and the breach of the JDLA and whether supply agreements were fulfilled or not, and was that a good thing or a bad thing.

That's what we're going to hear, Your Honor. And we're going to have to put up a bulwark, and we're going to have to put it up early in this case because, otherwise, what we're going to be doing is relitigating the JDLA. None of that is relevant here. Despite Mr. Sheasby's valiant efforts to try to dredge up issues to make it relevant, we have to allow this jury to decide the case that is going to be presented before them.

So as I was sitting here this morning and Your Honor was reminding us that we get our three-minute neutral statement at the outset of the case to tell the jury what this -- this trial is going to be all about, and it struck me that I don't know what Netlist is going to do. Are they going to stand up and say, this is a patent infringement case, we've got six patents over two disparate technologies, it's going to be, you

know, very bewildering to keep up with, but we're going to do our best to give you the information; or, are they going to say, this is about two companies that have had a long relationship and had a prior JDLA that Samsung, we believe, breached and are they going to present their breach of contract case? Because that's what I get from reading all of the mountain of papers that we presented to Your Honor.

So when we're confronting this particular issue, there are a few things that I think we can start with that make it pretty easy. If I can get this to work.

Can we have the next slide?

So this is a collection of the motions that are relevant to this particular issue. There are cross motions here. There is a motion that we filed asking the Court to enter judgment that, in fact, all accused products before July 15, 2020, are licensed so that we can avoid all of the contract issues that Mr. Sheasby just walked us through.

We have a cross motion that Netlist filed, after they saw ours, saying, well, we kind of agree with that; that, yes, there is a -- there is a termination that happened on that date. But then they phrased it in a particularly odd way. They said, we want -- want the Court to enter judgment that termination happened as of that date, as if it could have been that date or a date that was earlier. And we've seen that back and forth.

Next slide. Thank you.

So there's a little bit of good news here, Your Honor. In their reply brief, Netlist said that they agreed with us that termination as of July 15, 2020, means terminated on. And that is good news because that means that they're not disputing what Judge Scarsi did, they're not disputing the process that led the parties to that termination, and at least in my mind that means that we shouldn't be talking about what happened prior to that date. Those -- that is the judgment that Judge Scarsi entered.

And I will take the Court later through the issue preclusion and ultimately claim preclusion effects that -- that prevent that from happening.

Can I have the next slide?

THE COURT: Let me ask you this, Mr. Cordell. Like Netlist, Samsung doesn't contend that there's a license defense after July 15th, 2020.

MR. CORDELL: Correct. The parties are in agreement there.

THE COURT: Okay.

MR. CORDELL: So we start with the issue as to whether or not the JDLA covered the accused products, and we start with the JDLA itself and the language is very clear. And Your Honor is one of the best contract lawyers I've ever met so this is certainly probably nothing the Court needs.

Can I have the next slide?

But we have definitions in the agreement that tell us what's licensed, and it's Samsung's licensed products which are all semiconductor products manufactured by or for Samsung or its subsidiaries. And the exclusion that we're going to talk about here are foundry products, although Mr. Sheasby didn't spend much time on that. But that is an issue, and I'll touch on that.

Next slide, please.

Before the California court, Netlist argued that the JDLA was not limited to joint R & D. In the papers, Your Honor may have seen that there was an issue about whether or not we were going to limit -- the parties would limit the scope of the license to the JDLA only. And the JDLA by its terms was not limited -- the license under the JDLA was not so limited to the parties' joint development products.

You may recall that under the JDLA, the parties were to cooperate and create new products. And there was an issue that has been bandied back and forth about whether the license was somehow narrowly limited only to the jointly developed products.

I don't think that anybody is carrying that issue forward. And Netlist convinced Judge Scarsi that it was not so limited. And that was an issue. It was an issue that was presented to him, and they succeeded. There will be estoppel

and issue preclusion consequences for that.

Netlist has told the whole world that Samsung was licensed to everything. So back in 2017, they relied on their license to Samsung as their domestic industry in an ITC case. The Court's familiar with the fact that in an ITC case, the patentee has to prove that they or their licensees are practicing the patents.

And what did Netlist do in 2017? They went to the U.S. International Trade Commission and said, there's Samsung, our licensee, they are practicing these patents; therefore, we checked that box on what we have to prove.

We've seen in the parol evidence that in fact Netlist admits the license covered all of their patents and that related to all of Samsung's exposure on LRDIMMs, all of Samsung's memory that's at issue in this case.

We have had a little bit of an issue about the Korean Tax Tribunal, and I can cover that in more detail. But there's one fundamental fact that Netlist leaves out in all of their arguments here, and that is that Samsung was not a party to the Korean Tax Tribunal matter.

What happened with the Tax Tribunal is that Samsung was to pay Netlist some money. Korean law says, whenever you are paying money, you have to withhold taxes if it's for a patent royalty. There was a dispute before the Korean tax authority where Netlist said, hey, we want that money back. And they

had a case in front of the Korean tax authority. Samsung wasn't invited.

So all of these arguments that somehow Samsung was bound by that or estopped by that fail at the very first instance because Samsung wasn't invited to the party. They weren't there. So they can't be estopped; they can't be binding on them.

THE COURT: Let me stop you for a minute, Mr. Cordell. You made a statement a second ago that nobody is arguing here that the JDLA was limited to the joint development project. I saw that argument in Plaintiff's briefing over and over.

So has that been conceded to your knowledge, or did I misread their briefing, or --

MR. CORDELL: You are correct.

THE COURT: Revisit that statement for me.

MR. CORDELL: Sure. You are correct that that was -- and I hoped I acknowledged that, that it was in -- it was in the papers at the beginning. It just didn't --

THE COURT: You don't need to add anything to this, Mr. Sheasby.

MR. SHEASBY: I'm not going to say -- but this is sort of Netlist's confidential information. I'd like for the courtroom to be sealed because we are a public company. The terms of it both are material.

THE COURT: Are you talking about the terms of the JDLA?

MR. SHEASBY: Yes, Your Honor. I'm not trying to --

MR. CORDELL: I'm going to do it at a very high level.

THE COURT: I don't think we need to seal the courtroom at this point. You can reurge it later.

Go ahead, Counsel.

MR. CORDELL: My point is that as the briefing was developed, no one seemed to -- to be focused on whether or not these products were the jointly developed products or not. That issue seemed to kind of go by the boards, as it should, because the reality is that JDLA was in no way limited to jointly developed products.

Again, I can take you back to the language of the license, and the license --

THE COURT: I've read the JDLA.

MR. CORDELL: Yes. Thank you.

THE COURT: Several times unfortunately.

MR. CORDELL: I'm going to come back to it yet again.

So we do have an issue about the foundry products. And Mr. Sheasby didn't visit this, but I will and I'll do it in some detail. But the reality is people in the business know exactly what a foundry product is.

They know that a foundry product is when a company comes to a foundry, people that have a semiconductor line, and they have a tape in their hands, and they hand it over and they say, can you make this tape into a semiconductor product? They feed the tape into the machine, out come the semiconductor products.

It has to be designed by somebody else. It has to be their information. It has to be their product. That's the way the foundry product process works.

But, importantly here, and, again, I'll treat this in some detail, the problem that Netlist faces in this case is that they haven't accused any foundry products. So it's one thing to say that Samsung has a foundry product business. That is true. Samsung will make products for other companies.

But the only way it's relevant in this case is that if Netlist says product X28.48 is accused and it's a foundry product and therefore excluded from the license. And they just haven't done that. So we don't have any accusation of a foundry product which makes it very difficult.

What we know from the experts is that they say the products that actually are accused were designed by Samsung, not by some third party. So we have Doctor Brogioli, who is the '060 and '160 HBM patents' expert. And he affirmatively says that Samsung is the one who is designing and manufacturing these accused products. It's not incidental.

He says it -- says it right out because he wants to show direct infringement, so he's got a reason to make that statement.

Slide 16.

We have again another statement from him about the design and the -- the purposeful design of the products for the '060 and '160 Patents. It's not incidental. It's not off the cuff.

Doctor Mangione-Smith, who is the expert on the remainder of the patents, the LRDIMM patents, again makes the statement that Samsung is the one who designs these products. It's not incidental. It's not off the cuff.

We have Samsung employees at slide 18, Mr. Sri (ph) and Mr. Kim, that both say that there are no foundry products here, that the ones that have been accused in this case are Samsung products, which make sense, Your Honor, because these are memory products. These are kind of commodities.

These are the things that Samsung makes and sells widely, and it's -- would be unusual for a company to come to Samsung and say, I want you to make a commodity product based on our design. That would be -- that would be a fairly unusual thing. The documents in this case overwhelmingly show that these are Samsung products made according to Samsung's design for Samsung's purposes and they're sold as Samsung products.

The evidence that Netlist points to is very, very sparse,

Your Honor. They point to this document that I've got up on slide 20, which is Exhibit 5 to Docket 264, and they say, aha, look, you make foundry products. Well, yes, the HBM product is a product, but it is not a foundry product.

So the chart they point to does have some foundry products. There are -- there are processors, although I should say that this document was prospective. This was a business case that could be accomplished. So it was not something that had been accomplished, but it was something that could be accomplished.

But for HBM, because this is a memory product, this is a commodity product; it's just, even in this business case, designated as a Samsung product rather than a foundry product.

The one document they point to over and over again is on slide 21, Exhibit 4 from Docket 264. And they say, aha, look, there's a memory division product, HBM, that's combined with a foundry product.

Again, Your Honor, this is a prospective presentation. This is something that could be done rather than something that has been done. But what they're pointing to are things like an interposer, an additional product that could be added that a customer might want. But that's a different issue. That doesn't tell us that the HBM was a foundry product.

And, critically, they've got to point to something. They got to point to a product somewhere that they tell us, they

stand before you and say, this product, the X28.48 memory product, was a foundry product made for this company according to their specifications. And they haven't been able to do that.

So the termination date. There really -- it's hard for me to understand how this could still be an issue. The JDLA itself gives us the mechanism for termination. Good contract lawyers don't leave the parties to wonder about their relationship as they move through a contract. Good contract lawyers don't leave to chance whether or not a contract is still in existence or not.

And the good contract lawyers that wrote this one provided a specific termination mechanism, one that Judge Scarsi recognized. He recognized that this was a -- this was not an executory covenant subject to -- you know -- or I can't -- I can't quite dredge up all the words from law school.

This was -- this was a -- an existing covenant that could be defeated if you followed the following steps: You had the right to terminate the agreement by written notice. And that 13.2 termination provision was what Judge Scarsi relied on which binds Netlist, which means that none of this should be an issue in this case, because that's exactly what happened.

They sent us a letter on July 15, 2020, and they declared that they were terminating effective immediately, that is, of

that date. On that date, they elected that remedy. They could have sued and said, for past performance things weren't right. They could have made claims. But they elected termination. Election of remedies is an important part of this case.

And they said that we're terminating the agreement as of this date, including the patent license granted to Samsung. They didn't say the patent license that had been previously terminated; they didn't say the patent license that might be injured in some fashion. They said, we're terminating as of this date, which meant that it was in existence until that date.

The JDLA gave us the -- the process and the mechanism for termination. They followed it, and that's what Judge Scarsi ultimately decided on.

In the interrogatory responses in this court in this case just a couple of months ago, Halloween 2022, they put in an interrogatory response saying that the JDLA was terminated on this date, which ended the license granted to Samsung. The language could not have been clearer.

So part of the problem with Netlist's coming up with these new theories as we move through the case is they forget what they've said earlier. They could not have been clearer about this.

They told the California court this, of course, over and

over and over again. At slide 26, I have Exhibit 6 from the complaint in the Central District of California. They said that the licenses and rights previously granted to Samsung ceased on July 20, 2020. They had that wrong. It was July 15, but they were close.

We know that the arguments they made in getting summary judgment, they absolutely emphasized that the JDLA was in force between November 2015 and its termination in July 2020. And they urged the Court to focus on that because that was the extant period of the contract.

Mr. Sheasby told us that, well, it's an executory covenant, and somehow he called it an executory contract. I think he meant executory covenant. I'm not sure about that. I'm not sure that -- and this is an argument that they bubbled up in the reply section of the briefing. I'm not sure that a license is an executory covenant. I think if we went back to law school and we had a debate about this, that executory covenant is something that requires you to take action, not something that requires you to forebear action.

And according to Mr. Sheasby, he says that a license means you are going to be free from suit. So we have no time machine, we have no ability to go back and have them take an action that they had promised to forebear. But we can leave that debate aside because the reality is what they told us, what they told Your Honor, what they told the California

court, is that the JDLA was in force from November 2015 to July 2020. And that's what Judge Scarsi found.

Judge Scarsi found that, when they sent that letter on July 15, 2020, they terminated it, and he specifically found that Netlist followed the rules that were set forth in clause 13.1 of the JDLA. Their termination comported with the JDLA's termination method.

So their appeal--and I don't want to belabor this--repeats that in fact July 15, 2020, was the termination event -- terminating event for the JDLA. They have got this before the Ninth Circuit. They also told the Delaware court the same fact.

Slide 30 recites Exhibit 10 from Docket 255, that the letter of July 15, 2020, was the termination of the JDLA.

And then, finally, Your Honor, in this case, in this complaint, once again they tell us that they terminated the agreement on July 15, 2020.

At some point, at some point their statements to Your Honor, to us, to Judge Scarsi, to Delaware, to the Ninth Circuit, at some point those have to have meaning.

So with that, Your Honor, I'll turn to the preclusion arguments. We think issue preclusion does lie here. We think that Judge Scarsi was presented with this -- this issue. They could have advocated to him that the agreement had terminated much earlier.

Recall that the jury looked at this breach of contract claim in California and awarded zero dollars. Zero dollars. Judge Scarsi said, well, there was a breach; therefore, I have to enter nominal damages. They gave them $1. $1.

But they chose. They chose to ask for termination, they elected their remedies, and whatever damages they thought flowed from that -- that breach, that alleged breach, they were compensated for.

And so it's -- it's not accurate to suggest that somehow Judge Scarsi wasn't presented with this issue. They could have -- they could have made the repudiation argument that they make here, that somehow they were relieved of covenants earlier on and therefore whatever damages flowed from that -- that -- that point forward would have happened. But that's not what they did. They elected July 15, 2020.

The issue was adjudicated, it was fully and fairly litigated, the determination was necessary to judgment, and there was nothing to excuse them. And so the reality is we have two issues: We have whether or not it covers the products--that was before Judge Scarsi--and when the license was terminated--that was also before Judge Scarsi. All four, all four of the issue preclusion elements are present here.

THE COURT: Does it make a difference as to whether it was the identical issue, that that was a declaratory judgment action versus the structure of this action here?

MR. CORDELL: It would for claim preclusion. Your Honor, remember there are two -- we were taught in law school res judicata, and they kind of stopped right there. But now we -- we've refined it. The law has refined it so we have issue preclusion, which is when was the -- when was the contract terminated. It's a factual question. It follows the law, it was resolved, and we have a judgment that depended on it. That's issue preclusion.

Claim preclusion would suggest that no matter what claims the defenses could have been brought, they are now -- they are now barred under claim preclusion because the entirety of the transaction should have been barred. And in that case, it does matter what causes of action were brought.

And so claim preclusion we have from the Comer case in the Fifth Circuit that, in fact, if the same claim or cause of action was involved in both actions. And if you recall from the briefing, Netlist suggested that somehow our defenses were precluded because of this, because of claim preclusion.

Now it seems that Mr. Sheasby is arguing against claim preclusion which, frankly, we agree with, except for the scope of the JDLA. The idea that the JDLA didn't reach these products or somehow wasn't terminated on July 15, 2020, is subject to claim preclusion. Anything else they could have brought, any other remedy they could have sought, is now forestalled.

And then, finally, judicial estoppel, Your Honor. You can't make these repeated statements to courts all over the country and then come back and take them back. You just can't. We certainly can't do that here in Texas.

So with that, I'll -- I'll pass the -- pass the podium, Your Honor, unless you have other questions.

THE COURT: I don't, but I do think it would be effective, and perhaps I should have told you both this before we started, but Samsung has moved under Document 196 for summary judgment that the accused products sold before 2015 are licensed.

I think that's part and parcel of this same issue. It's the other side of the coin, and I want to hear your arguments on that motion in conjunction with the 201 motion that you've been arguing on.

MR. CORDELL: Well, Your Honor, they're really the same issue.

THE COURT: So as long as you're at the podium, take the affirmative side of your motion.

MR. CORDELL: I will.

So we think that they're -- and I won't go back through the whole -- the whole analysis, but the reality is if you recall the language that I took Your Honor through on the JDLA is that the products were licensed. There's no dispute about, you know, whether the LRDIMM, you know, DDR4 versus DDR5, they

never raised any distinction there. Those -- those products are licensed if, in fact, the license was terminated on July 15, 2020. I don't think there's a dispute about that.

The only dispute is the foundry products. And, again, if they can -- they can point us to a single accused product that is a foundry product, then I'll sit down. But they haven't done that. What their experts say is that all of the accused products were designed by Samsung, for Samsung, sold as Samsung products, which means that they're licensed.

So there's really nothing left other than Mr. Sheasby's repudiation argument that somehow earlier there was a termination of the license. But the law doesn't permit We've seen that from issue preclusion, from claim preclusion.

But to me the judicial estoppel is the most profound, the fact that they told us in the complaint in this case, in discovery just a couple of months ago, told Judge Scarsi, told the Delaware court, told the ITC, they have -- they have represented over and over again to virtually every judicial body they could find--the Ninth Circuit I left out--that in fact that license persisted until July 15, 2020, and they have to live up to that.

So with that, I think summary judgment does follow.

THE COURT: All right. Mr. Sheasby, let me hear from you.

MR. SHEASBY: Just briefly, Your Honor.

May I have my slides, Mr. Huynh? And if I can have slide 43.

THE COURT: And, of course, I want you to address your response to Samsung's motion for summary judgment, Document 196, as a part of this.

MR. SHEASBY: Yes, Your Honor.

So I want to start off with just some importance on precision, which is in the interrogatory response. When they asked us for the first time what is our view of the JDLA, we expressly said that the JDLA was repudiated at the time of substantial non-performance. That was in our interrogatory response.

We also said that the JDLA excluded foundry products, and we also said that they were bound -- Samsung was bound by the decision of the KTT. So just as a matter of precision, the idea that this is a surprise or that this came out in -- in a reply brief on a motion just has no connection to reality.

Mr. Cordell actually showed a portion of this document, but omitted to show the portion--this is our Interrogatory Response No. 16--that made clear our position on all these issues.

If I could go back to slide 4, Mr. Huynh.

There are really three issues, each of which are disputed questions of fact, two of which are disputed questions of fact, one is an issue of comity, all of which we think require

denial of their motion and the grant of our motion.

The first is that these language -- this election of remedies language that Mr. Cordell used, there's no time machine that goes back. That is just not New York law. The parties agreed to New York law.

Under New York law, there are two mechanisms. There's a contractual mechanism of termination which is exercised in a particular date. That termination did not just apply to whatever license grant we had given them. It applied to a host of obligations that both parties had, all of which ended on that date.

Under New York law, there's a second issue, which is, at what point in time do you have to -- are you entitled to cease performing on an executory contract. And the point of time is when there's substantial non-performance. That is a question of law under New York law, and the idea --

THE COURT: Why is this an executory contract?

MR. SHEASBY: So the Ninth Circuit speaks to that in *In re CFLC*. So the contract at issue is a contract that requires us to forego bringing a lawsuit and seeking patent damages against them. And so under settled law -- and the Federal Circuit views that as well. And so I should say that the executory aspect of it is something that I'm adding. The --

THE COURT: I don't know why the Ninth Circuit

ruling here would be enforceable precedent in this court.

MR. SHEASBY: I don't think it is. So two things. One, the Mega Group and the *Merrill Lynch* case make no reference to executory. I'm using that as just a handle, which is, our performance is excused at the time of substantial non-performance. Our performance is not to seek patent damages against them.

We have a right to seek patent damages against them for six years before suit that we forewent based on this license agreement. We did not have to perform that at the time of substantial non-performance. And so we are entitled to bring a suit for those damages that we otherwise forewent based on the license agreement.

So I think executory versus non-executory is probably the wrong question. It's the *Merrill Lynch* case and the *In re Mega* case which both speak about the specific issues. They don't use the word executory. So I may have gotten us off on a detour by talking about that. It's that we have a right to have six years of damages from the time of filing a lawsuit.

We committed not to seek damages under the license agreement. They lost the right to rely on that commitment when there was a substantial non-performance. The substantial non-performance occurred on a particular date. From that date forward, we are entitled to seek damages. That's what the Second Circuit and the Interim Court of Appeals decides.

THE COURT: And your view of the date by which Samsung substantially failed to meet its obligations that you believe gives rise to that damages period in your mind, that is what date?

MR. SHEASBY: This is -- if you go to slide 8.

We believe it's on -- in January of 2018, there was a lunch meeting. And at that lunch meeting, they said, you get zero more, you get nothing from us. And we told them that was not allowed under the contract, and then we tried to negotiate with them.

I should say that's the -- the first issue.

The second issue is that we've taken the opinion that they -- taken the position that they had a right to our -- to sort of be relieved from infringement. I want you to look at this date. It's January 2018 is when they stopped. February 18th is the date of this email.

But if you go to slide 38, the pleadings they're referring to were before the period of time in which they shut off the spigot. So this is an ITC proceeding from October of 2017 when we were still hustling with them to get some performance and some product. It was before the period of time when they substantially shut off their performance.

That's the substantial non-performance issue, which is a question of fact under New York law.

The second issue is the foundry issue. If we go to slide

10, foundry products are not licensed if they are based on designs, specifications, or working drawings owned by such third party. It doesn't say based solely on. It doesn't say it has to be just a tape they receive. And we believe that there is substantial evidence in this case that the high bandwidth memory products are foundry products.

First off, they describe their high bandwidth memory products as a foundry product in their pleadings -- sorry, in their publications. They describe that they are using foundry services for each of the major purchasers of the HBM products. AND, Xylex, Microsoft, Amazon, Intel, nVidia are the major purchasers of the HBM products which we contend are foundry products.

And I think Mr. Cordell invited me to give an example, one example of a product that is a foundry product, and he would sit down. This is Exhibit 24 to our opposition at 1, in which they are specifically designing HBM3 products for nVidia and are making the changes that nVidia instructs so they can be custom for nVidia.

THE COURT: Let me ask this question, Mr. Sheasby. And I don't mean to get it out of order with your argument, but there are so many moving parts here, if I don't ask it, I may not remember it.

Tell me why, after Netlist concluded in January of 2018 that Samsung had materially breached its obligation to perform

by refusing what you requested or what Netlist requested, tell me why two years effectively evolved between then and July 2020, and what was it that delayed the 2020 action?

And are you alleging that somehow Samsung strung you along with promises that they would do better to drag out the time period, or did you reach the conclusion, or Netlist, rather, reach the conclusion that Samsung had materially failed to perform and you just took your own sweet time about doing what you did in 2020?

I need some context as to that two-year gap and who's responsible for the expiration of that time.

MR. SHEASBY: Sure, I'm happy to explain it.

THE COURT: I'm sure I may hear a different story from Mr. Cordell, but what's your version of that?

MR. SHEASBY: It was called begging. So the substantial portion of this company's business -- and, in fact, I can show you the slides if you want me to go to another deck -- when they cut off our royalty stream and it's -- the slide -- it's from Steve Metz to another Samsung employee. Samsung was effectively a hundred percent of our revenue, was us using Samsung product. It was existential for our business.

There was a joint development relationship relating to substantial product associated with this, and we begged them for years and -- for over a year-and-a-half to do something

about this.

In addition, what they did was something clever in the sense that there was an internal email in which they decided that they would give us $500,000 of product over a certain period of time. And the email actually said, LOL, laugh out loud, we're going to give them $500,000 of product.

So what they did is they kept giving a little bit of product, allowing us to hope that we would be able to repair this relationship. The relationship was not reparable. It became clear there was never going to be a hope of this achieving again.

In other words, this is not just a license agreement. This was the future of Netlist. Right? This was a JDLA in which there was going to be a joint development of a product design, they had a right of first refusal on our product design, there was going to be exchange of technology. This was a robust agreement, the license of which was only one part of it.

And so I can show you the slides, but that's the short answer.

THE COURT: Short answer, your view is that they strung you along.

MR. SHEASBY: They strung us along.

THE COURT: Okay.

MR. SHEASBY: The short answer is that under New

York law, stringing along is not a defense. In other words, from the moment of substantial non-performance under New York law, our obligation ceased.

THE COURT: All right. What other argument do you have for me?

MR. SHEASBY: And the last issue is -- so we talked about the foundry issue, and Mr. Cordell invited me to show him the evidence. I've shown him the evidence.

The next issue is, what do you do about the Korean Tax Tribunal decision that the license that Samsung received was limited to products that were part of the joint development relationship. This is only relevant to the HBM products. The HBM products were the only products for which patents existed before the termination date.

If you go to slide 21, this is where the Fifth Circuit instructs you to rely on comity. Was it rendered by a court of competent jurisdiction? Was there some type of due process? Did Samsung have the opportunity to be heard and were procedural formalities followed? And was there a formal renewal?

This is the decision of the Korean Tax Tribunal making clear that the agreement was limited to licenses under the joint development relationship. HBM was not part of the joint development relationship.

Samsung's representation that it was not a party to this

proceeding or did not have its right to be heard is entirely inconsistent with both the KTT decision, which stated that the -- that they were involved and submitted proceedings.

And if you look at Judge Scarsi's decision on the subject--this is Exhibit 22 to our opposition to their motion for summary judgment--Judge Scarsi actually noted that the Korean Tax Tribunal was a venue in which Samsung made, quote, arguments adverse to Netlist's interests. And Judge Scarsi actually adopted the Korean Tax Tribunal's decision.

I suggest that there -- I think -- I've dealt with claim preclusion. There's no claim preclusion because it's a different claim. The claim in that case was a declaratory judgment of whether the termination was effective.

The issue in this case is a response to their raising the license agreement as a defense. But if you go to slide 31, the idea that there was judicial estoppel or we told anyone anything different is just factually incorrect, and Judge Scarsi noted that as well.

We've always taken a position that the JDLA is limited to the joint development and research. He just said that was irrelevant to his decision. It was not relevant to his decision what the scope of the license are.

So four issues: substantial non-performance, foundry products, and whether they're bound by the KTT decision, and then the fourth issue is there's no claim or issue preclusion

or judicial estoppel.

Thank you, Your Honor.

THE COURT: Do you have anything else on Documents 201 or 196?

MR. SHEASBY: Nothing, Your Honor.

THE COURT: All right. Mr. Cordell, I'll give you the last word on these, and then we'll move on.

MR. CORDELL: Thank you, Your Honor.

So let's start with the substantial non-performance issue.

Could I have slide 25? Sorry.

So Mr. Sheasby said, Wait a minute, we told them in interrogatory responses that we had other issues.

But, Your Honor, they could not have been clearer. They could not have been clearer. I put slide 25 back up, which is their response to Interrogatory 16, on July 15, 2020, Netlist sent Samsung a formal written notice terminating the JDLA which ended any license granted to Samsung.

Netlist needs to be true to its word. They can't tell us that early in the case and then later say, well, you know, maybe we really didn't mean that, maybe we thought it was repudiated much earlier, and that the license granted to Samsung somehow expired earlier.

Mr. Sheasby, once again, he seems to be retreating from the executory covenant argument, and now shows us New York law

suggesting that you might be excused from performance if someone commitments a material breach. That is an unremarkable proposition. That's one that we were all taught that can happen, but that has to happen in the motion. That's an extant excuse.

So had they sued us in January of 2018, they could have pled that. That might make sense. But you can't go back now and say, well, you know, we could have been relieved of our burdens, we could have abandoned the contract and all of its benefits back in January of 2018, but they didn't. They waited until July 15, 2020. But there's more.

THE COURT: What's your version of why the period of time from '18 to '20 transpired?

MR. CORDELL: So my version is that it was Samsung's position that Netlist was abusing the supply part of the contract; that, in fact, the parties had understood that there would be some parts that were provided primarily for this joint development project, but there may have been some others.

But beyond that, there was a general memory shortage, and it was Samsung's view that Netlist was taking advantage of that situation; that they were -- they were depriving other Samsung customers of parts so that they could be the resellers and reap some profits because of the memory shortage. That's our view of what happened.

But, importantly, Your Honor, we have to go back to the contract, and the contract provides for this. Again, good contract lawyers don't allow the parties to drift around at sea while they -- they decide whether there's been a termination or not. Good contract lawyers will -- will provide for it.

So we go back to section 13, term and termination. And section 13 tells us what happens. They tell us what you have to do to terminate the contract. They tell us you got to give written notice. And, importantly, there's a cure provision. So you got to give 30 days for the other side to cure whatever breach you allege happened.

So if in January of 2018 they believed that there was a supply disruption or something that caused a breach of the contract, the contract says Netlist should have sent us a letter saying, hey, we're not getting the parts that we need, you need to give us some more parts, or whatever they felt like the breach was.

So the cure provision is there in black and white, and it is only if it is not cured within 30 days may they terminate the agreement or declare a breach. That's what section 13.2.1 says. But there's more.

Then in 13.3, we have effective termination. And it tells us what happens to the license that we've been here arguing about for the last hour. What happens to the license

is that it terminates as of the effective date of the termination. There is no contractual right to this prior repudiation. There is no ability to go back in time and say, well, maybe we're going to pull the license back sometime earlier.

They agreed to this in black and white, that the rights granted to the defaulting party -- they say Samsung defaulted. Samsung's rights will cease forthwith as of the effective date of such termination. Not before, not when they thought there might be a problem, not when they thought, well, maybe we should write them a letter, but as of the effective date of the termination.

Judge Scarsi determined the effective date. They admitted that's the effective date. This has been an enormous expense of the Court's resources for no purpose.

Let me address the foundry issue. So I did challenge Mr. Sheasby to show us an accused product, not any product but an accused product that qualified as a foundry product. What did he show us? Well, he showed us a slide that said that Samsung does, in fact, have a foundry business. And they do. They will make parts for other companies when they have time. That is true. That doesn't tell us that there are any parts accused in this case that were so made.

He then showed us this slide and said, Aha, look, HBM -- some of the customers of HBM are listed on this slide.

Well, they are, Your Honor. But the fact is there are several businesses mentioned on this slide. There is foundry services; that is true.

I can zoom in on that.

But there's also memory, a completely different business or platform services. The fact is that Samsung does a lot of things for a lot of companies. Again, I'm waiting to see an accused product that actually qualifies under the definition of a foundry product.

And then his -- his big gun is this slide where he says -- let me get my notes -- he says, Aha, I found one, there we go. This HBM3 that's provided to nVidia, that's -- there you go. That's -- you wanted a product, he gave us a product.

Well, number one, Your Honor, that's not a product. Okay? There is no accused product called HBM3. That's a class of products. That's like saying, you know, a four-door sedan. That's not a, you know, Ford, I guess, Taurus would be a four-door sedan. It's hard for me to think of one. That would be a product. This is not a product. That's a class of products.

But, critically, critically, HBM3 was not introduced until 2021 at the earliest. That whole class of products didn't exist during the effective time period of the license. So not only do they not tell us a product, this class of

products can't qualify. It can't be an accused product because it didn't exist at the time the license was there.

So he's got to show us an accused product or he loses the foundry issue. You can't just rely on their evidence.

And, finally, Your Honor, with respect to the Korean tax authority, you have to read this very carefully. I think Mr. Sheasby may have misspoken because he seemed to suggest that somehow the -- that -- that Samsung was making arguments, that Samsung's arguments were being considered or rejected. But that's not true. These were Netlist's arguments.

Let me find the right slide that he showed us. Well, he did show us a slide. I don't think this is the right one. Yes, actually it is the right one, Your Honor, the top -- the top quote.

It talks about Netlist's -- though Samsung referred Netlist to appeal directly to the Korean tax authority, and its arguments before the tax tribunal were adverse to Netlist's interests. The its there are Netlist's arguments, not Samsung's arguments. Samsung was not a party to the Korean Tax Tribunal. The way it works is Samsung withholds the money; they then go to the Korean Tax Tribunal and try to get it back.

Samsung gives them some information, but they were never a party to that proceeding and they can't be bound by it even if it had any relevance here. This is not a 44.1

determination, and it has none.

So with that, unless there are questions, I will sit down.

THE COURT: All right. Thank you.

Well, I've heard considerable argument on these two competing summary judgment motions, and I felt it was important, given their impact and significance.

Considering the briefing and considering the argument that's presented today, I'm going to deny Plaintiff's motion for partial summary judgment. I don't think it's supported by Rule 56. I don't think the Court has a basis to grant it.

But I'm going to grant Defendants' motion for summary judgment. I think it's replete in the record time and time again, both from the terms of the JDLA itself and the affirmative statements made by Netlist, that until July 15th, 2020, Samsung had a license defense, and after July 15th, 2020, Samsung has no license defense.

And there just does not seem to me to be any material question of fact with regard to that issue. And the repeated affirmative statements by Netlist to various listeners and recipients that the license terminated on July 15th, 2020, just seems to me to be beyond any material dispute.

And I'm going to grant Defendants' motion for summary judgment, Document 196.

Now, we've been in here over an hour-and-a-half. We're

going to take a short recess. I'm aware that the ruling I've just given you may impact some of the other motions that are before the Court. I want you to meet and confer during the recess about that, and I'll hear input from you on that when I return.

Mr. Sheasby, you're on your feet walking to the podium. I assume you have a question.

MR. SHEASBY: Yes. So did you -- did the Court grant summary judgment that the HBM products are foundry products as well?

THE COURT: No. I don't think there's any fact issue there, either.

MR. SHEASBY: Okay. Thank you, Your Honor.

THE COURT: All right. Meet and confer on where this takes us after this ruling, and I'll hear from you on that when I return. We'll take a short recess at this time.

Court stands in recess.

(Brief recess.)

THE COURT: Be seated, please.

Counsel, two things. Number one, I've been reminded that during my housekeeping instructions to you, I didn't expressly say that you each had 30 minutes total to address the venire panel during jury selection. That's what you have--30 minutes per side. Nobody asked any questions about it because I assumed that you already knew that, but I'll just make it

completely clear that's what my expectation is with regard to voir dire.

Secondly, as I've had happen before, Mr. Sheasby got up and asked a question as I was about to get off the bench. I won't say that I was completely ready for it. I thought about the question he asked just before I recessed. I want to clarify the Court's ruling.

When I said no to his question, that is correct, and that is what I intended to say--that there are fact issues with respect to whether the HBM products are foundry products, which would determine whether they did or did not fall within the license that existed up until July 15th, 2020.

So I'm not granting summary judgment. There are fact issues at play there. So that is a live issue for the jury and an issue to be presented during the trial.

But to the extent they do fall within the license, the license is in place up until January 15, 2020, under the JDLA, and that's the basis upon which I'm granting Samsung's motion.

Both sides have conceded on the record that as of -- both sides have conceded on the record that after July 15th, 2020, there is no license defense and the JDLA was terminated, which terminated the license. So that's I think a complete review of what I've given you so far.

If anybody has any questions about any of those issues, now's the time to ask before we move on.

MR. SHEASBY: Nothing from Plaintiff, Your Honor.

THE COURT: Is that clear to your side, Mr. Cordell?

MR. CORDELL: It is, Your Honor. But I guess I'm a little curious as to what the Court's expectation is on the HBM issue. It turns out that that is the -- that is the gravamen of this motion, that the only two patents-in-suit that existed as of the time of the license are the HBM patents.

THE COURT: Well, I can tell you I'm persuaded, based on the evidence that's been presented, that there's no material question of fact that the termination letter executed by Netlist on July the 15th, 2020, effectively terminated your licenses -- Samsung's licenses as of that date.

But that does not leave open the look-back argument that Plaintiffs presented that somehow in 2018 when they say Samsung failed to substantially perform, that that somehow suspended, repudiated, or otherwise took away the license protection from the JDLA.

So you've got a license -- Samsung's got a license defense up until July 15th, 2020, and you don't have one after July 15th, 2020. What falls within that license as prescribed by the other terms of the JDLA is a fact question, and I don't find as a matter of summary judgment that the HBM products are foundry products and either would be within or without the license. That's a fact issue that the jury's going to have to

decide.

MR. CORDELL: Understood, Your Honor. Thank you.

THE COURT: Okay. All right. Let's move on to Samsung's motion for summary judgment of non-infringement as to the '060 and the '160 Patents. This is Document 197. Let me hear from Samsung as the moving party here first.

MR. McKEON: Good morning, Your Honor. Mike McKeon for Samsung.

THE COURT: Good morning, Mr. McKeon.

MR. McKEON: I previously handed up the slides, and you should have them. May I proceed?

THE COURT: You may proceed.

MR. McKEON: So, Your Honor, we strongly believe -- and we were just talking about the HBM products and the two relevant patents for that. We strongly believe that based on the undisputed facts we have here, that summary judgment of non-infringement is warranted.

THE COURT: Does strongly believe somehow equates to strenuously object?

MR. McKEON: Well, I guess so, Your Honor.

THE COURT: Okay.

MR. McKEON: But really actually, and before I jump to the slides, I think the procedural history how we got here is actually quite telling. You know, of course, we had the *Markman* process before Judge Payne, and you know we argued

during that process one of the key terms was a term called array die and all the independent claims of the patent.

And we argued, based on the prosecution history, that there was real clear prosecution history disclaimer with respect to that term of array die. And, of course, as you would expect, Netlist opposed that and vigorously argued that, no, there was no disclaimer, and if there was disclaimer, it was a really limited disclaimer.

And then we got the ruling from Judge Payne, and he agreed with our position 100 percent and he adopted this construction based on the disclaimer and the array die construction was array die that is different from DRAM circuit.

Then Netlist went and filed objections to Your Honor to review Judge Payne's analysis and Judge Payne's ruling. And, you know, we actually thought that they would concede the -- the non-infringement point and take a judgment and get that out of the case. That's how strongly we felt about this. But they moved on, and they filed the objections.

So we went through the expert report process. And then after the expert report process, they withdrew the objections. And those objections made the exact same arguments that they made before Judge Payne that Judge Payne rejected. And then when we got the expert reports, and we were going through them, the expert reports made the same arguments, saying that

there was infringement. And they were the arguments that Judge Payne rejected.

And what they want to do here, Your Honor, they want to go to the jury on this. This is clearly what they are doing. They want to go to the jury and talk about the prosecution history, talk about disclaimer, because they want to limit the disclaimer in a way that they tried to convince Judge Payne to do.

And we just think, Your Honor, you know, they get an A for persistence and temerity, but they don't get the same grade for substance in this issue. And for that, Your Honor, I just want to just jump in. Of course, these are the relevant docket numbers for the briefing on this motion. And I structured the presentation around these fundamental buckets. So let me jump in here.

You know, we really don't think there's a fact issue here. There's not any fact dispute, certainly not a material dispute, based on the presentation and the papers. And just to set the issue up, Your Honor, what I have on the left is this key prior art reference to Rajan that you saw on the briefing. And this was the subject of a disclaimer in the Patent Office.

You notice, Your Honor, the DRAM that's highlighted here, that's the DRAM in the Rajan reference. And the reason why this is an issue for -- on the infringement side of the case

because if you look at the right-hand side, I have a graphic of Samsung's DRAM. This is used in the HBM product. And in the briefing, part of it referred to as the core. So you have the DRAM core die.

And even on the slide, you see here it says HBM DRAM die, and it looks just like what's in Rajan in its DRAM, the stack. And that's why there's a problem in their infringement case is because they disclaim -- they disclaim what's on the left in the Patent Office.

And in the beginning of case, there was no question what their allegations were, you know, even in the preliminary contentions. But throughout the case, the -- the stack that I just showed you, there was eight stacks of DRAM array dies. That was the accusations. So when you came to the claim element array die, they were pointing to the eight stacks, and that was their case.

And -- and it was clear, in case there's any ambiguity on this, that when they say die, when we say die, what do we mean? Everyone refers to this as circuits, one or more circuit. And this is from their claim construction briefing. And they're clear about one of ordinary skill in the art, yeah, of course, the term die is a circuit.

So to the extent there's some dispute whether our core dies are circuits or not, DRAM circuits or not, of course they're circuits because it wouldn't work if it wasn't a

circuit. Nothing would flow. It would be a brick. And they've even taken the position in the briefing before Judge Payne, one or more circuits. So we've got DRAM circuits in the accused product. That's what they've accused.

But then we get the expert reports, and this is an excerpt from their expert here. And they -- they characterize this Rajan reference as, you know, wire bonding. And I'll get into that in a minute. And what they say here is the HBM products that they've accused, this is Doctor Brogioli, do not have eight DRAM circuits--this is the key--like those in Rajan.

And if you look through his report, you flip through it, the entire analysis is the HBM products don't have DRAM circuits like those in Rajan. Judge Payne rejected that argument. And they withdrew the objection to Your Honor on that very argument trying to limit DRAM circuits to like those in Rajan. And that's all the expert did in the analysis. The expert never said anything about DRAM circuits' plain and ordinary meaning. And we'll get to that, Your Honor.

In terms of resolving motions, this is the list -- if our motion is granted, you know, we think it should be, where we would resolve -- there's an error in the top one. As we've just argued, there was a segment that is prior to the July 15, 2020 date, but this would resolve a healthy set of motions in this complex case for these HBM products.

So the patents, Your Honor, just briefly, the array die term, just to give you context on that, the yellow that we've highlighted here, those are the array dies of the patent, and there's -- there's different elements about how these are all connected.

But for the purposes of this motion, Your Honor, you just see that array dies are stacked on top of each other in the patent disclosure.

And, of course, Rajan is the key prior art reference that was at issue in the Patent Office. You see the same stack there. These are DRAM stack.

And then the issue, Your Honor, that you've seen in the briefing is they say, Oh, Rajan's got wire bonding.

And when they say that, they are pointing to these control signals on the left and they are pointing to these data signals on the right, and they say that's wire bonding. So when Rajan is talking about DRAM circuits, they are talking about DRAM circuits connected through wire bonding.

But, first of all, we disagree with that characterization because there's no disclosure that they're connected in that way in the specification, but the -- the DRAM circuit disavowal, disclaimer, was limited to what's highlighted here, the DRAM.

And then, finally, looking at the specification, again this is from the prior art. This was even cited by Netlist in

the *Markman* process. The memory circuits that they're talking about here with this excerpt from figure 1 are identified and highlighted in yellow, correspond to the DRAM that we have in figure 2A.

And Rajan is very clear, you take any form, any form, you know, not limit it to whether they are wire bonding or whether they are TSVs that connect this, no limitation in that regard. And, of course, they call them memory circuits. So they refer to the DRAM as memory circuits. And everyone skilled in the art, of course, knows these are memory circuits.

So the *Markman* process, Your Honor, was a robust process. The big debate we had, of course, and this is the key excerpt that Judge Payne relied on, what they told the Patent Office, Rajan does not disclose a plurality of stacked array dies. Merely -- Rajan merely stacks DRAM circuits 206A-D, which are different from array die.

Notice how they're just pointing to the DRAM. They're not pointing to the address circuits or the -- the data circuits where they claim is wire bonding. They're only pointing to the DRAM here.

And, of course, this is the public notice issue. I'm not going to repeat the law. You know it very well. You know, prosecution disclaimer is based on the public being able to rely on -- rely on these statements to the Patent Office. And Judge Payne, you know, points that out as well.

And, of course, they -- they opposed the disclaimer argument. And they first say, well, there's no disclaimer at all, as you see here on slide 20. This is the argument they made to Judge Payne--the inventor did not disclaim DRAM circuits. That was their point.

And then they go on and say, if there is a disclaimer, if there is one, it should just be the circuits of Rajan. It's Rajan, how Rajan wires it up. This is what they told Judge Payne. And they continued on, Your Honor, with citing that part of Rajan where you can have any kind of DRAM. Well, you know, it helped them at that point in the case.

And so we had a statement here that they rely on where at the end of the argument Mr. Sheasby got up and asked the Court, well, whether the Court's ruling if -- the Court's ruling is a tentative, as you know, and if it's tentative the Court sided with Samsung, and he asked Judge Payne whether he intended to give guidance on what DRAM circuits are excluded. In his opinion -- and Judge Payne said, Well, that's for the experts to talk about DRAM circuits. And that was his guidance.

So, sure enough, on the ruling, he ruled in favor of Samsung, adopted precisely our -- our proposal which was consistent with the disclaimer. And we have that on slide 24.

And then in the opinion, you know, he -- in analyzing the prosecution history, he says the disclaimer's broad because

Netlist did not explain that structural difference between any array dies in DRAM circuits or even how they might be stacked differently.

So his -- he had a broad -- his ruling here was a very broad disclaimer with respect to DRAM circuits, not how they're hooked up or how they're stacked. He rejected that because in the Patent Office, it was a broad statement and they didn't explain to the examiner it was narrow, the distinction. And then he rejected the structural distinction argument. Judge Payne.

And then, finally, you know, he -- he makes the point to Netlist that to the extent the disclaimer might create later difficulty in proving infringement, that's not a basis to avoid applying clear and unambiguous disavowal. And, of course, that's -- that's certainly the law, and that's the public notice function.

So then, Your Honor, December 28th, we have the -- the objections that were filed. And if you read the objections, Your Honor -- I won't go through it in detail here, but the same arguments they made to Judge Payne, they are urging the Court to reject Judge Payne's narrow -- or Judge Payne's, I should say, broad disclaimer and wanted to urge the Court to make it a narrow disclaimer if the Court ruled that there in fact was a disclaimer.

And we see here in slide 30, this is their argument. As

a result, the Court should make clear that the claimed array die is one that is different from Rajan's DRAM circuits 206A-D.

So that's how they're setting this up. They urge the judge -- Your Honor, in reviewing Judge Payne's decision, to say it's limited to Rajan's DRAM circuits. And, of course, this is now the infringement case they want to present to the jury.

And they -- they pulled this objection, Your Honor, after the expert reports were filed, they filed a motion with Your Honor on February 2nd withdrawing their objection. And, of course, we know why they withdrew it because if Your Honor sustained Judge Payne's analysis, they're going to be done and they wanted -- they want to argue this in front of the jury. That's the analysis, and that's what happened in the *Markman* process, and that's what brought us here today.

So, of course, the tune changed. Prior to this ruling, they've accused DRAM -- the DRAMs in Samsung's HBM stack. Now, after the ruling, what they're saying is, oh, the DRAMs and Samsung's HBM stack are not the DRAM circuits of Rajan. That's their entire infringement report. They're not the DRAM circuits of Rajan.

They changed. They never mentioned, of course, that in any of the filings or submissions or contentions prior to the ruling from Judge Payne.

Now, this is a quote from the expert. The expert, you know, uses the statement, but Netlist did not explain the structural difference between array dies and DRAM circuits or even how they might be stacked differently.

And he -- Judge -- sorry, Doctor Brogioli in his expert report says, well, I understand what the Court's saying here is the structural difference between array dies and DRAM circuits or even how they might be stacked differently. This is an issue for -- somehow for the experts.

And he goes on, and this is again from Doctor Brogioli, he says the rationale for the prosecution argument presented by Netlist had nothing to do with the nature of the memory cell employed by Rajan. In other words, he's saying in his expert reports that this disclaimer wasn't broad, wasn't a broad disclaimer like Judge Payne held.

And then if you read his expert reports, he goes on about Rajan and how Rajan -- Samsung's misquoting, and it wasn't this broad disclaimer, it was focused on the Rajan -- the circuits of Rajan. We have here in slide 36 just quotes from his report.

And then, you know, again his conclusion is that, well, Rajan had the wire bonding. And if you look here on slide 37, what he says here is, the HBM products do not have DRAM circuits like those in Rajan. And that's -- that's his expert report, that's his analysis. He is linking it up to Rajan.

And they go through and they cite the deposition testimony. And during the deposition, it's a little unusual during deposition to have the technical experts -- the technical witnesses be presented with prior art. But they -- they put the Rajan reference in front of our technical witnesses and said, do your HBM circuits look like Rajan. This is what the question they asked.

And, you know, the technical witnesses will look at it and look at some of the differences and, you know, would testify that there's some differences between what we do in our HBM products and what Rajan has, the way they wired it up. But that has nothing to do with the real question, which is the DRAM circuits that are -- that are in our products under the plain and ordinary meaning, are those DRAM circuits that they have. And, of course, they are. Nothing to do with Rajan.

And, again, more evidence about Rajan. It was a very healthy topic I should say in the depositions of our witnesses, what Rajan disclosed and how it was different from -- different from our HBM products, which is not the right question.

And, you know, here in slide 40, you know, Doctor Brogioli reiterates the point, Many differences between HBM core dies and traditional DDR DRAM dies such as those in Rajan. And that's -- that's the basis of their analysis.

And, of course, in the briefing before Your Honor in the summary judgment briefing, you saw that they are just repeating this. And on slide 41, we have excerpts from the briefing. They are focusing on Rajan. They are focusing on the distinctions that they believe they made in the Patent Office and why the disclaimer shouldn't be so broad, which was -- which was the objections that they withdrew from Your Honor.

It was exactly the objections that they withdrew from Your Honor. They didn't want Your Honor to deal with this question. They want to present it to the jury.

And again, slide 42, the same -- same thing. They keep pointing to the DRAM circuits of Rajan, thinking that the -- the disclaimer was limited to that. And certainly it is not.

So let's get to the meat of it, Your Honor. So the other fact disputes here, and we don't think it is -- I'm not going to go through the case law. Your Honor is well aware of it. This is not case law driven issue, frankly. There's no fact disputes. And Your Honor knows Rule 56. No fact disputes, no material publication disputes, require that supplementary judgment be granted.

This is an interesting slide because this is an excerpt of a brief that Netlist filed in another motion that Your Honor's going to consider later today, and they remind Your Honor that the Federal Circuit has made this very clear, this

law about claim construction: It's improper to argue claim construction to the jury because the risk of confusing the jury is high when experts opine on claim construction.

And both these cases are about prosecution history disclaimer. And the Federal Circuit says you can't be -- that's something you do at the *Markman* process. You don't do that in front of the jury.

That's exactly what they intend to do. They intend to go through the prosecution history and prove a case based on the fact that our HBM circuits are not like the circuits of Rajan.

So what do we have here? Brogioli, that's what he does. You know, he argues the wrong issue. He doesn't give an opinion on the plain and ordinary meaning of DRAM circuit. He don't doesn't give you anything on that. He focuses on the disclaimer in the Patent Office as his argument of why there's infringement here.

In fact, we argued here what is the plain and ordinary meaning of the term DRAM circuit. We asked him that question, and he didn't recall giving one, giving a definition to that. And, of course, he can't because we know that DRAM circuit plain and ordinary meaning, you know, this is the -- these are parts of the DRAM and they're all connected through circuitry. And as -- as Netlist already conceded that a die, DRAM die has one or more circuits in it by definition.

And, you know, he had trouble articulating, frankly, the

word circuit, knowing that, you know, apparently it was an issue he was trying to avoid answering, but there's obviously a dispute about a DRAM circuit's clear, plain and ordinary meaning, and he didn't want do give an opinion on that. He refused to.

THE COURT: What else?

MR. McKEON: Okay, Your Honor. I'll just -- you know, the evidence here on the -- prior to the *Markman* ruling, it's all about these four or more stacked dies, the eight dies. This was their allegation throughout the entire case, and they've, you know, changed their tune now.

You know, there's cited -- throughout the entire case, they cited our documents that talked about the DRAMs and the DRAM dies within the HBM core, and that was their infringement case until, of course, the ruling.

So with that, Your Honor, we think summary judgment is warranted here.

THE COURT: Thank you, Counsel.

Let me hear a response from the Plaintiff, please.

MR. SHEASBY: May I approach, Your Honor?

THE COURT: You may.

MR. SHEASBY: Mr. Huynh, may we have deck 6, slide 22, please?

The claim at issues recite a stacked array die, and those array dies have two particular features. One, they host data

ports; and, two, they host interconnects which are, in fact, channels that exist within the array die.

So to give you a context, this is the array die of the patent. This is the chip. And within it, there's either a bypass or a data point, and then there's these TSVs that run through it. So that's what the claim requires.

Judge Payne construed it as, whatever it is, it cannot be a DRAM circuit, and he concluded that the -- it would be a question of fact as to whether what Samsung had is a DRAM circuit.

If you look at the testimony of Samsung witnesses--let's go to slide 25--the chips with TSV, the president of Samsung memory who actually launched the HBM product, testifies that they -- no engineer would refer to them as a DRAM circuit.

The corporate representative of Samsung on their HBM products was shown the same thing. The distinction of their design that's accused of infringement, which has TSVs and data ports in it, and he said that is not a DRAM circuit. And he expressly says that the DRAM circuit, which is on the left, and he accepts that as a DRAM circuit, is not the same as the TSV connections that are in the product at issue in this case.

There's more. They had another corporate representative on the issue of HBMs and asked if the DRAM dies in HBMs are interchangeable with DRAM circuits. And their corporate representative, Mr. Kim -- Doctor Kim, testified that they

were not.

Mr. McAlexander, who is their expert witness, who actually designed DRAM, was emphatic about this, that DRAM circuits are not the same as the chips that are used in HBM which used through-silicon via. And he testified -- let me go back one. He testified that through-silicon via, which is claimed in the patent, is entirely different from what's in the HBM products. There's clear disputes of fact on all these issues.

In terms of the evidence, there's this argument, this sleight of hand, that's what -- in which he indicates -- slide 35, please, Mr. Huynh.

Mr. Huynh, slide 35.

-- in which he indicates that there is not a single Samsung product sheet that refers to what's in their product as a data sheet -- as a DRAM circuit.

Slide 36, there's not a single DRAM Samsung internal document that refers to what's in their HBM product as a DRAM circuit. And there's not a single engineer who has ever described what they make as DRAM circuits in the HBM products. And DRAM circuit is not used by JEDEC to define what's in the HBM products.

There are DRAM dies in the HBM products. Those DRAM dies are not DRAM circuits. But there is overwhelming evidence to the following: There's been no testimony by any Samsung

engineer or any Netlist expert that they have DRAM circuits in them. The only party who believes there is DRAM circuits in their HBMs are -- is Doctor Robins, who was paid to take that position by Samsung.

The idea that Doctor Brogioli is focusing on comparing the HBM products to Rajan is incorrect. He does discuss Rajan in case the Court does allow some of that to go in in terms of a comparison of Rajan. But what he does more generally is appropriately explains why the array die products, which have data ports and TSVs, are not DRAM circuits. DRAM circuits do not have those features of having data ports and TSVs.

He discusses at length the -- the next slide, please.

He discusses at length the overwhelming testimony from Samsung's own internal documents that a DRAM circuit is different from the HBM which uses internal data ports and TSVs.

And go to slide 44.

He talks at length about the testimony of Samsung engineers as to why there's not DRAM circuits in the HBM products.

And if you go to slide 45, this is not an idiosyncratic view of Doctor Brogioli. The materials recognize that HBM stacks do not exist in traditional DRAMs.

There was some thought as to whether we wanted more clarification from Judge Payne as to whether it should be an

array die as used in Rajan, which is the exact language in our disclaimer. We ultimately determined that that should not happen. The reason why it should not happen is because it wouldn't solve the debate.

So the debate is to apply an array die that is different from a DRAM circuit. And what we noticed -- what I noticed actually when I came back from trial is what Samsung says and what they are planning to do is they're planning to tell the jury that array die or DRAM circuit means DRAM or any type of DRAM.

So basically we have the Judge's construction of DRAM circuit, Samsung is going to say it means DRAM or any type of DRAM, even though that's not the construction we have, and there's no additional construction that can solve that problem. The only thing that can solve that problem is the jury.

If I could have the slides -- go back to the slides.

If we go to slide 47.

One of the things I don't think is super productive in our short time is this he said-she said or he did that or they did that. But basically what Doctor Brogioli did was correct which is he analyzes the plain meaning as he has done.

What I think is defective, and we'll get into that in the motion to strike for Doctor Robins, is that what he's doing is he's literally saying a POSITA would understand Rajan to be

broad in scope.

And then if we go to slide 47, his argument for why Samsung does not infringe is because he takes a completely different reference called Matsui and he said that Matsui reference could look like the Rajan figure and that Rajan figure would therefore encompass Matsui, it would therefore encompass the HBM design.

So sometimes people who protest the loudest are people who are most guilty of a delict, and it's been my experience, but ultimately what is clear is Doctor Robins is applying the plain meaning of DRAM circuit that is consistent with the repeated renditions of Samsung engineers, the president of Samsung who launched the product, and Samsung's own expert that these are not DRAM circuits. There's a substantial question of fact.

I think the issue that has to be dealt with is whether Samsung is going to get up and be able to tell the jury that DRAM circuit equals DRAM or any type of DRAM because that's not the construction they sought from Judge Payne. If they would have sought that construction, we would have raised holy heck. But I think what we're going to see at trial is them saying DRAM circuit equals DRAM. And they have to say that because their engineers admit they don't have DRAM circuits.

Unless Your Honor has any further questions, I'll stand on the papers.

THE COURT: No, I don't have any additional questions.

MR. SHEASBY: Thank you, Your Honor.

MR. McKEON: Your Honor, very brief rebuttal?

THE COURT: Very brief, Mr. McKeon.

MR. McKEON: Your Honor, they haven't cited to any of the expert testimony from their expert that says that DRAM dies and Samsung's products are not DRAM circuits. They always point -- the expert always equates it to Rajan. And none of the alleged admissions of our witnesses said they're not DRAM circuits. We even know, of course, that -- could I get our slides, please?

Can I get slide -- the beginning of the presentation?

Sorry, Your Honor.

You know, the fact is we know that a die, everyone knows, is one or more circuits, and we have DRAM dies. That's our -- that's the core. And they -- they acknowledge this one or more circuits.

You know, this is -- I can hold a book here of DRAM circuit design, the book, it is the name of the book, you can go buy them, these are all DRAM circuits. And, you know, we think there's no fact dispute here, Your Honor, based on the evidence.

THE COURT: All right. Thank you, Counsel.

With regard to Document 197, Samsung's motion for summary

judgment of non-infringement regarding the '060 and the '160 Patents, despite the arguments, the Court's persuaded that there are outstanding questions of material fact that would preclude the granting of summary judgment, and I'm going to deny the Defendants' motion.

All right. Let's go next to Document 209, Plaintiff's motion for partial summary finding the JEDEC materials are not publicly accessible.

Let me hear from moving Plaintiff, please.

MR. SHEASBY: May I approach, Your Honor?

THE COURT: You may.

MR. SHEASBY: Mr. Huynh, deck 3, please.

I want to level set because I think there's been some confusion. The question of whether reference is publicly accessible is a question of law, not a question of fact. The facts in this case are undisputed. And we believe in light of that, the Judge should make this decision and not put it to the jury.

The evidence that is undisputed is the following.

If we go to slide 5, Mr. Huynh.

The meetings -- I want to separate two issues. There is JEDEC formal specifications. Those are prior art. Those are in. We get that. That's undisputed. What we're talking about, this is meeting minutes, draft proposals, presentations, and ballots. Those documents are not publicly

available and are not prior art. It's a question of law and the facts are undisputed.

First, there is no mechanism to search for those documents on the JEDEC website. It doesn't exist. They're not indexed and searchable. We know that from the Federal Circuit that documents that are uploaded to the database are not publicly accessible when they're not meaningfully indexed. And that is a question of law, not a question of fact, and this is undisputed. There is no evidence, zero in the record, that these documents could have been indexed, none, zero, or that they could have been searched.

The idea that documents that are not --

THE COURT: Let me ask you this, Mr. Sheasby. Is the issue there's no evidence that they could have been searched or is it the moving party's obligation to show that they could not have been searched?

MR. SHEASBY: So it is -- so --

THE COURT: Who has the burden to establish the lack of accessibility?

MR. SHEASBY: They do. They have the burden of establishing accessibility. So public availability is an issue on which they have the burden. I'm moving for summary judgment, but --

THE COURT: You're moving from summary judgment saying there's no question of fact that they were accessible.

MR. SHEASBY: That they were not accessible.

THE COURT: That's what I meant.

MR. SHEASBY: Yeah.

THE COURT: That they weren't available, they weren't -- they couldn't be reached.

MR. SHEASBY: Right. It's their burden. So I'm making -- I'm moving on a motion for summary judgment. When someone doesn't have the burden, you don't even have to put in evidence. You just have to say they didn't meet the burden.

So going back to slide 5, it's not indexed and searchable. There's no evidence that it's indexed and searchable. Not that there's evidence that could be elicited, there's no evidence that could ever be elicited that it's indexed and searchable, because we know what their experts are going to say and none of them contend it was indexed and searchable.

Let's go to slide 8.

The next issue is that even the times when -- when a document that is not indexed -- even if it was indexed and searchable, it would still not be public prior art. And the reason for that is because the documents are confidential. They can't be accessed by the public. They can only be accessed by parties who are members of JEDEC and require and comply with JEDEC obligations. Once again, an undisputed fact. No one disputes it.

In fact, if you go to slide 9, their own expert concedes that the documents they're relying on for public -- that they're saying they're prior art are documents that they would only be distributed within a given subset of groups, which is only the groups that -- only the relevant committee that would be at issue. This is Mr. Halbert's testimony at slide 9.

Go to slide 10.

He admits it again and again and again that they're only accessible to a particularized set of committee members who are subject to confidentiality obligations.

Slide 11. We've listed -- in addition to the fact that every document in JEDEC committee is per se confidential, we know for a fact that these documents were treated as confidential because if you look at appendix B to Docket 308-3, they actually have internal JEDEC confidential information designations on it.

So let me take a step back. On their burden --

THE COURT: Let me make sure you and I view the legal issues the same here. Whether or not something as a potential invalidating reference is a printed publication is a question of law, but whether that item, whether or not it's a printed publication that is publicly accessible is a question of fact. Do you agree with that?

MR. SHEASBY: So a printed publication and public accessibility under Federal Circuit law are the same standard.

So there are underlying questions --

THE COURT: So what is a public -- a printed publication and what is publicly accessible are the same thing?

MR. SHEASBY: Yes, Your Honor, and I'll give you the cite for that if Your Honor wants it.

We can go to slide 23.

This is *Carela versus Starlight Archery*, which is that the knowing use prong is not a different standard than the publication prong. *Carela versus Starlight Archery*, 804 F.2d 135. Both of them require a finding that they were accessible to the public. So if it's a publication, it's got to be printed and accessible to the public.

This is printed material that they have to establish is accessible to the public. Switching it from saying it's a publication to saying it's printed material doesn't alter the standard.

I should also note, and I think this is quite important, which is, go to slide 23.

THE COURT: There seems to be some blurring of the lines in the briefing here between these concepts, and I think it's important for the Court to be clear here.

What is or isn't a printed publication is not necessarily, in my view, what equates to a publication being publicly accessible. Those are different things in my view.

Now, maybe there's a similar approach or analysis or test, but those are not -- those are not the same thing.

Whether something's a printed publication, you can have a printed publication that's not publicly accessible, and you can have something that's publicly accessible that's not a printed publication. And I don't think -- I don't think that there is any -- there's not any authority that I know of that says what constitutes a printed publication is the same thing as to whether or not something is publicly accessible.

MR. SHEASBY: May I approach my --

THE COURT: Yes.

MR. SHEASBY: I think I understand what Your Honor is saying, and I think the most amount of guidance on it is from *GoPro*. And if you read the *GoPro*, what *GoPro* and its progeny are talking about is something is prior art if it's a printed publication. They can use it if it's a printed publication.

Whether something is a printed publication depends on whether it meets a multi-factor task in terms of accessibility to a person of ordinary skill in the art. So the -- the -- there is no other prior art standard other than printed publication.

Then what you have to do is you do an analysis, and it's a multifactorial analysis based on the undisputed facts -- on the facts in this case as to whether it is a printed

publication based -- that's how I understand it.

THE COURT: I think *GoPro* does stand for the proposition that what is or isn't publicly accessible is a question of fact and not a question of law. Do you agree with that?

MR. SHEASBY: I do agree with that.

THE COURT: Okay.

MR. SHEASBY: But if there's no disputed questions of fact, then it just becomes an ultimate question of law for the Judge.

THE COURT: If there are no disputed material questions of fact, then it's appropriate for treatment under Rule 56.

MR. SHEASBY: Right. And that's what I'm getting at.

If you go to slide 3.5. Sorry -- 6. -- are you in -- let's go to -- I think you're in the wrong one. Are you in 3 or 6, Mr. Huynh.

MR. HUYNH: 6.

MR. SHEASBY: Let's go to 3. While we're dealing with that -- yes, here we go. 3.5.

This is not a mysterial dispute of fact. They're not searchable. They haven't put in evidence that they are searchable. They haven't put in evidence that there is some index.

Go to slide 8.

This is not a dispute of fact. The materials are confidential. There is no dispute about that.

If you go to slide 9 and 10.

Their expert admits that they're confidential and are only shared within the committee.

If you go to slide 15.

The requirements about what you need to do to be a member of JEDEC, it's not just writing a check. You have to be a certain category of individual. It's a requirement, and you have to comply and commit to all these requirements, including confidentiality. That's not a dispute of material fact. There are no disputes of material fact regarding the underlying nature of these documents.

The sleight of hand and the argument that Samsung makes to get around this is to say, oh, well, these aren't printed publications, that's not the way we're using them, we're using them as, quote, knowledge in the art. But whether something is knowledge in the art or something --

Go to slide 22.

-- is the same standard. If it's knowledge in the art, it has to be accessible to the public. For it to be a printed publication, it has to be accessible to the public as well. And we know that this argument they're switching to, that you can treat all this stuff which is not accessible to the public

as prior art, if you go to slide 23, that is not in their infringement contentions.

In their infringement contentions, they refer to this as prior art systems and inventions. They have a separate --

THE COURT: That's infringement contentions, not invalidity.

MR. SHEASBY: Not invalidity contentions. I apologize, Your Honor. Thank you for the clarification.

If you look at their invalidity contentions, they never cite these documents of being knowledge of the art. They always cite -- the documents they cite as being knowledge of the art are completely different. So they never say the JEDEC materials are well known in the art or reflect the stated knowledge of a POSA. That is a new argument in reply to forestall summary judgment. So I'll come back to it.

Printed publication is a question of law. Something is not prior art unless it's printed publication. Public accessibility, which is the requirement for whether something is a printed publication, is a question of fact. If they're disputing material facts, it goes to the jury.

There are no disputed material facts. It's not searchable. It's treated as confidential. Their own expert admits it's only accessible within a committee. And if you join it, you have to commit to the committee's -- to JEDEC's confidentiality obligations, you have to commit to their

requirements for use, and you have to pay funds.

Those are the four undisputed facts, and Your Honor is in a position to make the determination. Thank you.

THE COURT: So your argument is because it's not publicly accessible, it couldn't been within the knowledge of a person of ordinary skill in the art.

MR. SHEASBY: So that's -- so, yes. And that is settled law. This is *Carela versus Starlight*. It is settled law that the known or used prong is subject to the same -- that's known or used by others means knowledge or use which is, quote, accessible to the public. Same standard.

THE COURT: All right. Let me hear from Samsung in response, please.

MR. McKEON: Good morning, again, Your Honor. Mike McKeon for Samsung. May I proceed?

THE COURT: Please do.

MR. McKEON: So we're basing our invalidity case, Your Honor, on §102(a), known and used by others. §102(a) also has an element of printed publication. We're not relying on that. We're relying on what was known to the others. (unintelligible) printed publication.

THE COURT: With regard to known or used by others, as I just heard, still is burdened with the public accessibility.

MR. McKEON: That is very true. That is very true.

THE COURT: Tell me about the public accessibility.

MR. McKEON: Okay, Your Honor. There's certainly enough fact disputes on that. In fact, we produced an overwhelming set of evidence on that.

But I was just going to say, Your Honor, on the printed publication point, additional requirements are indexing, the length of time that the publication was displayed, whether it could be copied. These are things in the cases for printed publications specifically. You don't see that kind of thing with known or used by others. But public accessibility certainly is the linchpin of both.

And these are the -- the JEDEC documents at issue, Your Honor, but let me just jump right into that. Your Honor's familiar with language known and used by others, and that is our case. The *Carela* case they cited as well. That has been our position in the expert reports and throughout the case.

This is an excerpt from his report relying on these JEDEC materials known by others. So there's no dispute that's been presented throughout the case. And Doctor Robins for the HBM patents on slide 10, same thing. We presented this as a known-by-others case.

Now, JEDEC is not some backwater association, Your Honor. This is a very, very, very familiar organization that's in the electronics business. It's the Who's Who are the members of

this and 350 members --

THE COURT: The Joint Electron Device Engineering Council. Is that correct?

MR. McKEON: That's it, Your Honor. Yep, that's it. And JEDEC -- you know, we abbreviate it for short, JEDEC. And it comes up a lot in these patents cases because it's really a robust standard organization, 350 members, and, you know, people interested in -- and other people who are interested. And, you know, here's the list of examples of the members of JEDEC that participate in the meetings, participate in the submissions. All this is part of JEDEC.

And to prove the point, Your Honor, we even look at their own patents in this case. They cite JEDEC. The '339 cites the JEDEC and discusses JEDEC. The other patents--'918, '054--cite JEDEC, discuss JEDEC. So JEDEC is a well known organization and is directly relevant to the patents-at-issue in this case. It's even cited in the specifications.

Even Mr. Murphy, who's the expert that Netlist has in the re-exams in the IPRs in the Patent Office, understands very well. And I'll quote from -- from him in his deposition: One of ordinary skill in the art would understand about the content that was available in the JEDEC proposals, independent of how they got that information.

So whether they were reading the actual submission or whether they were sitting in the meeting, Mr. Murphy is very

clear that these are the POSITAs and these are the people that are going to see and have that information available to them.

And you know, of course, accessibility, Your Honor, the GoPro case my colleague already cited to Your Honor, what is the standard? It's reasonable diligence to locate the information by interested members of the relevant public.

And the point here is we can't get confused by the word public. It's not everybody out there in the public literally. It's the relevant public. And I submit, Your Honor, and we submitted in our briefing, that JEDEC is actually the relevant public here. These other people that are in these organizations from all the major companies that are involved in the memory business, in the semiconductor business. And the Who's Who is the POSITA in this case. They are sitting in these rooms. So they have access to all this material.

THE COURT: So only a member of JEDEC could be a person of ordinary skill in the art in this context.

MR. McKEON: Well, I would say --

THE COURT: Do you agree?

MR. McKEON: I would say, Your Honor, that they define sort of the POSITA here. This is -- this includes a group of people who -- who participate in the creation of all this documentation, the creation of what ends up to be the standards.

THE COURT: But if the documentation is not

searchable, it's not indexed, it's not available except to members of this organization, then what I hear you saying is it may be known to a person of ordinary skill in the art, but what -- but who is a person of ordinary skill in the art must in this context include a member of JEDEC who has access through that membership to this material. Is that right?

MR. McKEON: That's right, Your Honor. I mean, certainly they would be included in that group. There's no doubt about it. And by the way, Your Honor, we do dispute the notion that you can't search or get ahold --

THE COURT: Show me that.

MR. McKEON: Yeah, I will show you that. Let's get right to that.

This is the -- this is the website here, and you can see here on slide 22, Online access to JEDEC confidential meeting minutes, presentations, ballots, and surveys. All members, not just the members in the subcommittees, but all members have access to the material, including the actual standards, the proposals, and the ballots.

And, of course, we have our expert submission on this. The comments and changes are uploaded to the JEDEC member website where they can be accessed by members. It's very clear that accessibility to these documents from any moment within JEDEC and even members from the public outside JEDEC, which I'll get to in a minute, there is no requirement under

the case law that particular members of the public -- the relevant public actually received the information. It's just about reasonable accessibility. That is the standard.

This is another portion of the member website, Your Honor. You can register -- non-members can register to receive these materials just by putting their email in and a password. And we've highlighted that portion here on the lower right for -- anybody can do this. And, of course, non-members who want to become members can pay a fee and join JEDEC. And they encourage that. The more, the merrier. That's for sure.

And non-members can also attend meetings. We see here from our -- from our expert Mr. Halbert about interested members bring your checkbook and, of course, they can sign up to join, parties that are interested.

And -- and we know, Your Honor, from the evidence that we presented in summary judgment documents that, in fact, non-members do go to the meeting. We have an example here where these two companies, Acuid and Hiroce Electric, attended this particular meeting, and they were non-members at the time they attended. And the idea here that they can come in, hear it, you know, try it out, see if they are -- if some of them are interested in joining. They have access to this.

And, of course, this is an example of a document that's created at every single meeting. They've created who's there,

what presentations were given, who gave them, and all the presentations are given a number and are labeled as such.

Let me skip this slide, Your Honor, and get right to the evidence here.

And, Your Honor, the website of course is finding -- again, we dispute the fact that you can't search for this material. It's very accessible. The JEDEC website, you just click on the committee and you have a search box.

This is an example of the website in slide 32, Your Honor. You see the search box there? You can search for these committee documents right in the search bar and access any of them. And there's also a portion of the website you see here where we had the standard, different topics within the standard that you can click on and get focused on a particular subject matter that you're interested in.

THE COURT: Let me ask you this. Going back to what you just said, if a member -- if someone who's not a dues-paying member of JEDEC goes to this website, takes that search bar, can they input a search that will yield the exact material that Samsung seeks to use as an invalidating reference here?

MR. McKEON: Well, what I can tell you --

THE COURT: Sometimes the organizations where you can get a certain level of information just by going to a search bar, but you can't get the same granularity of

information than if you are a member or you have other special credentials.

So you're telling me that the documents or the information or the data that you intend to use as an invalidating reference can be accessed by anyone who goes to that page online and puts in the right search query in that search bar. Is that what you're telling me?

MR. McKEON: Well, Your Honor, I want to be very specific.

THE COURT: I'm trying to be specific, too.

MR. McKEON: Yeah, yeah. So the only -- the best I can do on that in terms of the record is we have here, you can access by doing a password, your email, most site content is available for non-members at no charge. And then, of course, we have the notion that anyone can sign up.

THE COURT: Is that most site content the same content you want to use as an invalidating reference?

MR. McKEON: Well, Your Honor, we haven't sort of pulled that specifically together. I think the record we have here is that -- that certainly the JEDEC -- anyone within JEDEC has access to that material, anyone within JEDEC. And our -- our view is, Your Honor, that is the relevant public.

THE COURT: You are the one presenting the argument that someone who's not a member of JEDEC can simply go online and put in a search and get the same material you want to use

to invalidate the Plaintiff's patents. If you've got evidence of that, I want to know that.

If it's not that easy and if the precise invalidating references you want to present to the jury aren't accessible by non-JEDEC member, then I need to know that. That doesn't necessarily foreclose whether or not a member of JEDEC is a, quote, relevant member of the public with accessibility or not.

MR. McKEON: Yeah.

THE COURT: But you either have a narrow argument or broader argument and you're going back and forth between them, and I need to know does the broad one really hold or is it the more narrow one.

MR. McKEON: Yeah. No, it's a really good question, Your Honor, and what I -- the evidence I have on that link is what I show you here, and whether -- I don't have evidence that specifically links it to the documents at issue here.

THE COURT: Okay.

MR. McKEON: Other than -- other than we have here is, you know, most -- most content, is what they say here.

So, Your Honor, I don't think we need to resolve that issue because I think the JEDEC membership itself is the relevant public. 350 companies and -- and all the people involved in the JEDEC organization is the relevant public for the purposes of the prior art accessibility analysis. And

that's what the case law tells you.

Let me jump to a case, Your Honor, that I think is on point.

THE COURT: Let me ask you this. Is your intended witness to present this invalidating prior art or potentially invalidating prior art, is this witness a member of JEDEC who has access to this on their own, or are they relying on some third party who won't be in court who's told them, I'm a member, I've seen it, this is what it says and means? Exactly how is this going to be teed up for the jury?

MR. McKEON: This is our expert, Mr. Halbert, here. He worked for Intel for many, many years. Intel is a very robust, very --

THE COURT: I know who Intel is.

MR. McKEON: Yeah, you know. You know, of course. But the JEDEC -- you know they're involved in JEDEC is my point. And he put in expert reports into this case regarding the JEDEC process, who's involved. And, you know, we've cited to a bunch of that, Your Honor, in our summary judgment papers. But he will be doing the bulk of that at our trial in terms of the involvement in JEDEC, in addition to potentially a fact witness from Samsung who's involved in the JEDEC organization as well, as Your Honor knows.

But just -- just one more point about the public access, Your Honor. I mean, I think I said it, but I want to make

sure it's clear as you were asking the question. Anyone can pay to get access. Anybody. So you pay --

THE COURT: Access to everything, not access to most.

MR. McKEON: Access to everything, yeah, yeah. That's just to bring your checkbook is the expression that Mr. Halbert used.

And, finally, Your Honor, I don't think there's a dispute here so I'm not going to go through the slides. The date -- these documents are all before the critical dates. We've shown that through the evidence. But since that wasn't raised by Mr. Sheasby, I won't.

And likely confidential designations, Your Honor, wasn't raised by Mr. Sheasby, so I won't -- I won't talk about that. But unless you have any more questions, Your Honor, that was my presentation.

THE COURT: No, I don't think so, Mr. McKeon. Thank you.

MR. McKEON: Thank you.

MR. SHEASBY: Brief response, Your Honor?

THE COURT: Yes. Tell me, Mr. Sheasby, can anybody with a checkbook make a payment, join, and get access to everything online?

MR. SHEASBY: No. And don't take my word for it. Take the word of their expert, Mr. Halbert. The documents

presented in the JEDEC committee and task groups, including ballots, technology presentations, or meeting minutes, were normally not shared publicly outside of JEDEC. Rather, these documents were accessible -- only were accessible to members of the JEDEC committee, the JEDEC member website, or voting machines. It's not my words, not the words of my experts; that's the words of their expert. The expert goes on.

The other expert is Mr. McAlexander, slide 10.

THE COURT: Let's go back to the prior statement. I want to read that one more time.

MR. SHEASBY: Yes, Your Honor.

THE COURT: What's your understanding of the sentence that begins after the yellow highlighting, "Only once a standard has been finalized and approved for publication would non-JEDEC companies be able to view the standard's materials?" Doesn't this really tell me that at time of adoption, it may not be public, but upon adoption, it becomes public to the broader membership and perhaps others?

MR. SHEASBY: Yeah. That's what I was going to ask him. I thought that was an open issue as well, which is why Mr. McAlexander answers it, which is, as to the ballots, go to slide 10.

This is what Mr. McAlexander said, is that they're accessible only to the committee members and it's not for a temporal period of time.

And we know that they're accessible only to the committee members in perpetuity because, if we go to slide 14, JEDEC produced these materials. The way we were given access to these materials in an evidenciarily sort of acceptable forum was through JEDEC, and JEDEC treats them all perpetually as Attorneys' Eyes Only, Confidential.

THE COURT: Let me ask you this question. In your view, is the determination about what constitutes the relevant public a fact question or a question of law?

MR. SHEASBY: It's a question of law. And the reason for that is because they've answered the question themselves.

This is their -- I'm going to do it in pieces. It's going to be too big. This is their definition of a person of ordinary skill in the art.

It's not the right page.

This is their definition of a person of ordinary skill in the art. It's -- and it was -- it's the same for all of them. It's a Bachelor's degree in electrical engineering or the equivalent education, and from two or three more years of experience in integrated circuits, et cetera, et cetera.

So their definition of the relevant public is not a JEDEC membership. Their definition of the relevant public is a much, much lower standard. JEDEC is purely for companies and for the people based on its membership requirements.

If you go to slide --

THE COURT: But who might be a part or a member of the relevant public is not necessarily the same thing as who might be a person of ordinary skill in the art. Is that not true?

MR. SHEASBY: Well, I think if you look at the -- the case on the subject, and I think their best case on the subject is *SRI International.*

Go to slide 7.

So this was a summary judgment case. So the Defendant in that case was in a similar situation in which there were people who were, quote, in the relevant public who he had given it to -- he had given to them, but that still didn't make it publicly accessible. And so I don't think the fact that -- the fact that someone in the, quote, relevant public has access to it makes it publicly accessible.

In fact, if you look at their best cases on the subject, slide 19 and 20, this is the relevant public. The documents were shared with joint venture members that were maintained in confidence. Those would not be prior art. But the fact that most, if not all, of the information was available without restriction is what made it prior art.

So this -- they argued that that was the relevant public because the joint invention members were some of these same persons of ordinary skill in the art. And they said that was

not enough to make something a prior art reference.

So I think under the Federal Circuit case law, accessibility to the relevant public, the fact that individual members of the relevant public had access to the document is not the standard. The standard is, based on their definition of prior art, was it publicly accessible.

THE COURT: All right. Anything further?

MR. SHEASBY: Nothing further, Your Honor.

THE COURT: Okay.

MR. McKEON: Your Honor, brief rebuttal response?

THE COURT: I don't think so, Mr. McKeon.

MR. McKEON: Okay.

THE COURT: With regard to Document 209, Plaintiff's motion for partial summary judgment finding that the JEDEC materials are not publicly accessible, I'm going to deny that motion.

I think there is a fact question as to who is the relevant public, who has access. There's too much at issue here about the degree of difficulty or accessibility of reaching this material. I just don't believe it's appropriate for treatment under Rule 56. I'm going to deny the motion.

Counsel, before we break for lunch, because we're about five minutes from the noon hour, I asked you, when I recessed earlier, to confer with each other about what on the list of disputed motions might be impacted by the Court's ruling on

the license issue under the JDLA.

Have you been able to do that and do you have a consensus or do you need more time to talk about that?

MR. SHEASBY: Your Honor, nothing on the motions to strike or motions for summary judgment are impacted. It's possible that there would be some motion in limine practice that would be narrowed based on it, but we haven't identified that specifically yet.

THE COURT: Mr. Cordell?

MR. CORDELL: Well, we identified it, Your Honor. We just haven't come to an agreement, but we'll continue to work. We have a pretty good idea what that should be.

THE COURT: All right. Well, we're going to recess for lunch at this point. I'll take up Document 211 when we come back. That's Plaintiff's motion for summary judgment on Samsung's invalidity defense based on *Sotera* under §102. I'll take that up after we return from lunch.

Counsel, let's take about 45 minutes. I don't think we can take any more. That should put us back about 20 minutes -- about 40 minutes after 12:00. We'll reconvene about 40 minutes after 12:00.

Until then, the Court stands in recess.

(Lunch recess.)

THE COURT: Be seated, please.

All right. Let's next take up Plaintiff's motion for

partial summary judgment on Samsung's invalidity defenses. This is Document 211.

Let me hear from the Plaintiff first.

MR. SHEASBY: May it please the Court.

This motion is the mirror image of the motion Your Honor just heard previously. In the motion Your Honor just heard previously, there was an allegation that the JEDEC ballots and materials were publicly available. If they were publicly available, those materials could have been used in the IPR proceedings that were initiated by Samsung and --

THE COURT: That would be subject to estoppel.

MR. SHEASBY: It would be subject to it. And it's not just estoppel, but it's a *Sotera* stipulation in which they committed unequivocally to not present any invalidity defense that was or could have been raised.

If you go to slide 4.3.

And so I was listening very carefully to my brother in arguing to avoid summary judgment, and my brother was very clear that he maintains that these references and this information was publicly available and it's printed. If it's publicly available and it's printed, the statements that were made to avoid the last summary judgment, it is unequivocally encompassed within the *Sotera* stipulation.

If we go to slide 8.

What they've done, and they've done this for each of

these patents, and I'll just give some examples, is they have these core references that they've used in the IPR and they've appended them to things--JEDEC proposals, QBM, and APA. All of those things are the printed publications that they are now telling you were publicly available and therefore could have been prior art, and if they could have been prior art, they could have been brought.

So it is impossible for them to be taking two positions: one, the JEDEC materials were publicly available; and, two, that the JEDEC materials could not have been presented in the PTAB. And if the JEDEC materials could have been presented in the PTAB, the *Sotera* stipulation applies.

Let's go to slide 23.

So there's no doubt that JEDEC specifications can be used in the IPR proceedings, and we know that because Samsung actually used JEDEC materials in its IPR proceedings in this case.

The JEDEC materials that they rely on, they repeatedly, repeatedly, describe them as being freely disclosed, freely disclosed, freely disclosed. But if they were freely disclosed, that means they could have been presented to the PTO and under the IPR proceedings as, in fact, in certain circumstances Samsung did.

Let me give you an example. This is another one of Samsung's challenges in which it uses the exact materials, the

exact class of materials that it's claiming are publicly available and are publicly accessible in its IPR challenge. There's no conceivable way it can have it both ways. If they are publicly available, they were citable in the PTO. And we know they were citable in the PTO because Samsung cites them in the PTO.

There is one other move that Samsung makes to try to avoid this.

THE COURT: Let me stop you, Mr. Sheasby, before you go further.

The motion, as I read it, addresses the JEDEC materials, it addresses something called quad band memory, it addresses the APA, it addresses the '339 patent, the '506 patent, the '918, the '054. There are multiple patents at issue that seem to come and go.

If you can be precise with exactly how these pieces all connect to fall within the estoppel argument you're making, then that would be helpful.

MR. SHEASBY: Yes, Your Honor.

So the --

THE COURT: And if you're conceding that some of the pieces fall outside of this argument, tell me that as well.

MR. SHEASBY: Sure. So I'm -- if we go back to slide 8.

So for each of the three combinations, this

patent -- this -- the relates to the '339 Patent, the '506 Patent, and the '984 and the '105 Patents. Those are the patents that are at issue in this motion.

There are three buckets of -- Samsung used a common tactic in the mirror image IPR references it filed and then the validity attacks it is making in this case. So what it did is it originally filed the IPRs on, for example, Elsberry and Halbert. It filed a *Sotera* stipulation saying it would not present anything that it could have raised in the proceedings.

And the way to try to evade that stipulation was for it to append one of three classes of information--the JEDEC proposals, a set of QBM materials, or what it describes as admitted prior art. Those are the three things that it adds to each of its obviousness combinations to try to avoid the *Sotera* stipulations.

So you'll see that was for the '339. You see it does the same thing for the '506. It adds the JEDEC proposals, the QBA, and the APA to the previous grounds that it presented to the PTO. And it does the same thing for the '984 and '054 Patents.

So the relevant question on summary judgment is, could JEDEC materials, could what they describe as admitted prior art, and could what they describe as QBM, be presented in an IPR proceeding. And if it could, the stipulation they filed

on slide 3 triggers on grounds that were raised or reasonably could have been raised.

Your Honor has heard me make this argument previously. I have been an advocate of a much more vigorous and aggressive position on these *Sotera* stipulations. This doesn't strike me as a case in which it's a close case that the -- let's start with them, to begin with, the JEDEC materials. If we go to slide 12, we know -- sorry. Slide 23.

Slide 23, we know that the JEDEC materials can be used as prior art because Samsung brought them in other PTO proceedings, IPR proceedings, with these materials. This is an example of literally what Samsung does in the PTO. It's appended JEDEC materials, the exact identical class that it's using in this case, and appended them to its IPR petitions.

It cannot be possible for it to have it both ways. If the arguments are publicly -- if the documents were publicly available, which is what they're contending, they could have been presented to the PTO, and on that, the summary judgment should have been granted.

Now, there is a move that Samsung makes to try to have it both ways. And the move that Samsung makes to try to have it both ways is, even though the -- the invalidity materials -- invalidity reports rely on these materials just as a printed publication, no different than a patent--these are excerpts from Mr. McAlexander's expert reports--they

have -- they -- they try for a slight nuance. They say, Well, these materials couldn't have been presented to the PTO because we're using them to knowledge -- to acknowledge the state of knowledge in the field, not as printed publications.

But that argument is foreclosed because in their invalidity contentions, they never identify any of the JEDEC materials as knowledge in the art. They treat the JEDEC materials as what they are--disclosures. Those disclosures are printed. Those disclosures could have been publicly -- could have been used before the PTO.

And for Samsung to have prevailed on summary judgment by saying, no, no, these were publicly available, you could access through the JEDEC site, it's all good, if that's the case, they could have been presented.

I should also note that the Federal Circuit has made clear that this argument by transforming something into the state of the knowledge allows you to evade the PTO's *Sotera* stipulation. The PTO -- in a PTAB proceeding, you're allowed to rely on the knowledge of a person skilled in the art.

And so for these JEDEC materials, whether they are treated as they were truly disclosed in the invalidity contentions, which is as printed publications, or whether they are treated as, quote, knowledge in the art, under either circumstance they could have been presented in the PTO because the Federal Circuit said, the PTO has considered it and they

have their own behavior.

Now, to be clear, we do not believe the materials are publicly available. We believe summary judgment should have been granted on that issue. But Samsung cannot have it both ways. It cannot evade summary judgment by saying these materials are publicly available, and then flip and say, Oh, we don't have to do our *Sotera* stipulation because we didn't -- we couldn't have presented in the PTO. That is a logical impossibility.

So that's the JEDEC material information.

The next set of information is the quad band memory documents.

And if you go to slide 32, Mr. Huynh.

The Court has made clear that simply calling something a system when it's a printed publication does not allow you to evade the *Sotera* stipulation. And, in fact, Samsung agrees that a collection of printed publications don't transform something into system art.

In this case, the quad band memory details, literally what they rely on for quad band memory is four presentations, printed publications. Not testimony about how the system would have operated, not something additional confirmation -- confidential. Four printed publications. Those printed publications were all publicly available.

By the way, that's how they get them in as prior art in

this case--by saying they're publicly available.

THE COURT: Go back one slide, please.

Your heading on this slide says, "Samsung admits QBM system is described in its entirety in four documents." If this is system art, system art is not appropriate under the IPR structure, is it?

MR. SHEASBY: Yeah. So that's -- the reason why I put system in quotes was for that exact reason, Your Honor, which is that when you take -- you can't use the word system to launder art to avoid a *Sotera* stipulation, which is, if the, quote, system is four printed publications, those printed publications could have been presented to the PTO.

The -- the consequence of that would be that the *Sotera* stipulation would become meaningless because you could always treat a printed publication as a system and therefore the stipulation would have no meaning.

And in this case --

If you go to slide 45. I take that back. Let's go back to slide 40. Give me one second. Slide 16. Excuse me. 34. 34.

So Samsung agrees with this as well, which is that if system prior art refers to patent or printed publications, it's governed by the *Sotera* stipulation.

So there are two buckets here so far. There's JEDEC materials and there's QBM. JEDEC materials, there's no

allegation that they're a system. They're just printed materials. Arguing that they are -- count as knowledge in the art is irrelevant because that can be considered by the PTO.

QBM -- QBM, quote, system is not a system at all. It's four printed publications, two of which relate to completely different designs. One relates to a QBM generation 1, and then two relate to a QBM generation 2. So it's not even the same material that they're relying on.

The third issue is admitted prior art. And for admitted prior art, what they do is they take -- and we can go to slide -- slide 13.

What they do is they take what they quote admitted prior art such as patents that are incorporated by reference into our publications -- into our patents and then use them as prior art against us.

Slide 18 is an example. 18.

So here is what they rely on in the IPR. They rely on the incorporation by reference to the '386 Patent, and in this case they're relying on the same '386 Patent and saying, Oh, we couldn't have brought the '386 Patent in front of the PTO because it's admitted prior art and admitted prior art is not accessible -- is not usable in front of the PTO.

But if you go to slide 13.

On this slide, you'll see that the PTO has made clear that admitted prior art can be used as long as it's in

combination with some other reference.

And if you go to slide 15.

The Federal Circuit -- and this is the controlling decision from 2020 on this issue--it's the *Philips* case--is that the board can appropriately rely on general knowledge, admitted prior art in particular, to supply missing elements.

So one of the things that Samsung takes to evade this is saying, Oh, well, we use the APA in too robust of a form for it to be presented by the PTO. But that's not what the Federal Circuit states.

So we've now talked about in the *Qualcomm* case, which they cite, says the same thing--that it's appropriate to rely on admissions in a patent specification when assessing whether the patent's claims would have been obvious.

So there's three buckets here. JEDEC materials: If what they said is true in the last summary judgment that they evaded, they could have been presented to the PTO.

Two, QBM: They call it a system, but it's four printed publications. That's literally all it is, four printed publications.

And APA. APA, they're allowed to rely on because their APA is just printed publications.

So I'll now turn -- unless Your Honor has any questions, I'll now turn to the last ground for -- for summary judgment.

THE COURT: Let me ask you this before you do that.

On the admitted prior art, are you urging this motion with regard to admitted prior art as to all four of these patents or less than all four of them?

MR. SHEASBY: I believe they use APA in all four of the patents, Your Honor. No -- yes, in all four of the patents, they use APA.

THE COURT: Okay. Then go forward with your next point, please.

MR. SHEASBY: So the last point is a sort of nuanced point, and the issue is -- is the following, is they treat these long lists of ballots and publications and presentations as a single reference.

Let's go to slide 41.

This is an example of what they do. So Samsung admits that the JEDEC proposals include the following, and they are from five different committees. And they say the combination of those five different references from five different committees constitutes a single anticipatory reference. And there's no attempt to provide a motivation to combine between these five individual references; they just treat them as one omnibus reference.

So it's running an anticipation argument because we all know how juries are compelled by anticipation as opposed to obviousness, but it's surreptitious because it's combining five references.

So this is not to preclude them from presenting these references as prior art. It's to preclude them from presenting them as anticipatory references because they didn't show a motivation to combine each of the five because -- because they're five separate references. And they can't just convert obviousness into anticipation by combining stuff together, especially from different subcommittees.

Let's go to slide 45.

Their excuse for that is that, Oh, well, we just rely on each of the other ones for some purpose to sort of -- to sort of elucidate what's in one of them. But the instances in which extrinsic evidence can be used for anticipation is quite narrow.

It can be used to explain, to show why something is inherent in a reference, but not to expand. They're not using these combinations of five different prior art references to show inherency. They are using it to expand. They take these references -- they do the same thing with JEDEC materials in Kentron.

It's not that they're saying, Oh, well, that makes something inherent in that and that makes the same thing inherent and that makes the same thing inherent. Each of them they cite for different purposes. So that's the fourth basis for -- for a summary judgment of no anticipation.

So, to summarize, JEDEC materials, they could have been

presented. They should not be allowed to present anticipation -- invalidity defense on that. QBM materials are system -- are not system prior art. They are just printed publications. APA could have been presented, and they cannot run an anticipation case that just shoves together references from five separate committees and call them one reference.

Thank you, Your Honor.

THE COURT: All right. Let me hear from Samsung in response, please.

MR. McKEON: Good afternoon, Your Honor. Mike McKeon for Samsung.

THE COURT: Please proceed, Mr. McKeon.

MR. McKEON: So just to level set on what we're asserting in this case, again we're relying on the knowledge under §102(a), the knowledge of a POSITA, and we're relying on evidence. The evidence of that knowledge is certain JEDEC disclosures that occur within the JEDEC meetings.

And then there are instances in some of the -- some of the claims where we combine these for some element, there's some elements that are used in the combination. But the basis, the primary -- for each of these, the primary source of prior art is the JEDEC, it's the knowledge of a POSITA, for the '339, '506, and the '908, and '054. And that's -- that's very important, Your Honor.

So one important point, too, Your Honor, let's make sure

we understand what's going on in the Patent Office --

THE COURT: Let me ask you this, Mr. McKeon. How do we get away from the all you got to do is have a checkbook and you can get everything in JEDEC you want to see.

MR. McKEON: Yeah. So --

THE COURT: I mean, that's pretty far off from one end of the spectrum for you to now try to walk back and say, oh, no, only a person of ordinary skill would have that and that's not anybody with a checkbook and -- I mean, you're going to have to reconcile the prior argument with this situation.

MR. McKEON: Your Honor, I clearly see the tension there, but what I will --

THE COURT: So do I.

MR. McKEON: So do you and you're the most important person in the courtroom.

The Patent Office proceeding is just one single ballot, and it's identified here on the slide. And, of course, in the Patent Office, as Mr. Sheasby alluded to, they're arguing it's not a publication. And the Patent Office instituted the proceeding, and we're not going to get a ruling for that for some time, but they haven't decided the issue of whether it is a publication or not a publication. They held that into -- into their final decision. But in -- in that proceeding, Netlist is arguing that it's not a printed publication.

Now, of course, in this proceeding, we are talking about different -- we're are talking about different documents, and these documents are documents that evidence the knowledge of one of ordinary skill in the art. So we're using them in a very different way than we're using them in the Patent Office. It's the knowledge of one of ordinary skill in the art in the Patent Office under §102(a).

We're not -- we're not taking the position in this case and we're not going to argue to the jury that they're printed publications where we have all the indexing and all that. We're not -- we're not making that point to the jury; we're not going to argue that to the jury.

They're just evidence of what one of ordinary skill in the art would know during the relevant time frames. And these documents that we cite here, which are not at issue in the Patent Office, are evidence of that. The issue of whether they're printed publications or not, is not an issue the jury is going to have to hear about. And that's -- that's the difference between what's going on there and what's going on here.

With respect to the JEDEC issues, and I do, you know, acknowledge, Your Honor, they're the most complicated in terms of the issues before the Court now.

THE COURT: But this is not about -- as I appreciate it, this is not about how did you use it here versus how do

you intend to use it there.

It's did you know it, was it available to you, could you have used it at the PTAB in a way that you're now going to use it, or even if you're going to use it a different way, did you have it? Did you know that it could be used at that body and could it have been presented there? That's the test in my view for what should have been -- or what's potentially subject to estoppel and what's not.

But saying we knew about X, we used X in this way there, and because we're going to use X in another way here, we're not estopped, doesn't seem to me to be a strong argument.

MR. McKEON: I understand the point, Your Honor, and clearly see the tension here. But what I would say is what do we know from the case law here in the *Qualcomm* case is that in the Patent Office -- in the Patent Office you cannot use, you cannot use, as a basis knowledge of one of ordinary skill in the art. You can't make that argument in the Patent Office. We all understand that, it's printed, and, you know, the patents are printed publications.

But you can -- you can have sort of gap fillers with knowledge of skill in the art. And here in the district court, in light of the statute and the stipulation, it's the complete opposite. We can -- we can use knowledge as our basis here, and then we can use printed publications and patents as gap fillers. It's the complete opposite of what

happens in the Patent Office.

And that's all we're doing is we're -- we're looking at the precedent here, including the *Qualcomm* case, where the line was clearly drawn in the context of the APA, admitted prior art. But the -- the -- the lesson from the case applies here, which is knowledge.

You can use knowledge in the Patent Office as a gap filler; you can't use it as a basis. In the district court, we can use it as a basis, but we can't use it, you know, in the Patent Office as a basis, and we can use patents and printed publications here as gap fillers. That's the line we're walking, and we think the case law supports that.

Again, we're not going to say that these are printed publications, and they haven't asserted in this case at all, Your Honor, that they're printed publications, the documents at issue here. And that's really the story with the JEDEC. And it's clear that if you go through the expert report of Mr. McAlexander, the basis is the JEDEC material which is the evidence of what one of ordinary skill in the art would know at the relevant time. And that's right from the statute.

And this is all we're trying to do is present to the jury knowledge of one of ordinary skill in the art, and with respect to each of the patents we lay out here on the slides here, Your Honor, the basis is -- for the JEDEC issue is the JEDEC documents that we cite which informs the knowledge of

one of ordinary skill in the art. And that -- that is all we're arguing here.

And we cite the same cases that Netlist cited. We cite the *Philips* case here which talks about, you know, the knowledge, you know. The knowledge is evidence of -- of what one of ordinary skill in the art would understand, and you can do that. And we have the JEDEC, of course, here in this case informs one of ordinary skill in the art and what our expert says.

And that's the JEDEC argument, Your Honor, and, you know, we think we're using it in a different way here and we're using it in a way that we could not do that in the Patent Office.

I think the QBM issue, Your Honor, is -- is not close to the line. This is a situation where we have these four documents, and we're basing it on, of course, the evidence, the system. So this is about a system case.

And during discovery, Your Honor, we deposed the co-founder of the company called Kentron, and they testified about this QBM system. And he says here on his -- what we have on the slide here, that the system was being demonstrated. And what he talks about is the demonstration, and he talks about what QBM was doing at the time to JEDEC members and presenting this to JEDEC members, was the testimony we got from Mr. Kentron.

And we are certainly relying on, Your Honor, these four documents that Mr. Sheasby cites to. But, again, how we were using them under §102(a), this is a system prior art. System can't be presented to the Patent Office. There's no dispute about that. And that's all we're going to do is explain to the jury that one of ordinary skill in the art and particularly JEDEC people -- JEDEC would understand the QBM system through public knowledge.

And, Your Honor, your case here that you -- which you wrote back in 2016 really, you know, makes the point that, you know, defendants have considerable latitude in using prior art systems embodying the same patents or printed publications placed before the PTO in IPR proceedings.

And that's all we're doing here is -- now, and I'll note here for the record, Your Honor, the QBM documents are not at issue in the Patent Office, in IPR. So no one here in the courtroom is telling you that these are to satisfy the printed publication standard.

What we're saying is they satisfy the standard under §102(a) in terms of supporting the evidence of what one of ordinary skill in the art would -- would know during the relevant time frame. And the fact that these underlying documents may be printed publications that could have been used at the PTO, and, again, there's no -- there's no evidence that that's the case, but if -- even if that was true, I think

Your Honor's prior analysis in the IV cases is spot on with that.

And, finally, Your Honor, APA, you know, similar point here. You know, it's focusing on the basis, and we're only considering APA as the -- the basis, the main -- the main deal, if you will, in terms of the prior art case. And, yes, there will be some other references that are gap fillers, but that's exactly permissible. And we look at the *Qualcomm* case. You know, at the Patent Office, APA can be used as a gap filler. It can't be the basis. Here in this district court, it can be the basis, and that's all we're using it for, Your Honor.

We can go through each patent, but we have here APA for the '506. And, as Mr. Sheasby said, for two of these patents, we're using incorporation by reference, which means you're entitled to do that because it's as though that same disclosure is in the Patent Office. So it's an incorporation by reference, and we're entitled to use that as APA because that disclosure is literally as though it was printed out and disclosed in the patent application in the published patent.

And we do the same, Your Honor, for '339. Again, there's incorporation by reference issue. The fact that we incorporate by reference has no impact on the fact that you can't use APA in the Patent Office and we can use APA here in district court.

And the same is true with '918 and '054, Your Honor. It forms the basis -- if you look at Mr. McAlexander's report, we have some cites here. It is the main deal for these elements for the main set of claims. And, you know, I think for the independent claims, in most cases it's the vast majority of elements. And there might be occasions where one or two are used in combination, but the APA is the main basis of the rejection.

One more point about a reference in the briefing, and I just alluded to it, but if you look here on the -- on the slide, this is an excerpt from Doctor Mangione-Smith's report, and each of these highlighted in yellow is an independent claim. And what he's demonstrating here is that, with respect to the independent claims, it's all APA except for the one element -- for that one element in the claim. There's another reference that's used. So the basis clearly is -- clearly is APA.

Okay. And with that, Your Honor, I'll move on to the last -- the last issue is the anticipation point here.

So the question is whether when we have a knowledge case that we're presenting to the jury, whether we can rely on a multitude of -- of different references to demonstrate the knowledge. And we submit, Your Honor, that you can do that. We cite the -- we cite the *monsanto* case for the notion that extrinsic evidence may be used to interpret the allegedly

anticipating reference to shed light on what it would mean to a POSITA.

And that's what we're doing with respect to JEDEC. We're looking at the documentation, and we're -- we're going to present to the jury that this is evidence of one of ordinary skill in the art would think. And it's in an anticipation case for some claims, not all of them, and we're only relying on the knowledge of a POSITA as evidenced by the different documentation. And --

THE COURT: Let me ask you this. Hasn't Samsung conceded that there's not a §102 defense as to the '918 and '054?

MR. McKEON: With respect to the '918 and '054, Your Honor, I believe -- let me just -- give me a moment here -- I believe it is for '506 and '339 only.

THE COURT: Right, which means it's not as to the '918 or the '054?

MR. McKEON: That's correct, yes.

THE COURT: Okay.

MR. McKEON: Okay. Thank you for that clarification. And so that's all we're talking about with respect to JEDEC here with those patents.

And, again, we cite the -- the *Iridescent* case where, you know, the Court says that -- you know, the Court finds the defendants will be allowed to use the standards document as a

single reference. And this is a situation where it had an LTE standard and then there was multiple documents, and the Court acknowledged that, you know, you can use that as -- as evidence of -- of, you know, the general knowledge of the anticipating network.

And, again, more from the *In re Baxter* case, this is a case from the Federal Circuit, Your Honor, where there was a prior art reference, and one of the key limitations in the claim was whether this blood bag was plasticized with DEHP.

That wasn't in the reference, but the Federal Circuit says you can use, you know, the general knowledge of one of ordinary skill in the art. And it was well understood that Becker had commercial embodiments of these bags that did use plasticized DEHP. And they -- they said that that could be used as part of the -- even though it's anticipation, you can use that knowledge to fill the gap. And that's the *Baxter* case.

The *Kyocera* case that is cited, you know, was a -- a public -- you know, a printed publication case, not a knowledge case, a §102(a) case. It's a little different in terms of its application. And, you know, we have demonstrated, we believe, that the 2011-2012 JEDEC proposals, you know, the ones we're relying on at least, you know, those are evidence of the knowledge, and that's all we're relying on for here.

And, Your Honor, the QBM system, I think the law's pretty clear, you know, in this district at least, that in our cite here from a case, *IP Innovation*, this Court sees no error in using multiple references to describe a single prior art system for the purposes of showing anticipation. And that's a case that -- from John Love that came out in 2010, and it's consistent, I think, with Your Honor's cases as well.

And that's all we're doing here is -- is that we're using, you know, evidence -- these documents as evidencing what the actual QB system is. And the system is the only thing that's going to be argued to the jury. This is the prior art system based on the testimony of the founder of the company and the supporting documentation.

And with that, Your Honor, I'm done.

THE COURT: All right. Anything else, Mr. Sheasby?

MR. SHEASBY: Contracts need to be enforced according to their terms. There was a *Sotera* stipulation that was filed in this case.

If we go to slide 25. Slide 25, please.

Every single reference that they're relying on for QBM, for APA, and for JEDEC, was a reference that they claim is a printed publication that could have been before the PTO. Every single one. We just had an hour argument before lunch in which they were swearing up and down that these JEDEC materials were publicly accessible to anyone who wants them.

They're printed, and they claim they're publicly accessible.

They evaded summary judgment on this point, and they need to accept the consequences of that.

This argument -- if we go to slide 28.

This argument that this is, Oh, we're just now using it as for knowledge in the art, that is nowhere in their invalidity contentions. Their invalidity contentions do not list any of the JEDEC materials, any of the QBM materials, or any of the APA as knowledge in the art. You will find it nowhere in their section on knowledge in the art.

This argument about knowledge in the art is an evasion that was created to avoid the consequences of the *Sotera* stipulation. We urge Your Honor to hold them to their word before the Patent Office.

THE COURT: All right. Thank you.

With regard to Document 211 and Plaintiff's motion for partial summary judgment on Samsung's invalidity defenses, I'll try to make this as clear as I can.

With regard to the JEDEC materials, I'm going to grant the motion in favor of Plaintiff.

With regard to the APA, the admitted prior art materials, I'm going to grant the motion.

With regard to the QBM materials, I'm going to deny the motion.

And I'm going to grant as unopposed that Samsung has no

anticipation defense as to the '918 and the '054.

Are there questions?

MR. SHEASBY: Nothing for Plaintiffs, Your Honor.

MR. McKEON: Nothing from the Defendant, Your Honor. Thank you.

THE COURT: All right. All right. Let's go to the marking issues, and we'll begin with Document 198, which is Samsung's motion for summary judgment.

And I think, Counsel, there's some efficiencies to be gained by arguing 199 concurrently with it, which is Netlist's motion. There seems to be a significant amount of commonality here, in my view. So I'd like to take these two up together.

We'll still start with Samsung, though.

MS. DEGNAN: Thank you, Your Honor. Lauren Degnan for Samsung.

THE COURT: Please proceed.

MS. DEGNAN: May I start with slide 6, please?

All right. Usually I don't spend a lot of time on the law slides, Your Honor, but I want to just highlight the second quote from *Amsted* here, that it's irrelevant for the §287(a) actual notice statute provision. It's irrelevant as to what the Defendant may or may not know about the patents or alleged infringement.

What we're looking for here is an affirmative act by Netlist to identify the asserted patents to Samsung and

include allegations of infringement. And I mention that because what we see in the briefing from Netlist is a whole lot of evidence that's really shifting the focus to what did Samsung know or should have known. And that's utterly irrelevant. So whether they knew of a family of patents is not relevant to §287(a); whether they knew about technology or portions of the specifications for family patents, not relevant.

I want to focus on §287(a), actual notice, and the focus is really on what did Netlist do to provide this notice.

If we could go to the next slide, please.

What we have here is Netlist's CEO, Chuck Hong, testifying in deposition that Netlist never sent any kind of notice of infringement or indication of infringement by Samsung prior to the termination of the JDLA we spoke about before lunch.

This is a binding admission by Netlist, and there's no genuine issue of material fact based on this admission that Netlist did nothing to provide actual notice under §287(a).

Can we go to the next slide, please?

And in his testimony, he mentions that there was an affirmative act for communication made to Samsung post-termination. We show that on slide 8. But what we see is that the affirmative act post-termination of the JDLA has nothing to do with the asserted patents in this case. So,

again, we've got no genuine issue of fact, there was no actual notice under §287(a). So that really should be the end of the matter.

Given these facts, I think Netlist, you know, does identify some documents, and I'm going to look at a few of them with Your Honor because even the documents they've identified, almost none of them -- there's five patents that don't mention -- aren't mentioned at all, and then the '060 Patent is mentioned. And we can talk about a couple of documents there and explain why that also does not satisfy Netlist's burdens for actual notice.

Can we have the next slide, please?

THE COURT: Let me ask you this. In light of the Court's ruling on the license defense before July 15th, 2020, how does that impact what's at issue here, if at all, in your view?

I mean, are we talking about -- I mean, if we're talking about either actual or constructive notice, pre-suit knowledge at a point in time when there's clearly a license defense as well, what do you see as the intersection between these two things?

MS. DEGNAN: So I think that's an excellent question, Your Honor.

So, again, for actual notice, there's needs to be affirmative communication by Netlist. And during the time

where there was a license, they really would have been precluded from even providing that allegation of infringement, which is I think why you see Mr. Hong admitting it so freely. Okay?

So for that, again, we're talking notice. We'll get to willfulness later where I think it might be different. But for actual notice, there really couldn't have been a charge of infringement during the term of the license, and there wasn't. So that covers a long period of time.

If I could have slide 9.

And this is one of the documents Netlist identifies. It's actually from an internal email of a slide perhaps they were thinking of presenting to Samsung. I don't think they've established that it was presented. But in any event, although it mentions the '060 Patent, there's no charge of infringement of any kind, no mention of Samsung products. So under §287(a), that's the end of the inquiry.

Can we have the next slide?

You'll see this chart, this patent list that Netlist created in June of 2016, and they cite as being sent to Samsung in November 2016, that does mention the '060 Patent. But, again, there's no allegation of any -- of accused products infringing this patent. And that is important under Section 287(a).

I will note that at this time the '160 Patent had issued,

but it is not on the list. So they can't rely on this list for the '160 Patent, and they shouldn't. They take a lot of solace in the idea that they're saying the '060 Patent covers -- covers HBM without further explanation. Of course, at this time, as you noted, Samsung's licensed so there couldn't be an allegation of infringement.

THE COURT: Where is the -- where is the element of asking or demanding an abatement of the alleged infringement on this?

MS. DEGNAN: It's not there at all. And with that, I'll move on, Your Honor.

We can go to the next slide.

We say this in the brief, and I won't belabor it, simply saying HBM is not an accusation of infringement of any of these products because there are multiple versions of the standard. HBM is not accused of infringement. Netlist has admitted it doesn't infringe the first HBM. The only kinds of HBM products that are even accused in this case are HBM2, HBM2E, and HBM3.

So with that, I will move on to slide 14, please.

And we'll cover the other -- other four patents, Your Honor. Again, they're not even listed, so this can't provide notice. But what we show here is Netlist identifies a family member of the '339 and '506 Patents. Again, that would be insufficient.

We go to slide 16.

It's the same chart. This time they're identifying a family member of the '918 and '054 Patent, again legally insufficient for actual notice.

I do want to stop on slide 19 because, in the briefing, for the first time Netlist makes this assertion, that they are seeking pre-issuance damages for the '054 Patent under Section 154(d). Not used that often, but it does allow for pre-issuance damages when certain criteria are met.

Importantly, this was never pled, it was never discussed in interrogatories. There was no discovery taken on this point, and we heard about it only in the briefing. So under Rule 37(c), this issue really should be out of the case, and they can't seek pre-suit -- any sort of pre-suit damages based on Section 154(d).

But we can go to page 20 of our slides, and we see this is all they said in the complaint. This would be insufficient because the complaint doesn't -- this does not evidence actual knowledge by Samsung of the published application. Again, not surprisingly, they didn't plead it. They didn't take discovery on it. They've got no evidence of actual knowledge of the published application.

They only have mentioned here and attached to their complaint certain claims from the prosecution history, and that is insufficient even under §154(d). But, first and

foremost, this issue should not be in the case at all.

There's one more document I wanted to mention because they do point to it on slide 21. And it's a particularly creative use of 287(a). This is a document, a list that Netlist produced as confidential -- not confidential in the CDCAL contract breach of contract action to Samsung's outside counsel. So this is a list produced to outside counsel that was marked confidential.

There's no evidence whatsoever that Samsung -- that Netlist has identified that anyone at Samsung ever received this, and under the protective order there's reason to think that they didn't. It doesn't list all the patents, and it certainly doesn't have any sort of specific accusation of infringement for the accused products in this case.

So I think the evidence is undisputed that there's no actual knowledge, actual -- excuse me, actual notice.

So let's switch to marking. I'll go to slide 24, please.

Now, this is where we have some overlap, Your Honor, with Netlist's motion. What we have is just a little quote from *Arctic Cat*. I'm sure Your Honor is familiar with it. We had a burden to discharge, and if we go to slide 25, we did.

I want to take this in a couple of pieces. I'm sure you read the letter, Your Honor. The letter starts out by identifying all the patents and saying, this is an *Arctic Cat* letter, we're going to tell you unmarked products that we've

identified so you can discharge your duty under *Arctic Cat*.

We then explain with reference to Netlist's allegations what kind of products they've accused of infringement so you can marry up what -- what products classifications go with which patents. It's not a mystery. We're working in the context of the litigation, and Netlist knows what we're talking about.

Then we go forth and we discuss specific unmarked products. We identified them to Netlist. The first one is Netlist's products itself, DDR4 LRDIMMs. And if you can look in all their briefing, they never dispute that they did nothing to respond with respect to their own products.

As a reminder, the DDR4 LRDIMM products have been accused in this case of infringing the '506 and '339 Patents. Nothing, complete silence, from Netlist with respect to its own products. They focus on the second part of that paragraph, SK hynix.

Go to the next slide, please.

But the letter does more than -- everything it needed to do with respect to *Arctic Cat* for the '060 and '160 Patents with respect to the SK hynix products. And then this is the part -- the issue that overlaps with Netlist's motion, the '060 and '160 Patents.

So, first, the letter identifies the patents. Now, by the way, Netlist consistently calls these patents the HBM

patents. It's not a mystery which products they're referring to with respect to these patents.

Then it goes forth in this paragraph we've highlighted and identifies specific SK hynix products that are HBM, first by class and then by specific models which are highlighted on the screen here. That is sufficient. That satisfies what *Arctic Cat* has told us to do to have a defined universe of unmarked products, which we did. It's not that big.

And that point in *Arctic Cat* is to avoid kind of a fishing expedition or gamesmanship, give the other side a universe to get back to you on, which is exactly what we did here. And what Netlist did was nothing--never gave back, never tried to prove non-infringement or non-coverage by those products, doesn't establish that SK hynix mark. There's no marking obligation in this contract.

What it does is it tries to shift the burden to Samsung, saying, Well, there's no claim charts, you haven't proven infringement. That is not our job under *Arctic Cat*. We have to give them a defined list. We don't have to prove infringement; we didn't.

They do talk about the next paragraph of the letter where we don't mention the '060 and the '160 Patents again in the next paragraph, but we're not required to repeat ourselves over and over again. We gave special emphasis in that last paragraph to the '506 and the '339 Patents because those are

the patents that Netlist asserts against LRDIMM DDR4 products, which are its own products. So Netlist deserves a little extra. Netlist itself sells DDR4 LRDIMMs. Emphasizing some other patents does not take away from the compliance with *Arctic Cat* with respect to the '060 and the '160 Patent.

So for the '060, '160, there's no genuine issue of material fact. There's nothing for the jury to resolve here. The letter says what it says, and it's really applying the law to those undisputed facts as set forth in the letter. So we would say this is ripe for summary judgment.

I can quickly finish up with respect to the other patents.

If we go to slide 28.

Again, this is sort of the slide that backs up they're trying to shift the burden to us to prove infringement, but really it was on them to identify specific limitations that were missing in these products. They never did so.

What it says in this briefing now is that, well, since none of the experts are going to be arguing about whether the standards practice the patents, the jury can find non-infringement by SK hynix. That's utterly insufficient for summary judgment when it's their burden to prove limitations are missing.

They took no discovery from SK hynix. We have no idea what limitation they are going to say is missing because their

experts don't say so, and they took no discovery. So, again, they have not met their burden with respect to proving that either these products did not have to be marked or they were marked.

If we can go to the next slide.

The same is equally true with respect to the other SK hynix products. They have nothing to establish that they don't practice.

And with that, I will sit down unless you have questions, Your Honor.

THE COURT: Not at this point, Counsel. Thank you.

Let me hear from Netlist on this, please.

MR. SHEASBY: Mr. Huynh, if I can have deck 7, please.

THE COURT: Go ahead, Mr. Sheasby.

MR. SHEASBY: Yes, Your Honor.

If I could have slide 31.

This is their *Arctic Cat* letter. It says that there may be products practicing the patents, one or more from the following list, and then it gives a list of products by SK hynix. As to our own products, as we previously disclosed in fact discovery, we stopped selling DDR4 LRDIMM long before these patents issued. So our products could never be a basis for a marking claim.

As to SK hynix products, they list a set of SK hynix

products. And then they, in the next slide, say that they're unmarked under the '506 and '339 Patents. They don't say that they are unmarked as to any other patents.

I understand that what they're saying is that that's a typographical error, that what they meant to say is that they're unmarked under all the patents. At trial, there is sufficient evidence, despite what they said. I don't think this was an accident. I think it was intentional only for them to present the marking defense on the '506 and '339. I understand they're now changing tune and they intend to present their *Arctic Cat* letter for all the patents.

And so I'll go through them in sequence, which is the --

THE COURT: Let's stop on this slide, Mr. Sheasby. Look at the last line of this paragraph. Don't they list all the patents there?

MR. SHEASBY: Yeah. They -- they do. I believe there's an ambiguity there. That's why, for the purpose of this argument, I'm assuming they complied with *Arctic Cat* under all the patents. I don't think they -- I think they did it intentionally, but for my argument I'm assuming that I should treat this as an *Arctic Cat* for -- for everything.

THE COURT: All right.

MR. SHEASBY: And so the evidence that we will show is that none of these patents are standard essential patents. These patents are read on an implementation-specific level.

And I'll go through the evidence that we have on that, and I'll start with the HBM because that's the most critical.

If you go to slide 37.

If you go to slide 37, you'll see that the -- what Doctor Brogioli talks about is the fact that the HBM product that SK hynix launched into implement something called a multi-drop procedure. Samsung's HBM implements something called a point-to-point or a twist-and-spiral design. The multi-drop procedure does not -- is not covered by our patents. The -- the twist-and-spiral design is covered by our patents.

And we will present that Samsung chose to go with our patented point-to-point design because of the deficiencies in SK hynix's HBM design.

The idea that just because something says HBM2 or HBM3, that that says anything about whether it's an internals user product. HBM -- the inner workings of HBM are not specified by JEDEC, they are not central, they are not standardized, and we're dealing with the -- with the internals of these designs.

So, in fact, what Doctor Brogioli is going to say at trial is SK hynix had an old design. It used -- it used multi-drop. Multi-drop was no good. Samsung switched to a point-to-point design, which is our patent. That is exactly what our burden is under *Arctic Cat*.

They can argue in response to that, but to suggest that there's not evidence in the record on the difference between

SK hynix and the Samsung design, the whole story here is that the multi-drop, which is what SK hynix innovated, is not used.

I should say that, as to the LRDIMM and DDR5 patents, it's the same thing. We have admissions from John Halbert, who's their JEDEC expert, that our patent relates to on-module power distribution, and the specific on-module power distribution that we claim, we are not asserting that standard. And, in fact, Mr. Halbert admits as well that the -- that how you design the internal blocks for your power management are unique to Samsung. And that's the argument we're making.

The same is the case for DDR4 LRDIMMs, which is that for the buffer designs, we're not arguing that they're standard essential; we're arguing that there is specific buffer designs that Samsung has selected based on our presentations to them to go in.

So I'm not going to belabor the point on the actual notice issue. The record is what is record is. I think if Your Honor doesn't believe there's been -- I think the reason why we think there's been actual notice is because it occurred in two phases, which is, first we said these are the designs we used, and then once the license terminated, we said you have to take a license. And so we think the combination of those two could go to the jury as actual notice.

But regardless of whether those two go to the jury on

actual notice, we think we -- we intend to present a robust defense to marking because we believe these patents are not standard essential and -- and rely -- and Samsung is using them specifically.

And it's -- the HBM is the exact story.

THE COURT: How does standard essentiality relate to marking requirement?

MR. SHEASBY: Yeah. So -- so the reason being is if we were saying these patents were standard essential, I could understand how it would be impossible for us to say that Samsung infringes and SK hynix had to mark -- didn't have to mark. That would be an impossibility--right?--because SK hynix practices the standard as well.

We're not saying they're standard essential. We're saying they're implemented specifically. And in particular as to SK hynix, we are saying the SK hynix multi-drop design is not the design that Samsung used. It's an old and weak design. The design that we use is point to point.

So now their next argument will be for them to present evidence from that SK hynix uses point-to-point, and at that point the jury will make the decision. But to be clear, this issue is joined. We don't think that HBM is all size -- one size fits all. Multi-drop is the alternative that was used by SK hynix, and Samsung presented our design -- Samsung decided to use our design.

THE COURT: Let me ask you this. In your mind, is Samsung alleging any application of standard essentiality to the patents-in-suit?

MR. SHEASBY: They're not. Samsung is --

THE COURT: SEP is just off the table here, is it not?

MR. SHEASBY: Samsung claims they're not essential.

THE COURT: Okay.

MR. SHEASBY: So I -- I want to be very clear, which is that actual notice -- I understand where the Judge's views are on that, I think, and that you think it's not specific enough. I'm focused on something more basic, which is, we don't think all these products are equal. Our whole story for why Samsung infringed the HBM patents is because SK hynix multi-drop was no good. And it's exactly what Doctor Brogioli is going to present.

Have I answered Your Honor's questions?

THE COURT: I think so.

MR. SHEASBY: Thank you, Your Honor.

THE COURT: All right. Ms. Degnan, do you have something further?

MS. DEGNAN: I do, briefly. Let's start with slide 3.

Just to reiterate, there's nothing that they disclose in discovery about this theory about multi-drop being something

that SK hynix does. And even in its briefing, it says that's something that SK hynix HBM1 products do. They made no showing, even slide 37 from -- that they just showed you from Mr. Brogioli, there's nothing in his expert report -- oh, excuse me, sir, Your Honor.

There's nothing in his -- in this slide or in his expert report that ties this multi-drop thing to the specific models that we identified in our *Arctic Cat* letter. It doesn't say that. Again, the -- the briefing talks about HBM1, SK hynix is not practicing the patents. We identified HBM2E, HBM3. There's no indication of what limitations are not being practiced by those products as is their burden to prove.

So I just wanted to make that point that this idea -- again, they're trying to shift the burden to us to prove that the products practice. That's not our burden. We've identified, given them a list. They've got to tell us why -- what limitations are missing. None of that has been talked about today in the briefing or in the expert reports.

And with that, that's really the one point I wanted to make for Your Honor.

THE COURT: All right. With regard to Documents 198 and 199, particularly with regard to Samsung's motion for summary judgment of no damages prior to the filing of the complaints, Document 198, the Court's persuaded that Samsung met its burden under *Arctic Cat*, which is relatively low, and

shifted the burden to Netlist, and that burden having been shifted, I'm persuaded that the result of the motion is as follows:

With regard to the '506 Patent, I'm going to grant Samsung's motion. I don't find any notice, actual or constructive, prior to filing the original complaint on December the 12th, 2021. Hence, I'm going to find as a matter of summary judgment there are no damages prior to that date.

On the '339, it's the same ruling by the Court--no damages prior to December the 12th, 2021. And it's the same on the '918.

On the '054, Samsung asks for a finding via summary judgment of no damages prior to filing the first amended complaint on May the 3rd, 2022, but I am persuaded there's an issue as to potential pre-suit damages prior to the filing of the first amended complaint but not subsequent to the original complaint.

So I'm going to grant summary judgment that there's no pre-suit damages or there's no damages prior to the filing of the original complaint on December the 12th, 2021, as to the '054 Patent.

On the '060, there's no notice prior to filing the first amended complaint on May the 3rd, 2022. So it's the Court's ruling via summary judgment there are no damages prior to May the 3rd, 2022, on the '060.

On the '160, there appears to be no notice prior to the original complaint, so there are no damages by way of summary judgment finding by the Court prior to December the 12th, 2021, on the '160 Patent.

As to Document 199 and Netlist's motion to strike Samsung's affirmative defenses, I'm going to deny that motion.

Now, it's my understanding that Samsung has withdrawn both its exhaustion defense and its implied license defense, and either by withdrawal or failure to come forward with adequate evidence, I'm going to grant summary judgment in favor of Netlist with regard to both exhaustion and implied license as defenses otherwise available to Samsung.

Are there questions from either party with regard to that ruling?

MR. SHEASBY: Your Honor, just to clarify, are we entitled to present evidence that we did not have a marking obligation as to SK hynix HBMs to the jury?

THE COURT: Say that again, Mr. Sheasby?

MR. SHEASBY: Are we entitled to present to the jury that we did not have an obligation to mark the SK hynix HBMs?

THE COURT: I don't think so based on the ruling I've just given you.

MR. SHEASBY: Okay. That was the clarification I wanted. Thank you, Your Honor.

THE COURT: Any questions from the Defendant as to

the Court's ruling?

MS. DEGNAN: Forgive me, Your Honor. Could we get a little more explanation as to why the '160 is being treated differently than the '060 with respect to the original complaint? I'm not -- it's not immediately apparent to me why that's the case. We respect your ruling absolutely. But so we can plan our evidence accordingly, if we could understand that, if there's any additional explanation you can give.

MR. CORDELL: The question is pretty precise, Your Honor. The '060 and '160 were presented together in the amended complaint.

THE COURT: Clarify for me. Were those two -- were either of those two in the original complaint or were they both added in the first amended complaint?

MS. DEGNAN: They were both added in the first amended complaint.

THE COURT: Okay. So I'll -- based on that clarification, I'll revise my ruling that the '060 and the '160 are both situated such that there are no damages prior to the date of the first amended complaint on May the 3rd, 2022.

MS. DEGNAN: Thank you, Your Honor.

THE COURT: All right. That brings us to Samsung's motion for summary judgment of no willfulness, Document 200.

Let me hear from Samsung, please.

MS. DEGNAN: Thank you, Your Honor.

So this is slightly different. As you know, the standard is a bit different. We do look at Samsung's own knowledge here.

The way the briefing is set up is we talk about knowledge of the patents, the accepted prerequisite, because if you don't have knowledge of the patents, you can't have the -- you can't have the deliberate and intentional infringement of those patents. Okay? So we do focus on the lack of knowledge of the patents here because for five of them there's no evidence that there is actual knowledge of the five patents.

Of course, we do talk about the good faith defenses, and we'll get into that, but I just wanted to start with that kind of overview.

If we go to slide 6, please.

As I mentioned, for five of the patents, there's no evidence of pre-suit knowledge of the asserted patents. There is that some -- that patent list we looked at just a moment ago--if we move to slide 7--with respect to the '060 Patent.

Can we go to the next slide?

Here I am going to just take you through -- again, this is a chart. It was created in June of 2016. At that time Samsung had not even made an HBM2 accused product. Those were launched later. And just to point out, that HBM1 is not the same as HBM 2. This is a chart that talks about the '060 Patent, but it does not, in fact, really discuss why Samsung

would be infringing.

And, of course, as Your Honor pointed out, during this time period, we are licensed. And I think that's an important issue for the willfulness arguments that, since we are licensed, there's no reason to -- for us to believe that there would be -- that there literally could not be deliberate, intentional infringement in this time frame.

So if we move on to --

THE COURT: Just a minute, Counsel.

MS. DEGNAN: Sure.

THE COURT: This chart on the right side of the slide that you have, that was sent to Samsung by Netlist. Correct?

MS. DEGNAN: It was.

THE COURT: And what is your understanding of the rationale or the impetus behind the delivery of that chart?

MS. DEGNAN: So I think at this time we are licensed, Samsung is licensed. We are engaged in the collaboration. Remember the joint development collaboration process? Towards that end, I think it was maybe -- Netlist was, you know, happy to be working, wanting to show -- I'm not sure why.

Samsung was interested. They sent some information. But it was in the collaboration -- it was the part of this collaboration of the joint development effort. It could be,

you know, as we're working together, it's just sort of a sharing of information. It was not obviously accusing anyone of infringement because we were licensed.

THE COURT: All right. Now, if you consider the Court's prior ruling that up to the termination of the JDLA on July the 15th, 2020, Samsung was licensed pursuant to the terms of the JDLA, then we're not talking about willful infringement prior to July 15th, 2020, under any scenario.

Is that your understanding?

MS. DEGNAN: You are absolutely right, Your Honor.

Could we have slide 27? There's a timeline I think --

THE COURT: So really we're talking about 2020 until suit's filed.

MS. DEGNAN: We are. And so let's just jump into this, what is the good faith defense, because that is, of course, a reason that would prevent a willful finding here. And we think no reasonable juror could find we're willful infringers given -- given this license. So, yes, certainly through the termination on July 15th, 2020.

But I think it's important to recognize is that between July 15th, 2020, and the final judgment in the California case, this was a hotly-disputed issue. They were claiming termination, we were saying it was unlawful, we thought the termination even if there was a breach, it was not material justifying termination.

So during this period of time, subjective intent was of -- of Samsung was that we are being wrongfully denied our license and we're fighting to keep it.

And during this period, we should point out that there was a jury trial, and the jury awarded on December 3rd, 2021, no damages for this breach that Netlist said happened. That, again, looks like it's not material. If there's no damages, then there shouldn't have been a termination.

And so we would say that this entire period through at least the end of the action in California, no reasonable juror could find that we were -- had the subjective intent to deliberately infringe because we had a license, we were fighting to prevent Netlist from taking it away during that entire period.

And I would submit even after, while we appealed what we would consider a miscarriage of justice, to the Ninth Circuit we still had a good faith belief that, you know, we have been denied a license that we bought and paid for, it was a perpetual license, and that that prevents any sort of reasonable jury from finding infringement both pre- and post-suit.

THE COURT: Let me go back a minute since you've got this timeline on the screen.

Now, I don't want to limit or backtrack on my prior statement that, during the period of licensure, Samsung really

can't infringe if they're licensed, and if they can't infringe, they can't willfully infringe.

But I have left open the issue for the jury to decide whether these HBM products are foundry products which would not be covered by the license and, therefore, could be subject to infringement. And if those can be subject to infringement under that construction or scenario, then there's potentially willful infringement.

So as to what's covered by the license prior to 2020, I'm satisfied that willfulness is out. What's not covered by the license prior to 2020 is fair game for willfulness. And then we had the issue of post-July 15, 2020, up until the time suit's filed.

And I understand your arguments about your sense of justice in the Central District of California and the Ninth Circuit. You are entitled to your opinions. I'm not sure that's dispositive on the issue of willfulness or not.

MS. DEGNAN: So let me just respond to this issue of products that were foundry products. And so I think what -- I think what we're going to see --

THE COURT: I can see that fight coming.

MS. DEGNAN: I can see --

THE COURT: I think everybody in the room can see that fight coming.

MS. DEGNAN: But to clarify, it's not really the

whole class of HBM products were foundry products, but whether any specific model was designed for somebody else and then sold without our brand on it. So there will be huge swath of --

THE COURT: Oh, there will be factual disputes about all that.

MS. DEGNAN: And then to the extent -- to the extent there's evidence that Samsung specifically knew it had a foundry product and sold it anyway, knowing it was unlicensed, for that small amount of sales, if there are any, which we would say there aren't, there could be willfulness. But it's not sort of a sweeping we're going to have willfulness for the entire class of HBM products. You have to prove that we snuck a foundry product out and it infringed and it wasn't licensed.

THE COURT: I don't know that there's any element of sneakiness required, but if there is a product that doesn't fall within the license and the jury's persuaded that it doesn't, then that opens the door to a potential finding of willfulness and a potential finding of -- excuse me, a potential finding of infringement and, subsequent to that, a potential finding of willful infringement.

What I'd really like to focus on now is post-July 15, 2020.

MS. DEGNAN: Absolutely. So what I would say -- what I would say there is prior to the

complaint--right?--we're just talking pre-complaint, is because we -- the -- the subjective intent. We can't be deliberate infringers if we subjectively believed that we are still licensed because the termination is unlawful. That takes us all the way through February 15, 2022, at which time the original complaint had already been filed.

So, you know, we're talking now post-suit, and we can talk about the post-suit stuff. There's a couple of patents where the '060 and the '160 that are in the amended complaint, so that would be post-suit.

So this sort of covers the landscape on pre-suit.

If we can go to slide 30, we can talk about post-suit.

Again, because of the amended complaints, there's a little bit of, you know, time difference. But, again, what they're pointing to post-original complaint and post-amended complaint is really the garden variety continuing to sell your product.

And the Federal Circuit said that the deliberate, intentional infringement, it's not something in every single case, that garden variety continuing to sell products is enough that kind of upsets the entire balance of what willful infringement is about.

If we can have the next slide, 31.

And, again, we have a variety of differences. We've got the license defense we've already discussed, you know, both

during the litigation in CDCAL and on appeal. That is evidence of Samsung's good faith defense that it should still be licensed and the termination was wrong.

We also pled invalidity and non-infringement, and we've developed those defenses during discovery. And, importantly, we have invalidity that is evidenced by our IPR petitions that were filed, four of which are already instituted. That is further evidence that we have a good faith defense in invalidity which would bar any post-suit finding of willfulness. And so we would point to that.

And with that, if Your Honor doesn't have questions, I might sit down and hear what my colleague has to say.

THE COURT: All right. Let me hear from Netlist, please.

MR. SHEASBY: The issues that we can speak about, and I'll speak about them in order, which is, from the time of termination, the question is, was Samsung's conduct willful or was there evidence from which a jury could find that the conduct was willful. And I'll focus on that to begin with.

And so the question is upon that time of termination, did Samsung have knowledge or was it reckless that it was using Netlist patents. And I would submit to you that there's overwhelming evidence that it knew that it was using Netlist patents.

So for the '054 and '918 Patents, in July 2014 we

actually presented them with information that we had the -- we presented them with information that we had intelligent on-module power distribution patents pending.

And to be clear, they had hired us and contracted with us to assist them in designing their NVDIMM product which had power management on-modules. And what they did is that they transitioned that design that we were presenting to them to DDR5. And this is Jung Bae Lee, the president of Samsung memory, speaking about that fact.

So the NVDIMM technology, which is the on-module power management, was ported into the DDR5 products. And DDR5 products are not just for NVDIMM, but now for RDIMM, for LRDIMM, for NVDIMM. But they all use the similar on-module power management that they were on notice we had patents pending on at the time.

In 2019, they candidly admitted that they needed our technology because there was a patent risk associated with RDIMM and LRDIMM products that they were selling. The DDR5 patents-at-issue, the on-module power management that we design that -- that we designed, they use now in its generic to RDIMM, LRDIMM, and NVDIMM.

So there's clear knowledge that they were conscious of the fact that they needed a license to our patents. And, in fact, they admitted that they took a license to our patents or they entered into the JDLA immediately because they might

pursue lawsuits against them for using the patents.

We were very upfront with them that we had family members that related to hybrid technology. The PMIC, power management on-module that's used in DDR5, is a hybrid technology. The patents we asserted are children of the patents we disclosed to them.

And we particularly, in 2016, sent them a chart saying these are the products that are covered by our patents. And you see that we are very emphatic about that. It's the products covered by our patents.

Why were we telling them that in 2016? The reason why we were telling them that in 2016 is because we were trying to get them to cooperate under the agreement.

This is one of the documents that will be shown at trial. Our EVP went to them and -- went to Samsung and said, company N--this is an internal Samsung document--company N's patents is applicable to RDIMM, HBM, and Enterprise SSED as well as, in addition to, LRDIMM. So we plan to obtain relevant information and review it.

So this was part of us trying to tell them, listen, you guys have to change your ways; you have serious, serious problems.

The patents that we disclosed there, the '833 and the '831, we list that those are the parents of the patents that are asserted in this case. And they knew we had technology in

this space. They actually came to us and asked us to help design their on-module power management that was at issue in this case.

THE COURT: Go back two slides, please. No. Right there. No, that's not what I'm thinking of. I'm thinking of the similar chart that's got the '060 in it?

MR. SHEASBY: Yes, sir. It's the same chart. The '060 is on the bottom.

THE COURT: I asked opposing counsel what her understanding of from Samsung's standpoint as to why that was sent from Netlist to Samsung. What's Netlist's --

MR. SHEASBY: If I could have -- it's this, Your Honor. It's us telling them, you need to get your -- your act together.

THE COURT: I did say Netlist, didn't I, not Netflix?

MR. SHEASBY: I heard Netlist.

THE COURT: Okay. I'm conscious of not misspeaking.

MR. SHEASBY: I wish I was more conscious of that, Your Honor. It would help.

This is us telling them, the reason why they got this chart is we told them, you need to get your act together, you need to start complying with this agreement, you need to act with us in good faith and not stiff us on product. Our EVP went and told them that.

And then this -- let's go back.

This is the document that was after this. So this is them being on notice, your HBM product, it says covered product, for example, and it lists the '060 Patent. It says, covered product, and it lists the parents of the '054 and '918 Patent.

And so what happened is they let the agreement terminate. And from that moment in time, knowing our position was explicit that their LRDIMM hybrid design products and HBM products were using our patents, they continued to use them and didn't abate. And that, I would say, is the quintessence of willfulness.

And let me give you another example of this. They said they entered into the JDLA because they were afraid they were infringing our patents. If they entered into the JDLA because they were afraid they were infringing our patents, upon the termination of the JDLA and letting it terminate is the epitome of recklessness as to this idea.

This is not a situation in which, oh, we're shocked that you had these patents. They admitted candidly they entered into this agreement because they thought they infringed. When they lost the agreement, that makes clear that their conduct at that point in time becomes reckless. So that's as to the DDR5 patents.

As to the HBM patents, once again we list it as covered

products. I want to be very clear about something. This distinction between HBM1, HBM2, HBM3, Samsung only sells one class of HBM--it's point to point. We cover point to point. And in November of 2016, they were preparing to launch and announce their point-to-point technology.

So this -- and there's a debate as to whether HBM patents are licensed. But even if they were licensed, when they let it terminate, they became willful, having been put on notice that our products cover their -- our patents cover their HBM products.

THE COURT: You keep saying they allowed the JDLA to go away. Netlist is the party that terminated the agreement. Correct?

MR. SHEASBY: Oh, we asked them to cure. They refused to cure.

THE COURT: That's what you mean by that statement?

MR. SHEASBY: Yes, sir. Yes, Your Honor.

THE COURT: Okay. That clarifies for me.

MR. SHEASBY: So we did not raise that and terminate it. We asked them to cure. They refused to cure. When the cure period ended, we terminated.

So then as to the HBM patents, we list it repeatedly as being relevant to the hybrid technology that we're accusing in this case. Jung Bae Lee, the president of Samsung, concedes that he received this presentation in which we talk about our

patents covering TSV. This is the exact design of their product. Their product uses TSV, and it has each of the bullet points that we list as being relevant to the -- our patent. And they have them all.

So it's the HBM -- and, by the way, our inventor, Mr. Hung Lee, testifies under oath that he showed this design information to Samsung. So as to the '060 and '160 Patents, there is clear and, I think, very explicit evidence that for them not to cure, for them to continue to not comply with the agreement and to let it terminate, was reckless.

The next issue is LRDIMM. This relates to the '309 and '506. They entered into the agreement because of patent risk related to LRDIMM. We are accusing their LRDIMM products. We list the parent of our patents-in-suit as covering specifically LRDIMM.

Samsung began distributing the LRDIMM with distributed buffers after we put them on notice that we had patent -- we had coverage on this technology. We list the parents of the '056 and '339 as covering LRDIMM.

And I want to focus specifically on this. This is what they told us in a questionnaire that they sent -- that they copied us on that they sent to the Korean equivalent of the IRS. "It is well known in the industry as a company that created a system for LRDIMM technology. What did you mean by that?"

It's -- and what he goes on to say, he's the person that did the JDLA with us, it means distributed buffer. The Samsung design has distributed buffer. What we're accusing of infringement in this case has distributed buffer.

So in 2020, we terminated the license agreement and warned them that they needed to take a license. We then sent them a list of documents to the people who subsequently filed a lawsuit against us based on the letter that -- that -- that we sent talking about our portfolio.

So it was their licensing counsel. And that licensing counsel received the list that included a large number of the patents-at-issue or their parents. So that list included the parent of the -- of the '339, it included the '506, it included the '160, it included the '339, and it included the '060.

So in my mind, the way you look at it is there was this relationship that was created. It was created, Samsung candidly admits, because they were afraid of patent infringement. They stopped complying with the agreement.

We warned them, Guys, you have got to get your act together, we have patents that cover your HBM, we have patents that cover your NVDIMM, we have patents that cover your LRDIMM.

They did nothing, they did not cure, they let the agreement terminate, and at that point in time they became

reckless with that knowledge that they needed our patents, and they continued without abatement.

Does that answer Your Honor's questions?

THE COURT: I don't have any additional questions, Mr. Sheasby.

MR. SHEASBY: Thank you, Your Honor.

THE COURT: I'll allow you to respond, Ms. Degnan, and try to be brief.

MS. DEGNAN: I will. I'll just try to -- know if I can get through all the slides, but I'm going to at least talk big picture.

THE COURT: That's probably a good idea.

MS. DEGNAN: Yes. So I want to talk -- maybe I'll start with this idea of the cure. And I think it's telling what Mr. Sheasby has said is that this list which they say was sent after they told us to get our act together was in 2016. And earlier Mr. Sheasby is like, oh, no, they stopped -- they stopped complying in 2018. So the timeline makes no sense.

There is no -- and, again, they sent us a termination letter in 2020. How could we cure anything in 2018 if we got notice of it in 2020? And that's just the supply provision. There's another part of the agreement they said we breached having to do with the tax withholding. Again, we were put on notice long after the refund had already been made.

So this idea that we let the agreement expire because we

didn't cure is, frankly, dishonest because we had no opportunity to cure things that happened multiple years in the past. So that's just not right at all. I just needed to clarify the record on this.

I also want to talk about this notion that -- you know, they love to talk about our statements about why we entered into the JDLA was to avoid the risk of being sued, much like we're being sued today, not because there was a risk of actual patent infringement.

They were a litigious company at that time. They made veiled -- you know, they made it known that that's what they did. And so there's no -- nothing to be seen by the idea, there was no admission that we needed to get a license for infringement.

THE COURT: What's all this got to do with Rule 56? What's your response to that?

MS. DEGNAN: Well, I'm just sort of responding to these arguments at a high level.

THE COURT: Well, I understand opposing counsel's arguments may elicit a response, sometimes a visceral response, from the other side, but we've got to focus on the issue.

MS. DEGNAN: Sure. And let me tie it to the issue. They're saying, because what we saw back in 2015, when the license was terminated in 2020, we somehow then had the mens

rea subjective intent to infringe patents that hadn't even issued during the negotiations and the license for the JDLA.

So you can't tie this activity that they spent many minutes showing you in the slides to the mens rea requirement post-termination of the license. So there's nothing that springs back into effect at that time.

And so it's my way of explaining why all the slides you just looked at, all those evidence, is actually not relevant to the issue of willfulness because those patents had not issued. And by the time they had issued, even if there was some desire to license a patent that was in existence in 2015, that does not mean we had the subjective intent to infringe completely different patents in 2020.

So I guess I'll leave it with that. That's sort of high level response to counsel's remarks.

THE COURT: All right. Thank you, Counsel.

All right. With regard to Document 200 and Samsung's motion for summary judgment of no willfulness, as to the period prior to July 15th, 2020, and as to what was covered by the license granted through the JDLA, the Court grants summary judgment that there was no willful infringement during the period that whatever was licensed by the JDLA remained licensed by the JDLA up until and through or as of July 15th, 2020.

Post-July 15th, 2020, the Court finds that there are

material questions of fact under the totality of the circumstances standard imposed on the willfulness determination that preclude the entry of summary judgment, and I'm otherwise going to deny the Defendants' motion.

Now, that should cover all of document 200. We've got a series of motions to strike coming up next. We're going to take a short recess, and we'll turn to those when I return to the bench.

During this recess, I'd like to see Mr. Cordell and Ms. Smith together with Mr. Sheasby and Ms. Truelove in chambers.

The Court stands in recess.

(Brief recess.)

THE COURT: Be seated, please.

All right. As I said, Counsel, there's a series of several motions to strike. I don't know that the order is particularly critical.

Why don't we jump into the deep end of the pool and start with Document 205, Samsung's motion to strike Kennedy reports, Kennedy being Netlist's damages expert.

Let me hear from Samsung first.

MR. CORDELL: Thank you, Your Honor. Ruffin Cordell for Samsung. May I proceed?

THE COURT: Please do.

MR. CORDELL: Mindful of the time that we have, I will offer my apologies on the front end that this is not the

simplest of motions. Mr. Kennedy's, you know, thousand-page report departs from precedent, it is filled with never tried before approaches, and is justifiably a bit over its skis on almost every conclusion it reaches.

So, again, with my apologies on the front end, let me see if I can dive in.

The nine issues I have on slide 3 are our primary objections to his reports, and there are nine separate substantive objections. And I wish I could tell Your Honor that I could easily dispose of two or three, but I can't. So let me at least attempt to begin at a higher level and explain what our position is.

So Mr. Kennedy does not really employ any of the typical damages approaches that we've seen. We typically in this court see damages experts take the stand, they look at four or five comparable agreements, they explain the value of the technology, and they come to a damages number. Mr. Kennedy really doesn't do that.

And we challenged -- in our motion practice, we challenged Netlist to explain where he got this stuff, and what we got were explanations from the lawyers rather than from Mr. Kennedy. And so we're forced then to try to deduce exactly where he got these things and where -- you know, essentially what the basis of his opinions ultimately are.

So there are a few unifying principles. Mr. Kennedy,

time and time again, starts with a non-infringing alternative and says that non-infringing alternative is the baseline from which his damages calculation to be judged. He looks at a Lexus that might be patented and he says, I'm using a Toyota Corolla as a non-infringing alternative. You can drive it; therefore, it's a non-infringing alternative. And I'm going to phrase all of my damages in the differences between a Corolla and a Lexus.

He never -- and Netlist never offers us any evidence that the non-infringing alternatives are acceptable or viable, and that's what the law tells us. If you are going to do a non-infringing alternative analysis, you're supposed to tell us why it's acceptable, why customers would buy it.

They also never disclosed in I of this in discovery. In a moment I'll show Your Honor an interrogatory response when we challenge them on the bases of their damages case and what they gave us with a couple of sentences per patent family, and that's it. And none of what Mr. Kennedy says is in that interrogatory response.

He then takes the difference between a Corolla and a Lexus, and he applies what we now call the hundred percent rule. So he said, well, Samsung, you are making more money selling the patented product than you do what I say is a non-infringing alternative. The Lexus is selling for more than the Corolla. I'm going to take 100 percent of the

difference between those two products, not in terms of profit but in terms of revenue, and I'm going to award that to Netlist.

So recall where we're supposed to be in the damages case. We're supposed to be in a hypothetical negotiation. We're supposed to be in a situation where rational people are sitting across the table. And what Mr. Kennedy's fundamental opinion is, is that Samsung would agree to give 100 percent of whatever increased revenue they get from using the patented invention, and not just the patented invention, the patented product, the entirety of the product, I'm going to give a hundred percent of that revenue to Netlist.

There has never been a case that we've found where an expert has suggested that that was an appropriate approach.

Perhaps as a stalking horse and understanding that his 100 percent rule analysis is flawed, Mr. Kennedy gives us a couple of fallbacks where he does actually at least say he's going to apportion that 100 percent. He's going to apportion the difference between a Lexus and a Corolla not by looking at the contribution of the patent and saying, well, that's 1 out of 10,000 developments that allowed this new memory product to be put on the market.

He doesn't look -- he doesn't do a survey to find out if -- you know, we -- we heard Mr. Sheasby tell us the direct through approach. He said that over and over again. You

know, what's that worth? He didn't -- we have none of that.

What we have instead are two odd approaches, one where he compares Samsung and Netlist's R&D investments company-wide and then uses that as a proportion to award damages.

So imagine that for a second. He said, Well, you know, the Netlist R&D investment as a proportion of its overall monetary health is X, and Samsung is Y, including their R&D investment in washing machines and completely irrelevant products, and he uses that as a metric to do an apportionment. But it has nothing to do with the case. It's almost like he's struggling just to not award a hundred percent so he came up with these arbitrary methodologies.

It's unprecedented, Your Honor. It exposes this Court to error. So if we all go along with this, we're going to be at the Federal Circuit trying to explain this, and it's not useful for anyone. Mr. Kennedy's approach simply doesn't work.

THE COURT: I suspect there's a very high probability that we're going to be at the Federal Circuit at some point regardless, so that's not a big persuader.

MR. CORDELL: Understood. Just as officers of the Court, I feel like it's our obligation to make sure that we put good sound legal arguments before Your Honor and before a jury, and that's not what Mr. Kennedy is doing.

THE COURT: All right.

MR. CORDELL: So let me show you that interrogatory response. In Interrogatory No. 11, we asked them tell us about your damages case. Tell us, you know, if it's going to be a royalty, tell us how you're doing it, give us a royalty base, the measure and proper calculation, how you came up with it, your contention of the terms and the conditions. You know, tell us about your damages case.

What do they tell us? Well, remember there are three patent families in this case. For the DDR5 patents, they began with that one, and they say, well, these -- these inventions have tremendous technical value. Okay. Well, that's a platitude, but at least we have something.

And they say, without voltage regulation on board, the required tolerances would not be possible. Okay. Well, that's probably worth something. Maybe we can work with that.

They then say that, as acknowledged by Samsung, voltage regulation enabled 30 percent DIMM level power savings. Okay. So power savings. So that's where we thought we were. That's what we thought we were going to get from their damages case. We were going to get an analysis of what these tolerances were worth and what some power savings were worth. And there's an exposition about fewer ground pins or supply pins, and that might be something, too. That's where we thought we were. And that's it. That's it.

We heard nothing about a non-infringing alternative. We

heard nothing about the Lexus versus the Corolla. This is what we were told in discovery.

When we get to the HBM patents, it was particularly terse, and they said, well, the HBM patents are relating to reducing capacitive loading in 3D stacked array dies. Okay. Well, reducing capacitive loading, that might be worth something. That's what our expectation was. We expected Mr. Kennedy to get with the technical experts, find out what that really means, and then figure out what that was worth.

And then for the DDR4 patents, we had PVT variance in DIMMs from a distributed-buffer architecture and that -- again, that's going to be worth something. That will enable buffer paths only during specific time periods, and that that might be worth something. But that's what we were told.

We heard none of this about -- about a comparison with non-infringing alternatives and, critically, Your Honor, they never identified these non-infringing alternatives.

And what does that roll up to? Well, it rolls up into massive, massive numbers. So instead of finding a reasonable royalty, for example, for the reduced capacitance and HBM stacked memory, a reasonable royalty that would be attached to the footprint of that invention, what they've done by using this non-infringing alternative fantasy is come up with numbers that run into the billions of dollars.

So for the HBM license, now, Your Honor, I admit these go

back to the beginning a bit, so they're going to be changed based on today's rulings. But the reality is they will run into the billions.

And then, importantly, Mr. Kennedy doesn't just stop with past damages. His numbers include damages going into the future. He assumes that whatever the jury finds will persist until the expiration of these patents which will be for some cases decades into the future. And he rolls all of that up into a number that is almost $3 billion. And if you add the past with the future, you get to $4.2 billion. Again, unprecedented.

The Court is to set the forward-looking royalty after trial. We've seen that in other cases. The *Pace* case set it -- set it up for all of us. It's not the damages expert who is here to ask this jury to award damages as if we know what those sales are going to be, that we're able to predict into the future that there aren't going to be further generations of memory, for example, that replace DDR5, as DDR4 was replaced by DDR5, et cetera. So there's a fundamental flaw in his approach.

All right. Well, let me get into my issues. And, again, Your Honor, I will do my best to clip through these quickly.

We know the law from the *CSIRO* case, from the *Uniloc* case. I won't belabor this. We're supposed to tie the damages opinion to the actual footprint of the invention,

whatever the invention contributes. And this is a crowded art. Make no mistake. That JEDEC committee is not working on just one concept at a time. They're taking patents, thousands and thousands of patents, and incorporating them into each of the -- the new -- new generations.

We know that apportionment is not optional. This is not a case where we can look at the Lexus and say, okay, we've got a patent on a new fuel injector and there's going to be six of them in the car and that fuel injector is going to -- is going to help with the power of the engine. Okay. That's all great.

But the Federal Circuit says we have to apportion. We can't just look at the value difference between the Corolla and the Lexus because there's a lot that went into that Lexus that has nothing to do with the fuel injector, and the same is true here.

THE COURT: I understand the principle of apportionment.

MR. CORDELL: Okay. Sorry.

So what we start with is this non-infringing alternative. And the non-infringing alternative for the HBM patents are, he says, the HBM2 patents. And he says, well, you know, those were slower, they weren't as good, they were a Corolla compared to the -- the HBM3 patents and higher, or 4 or higher, that -- that move quickly and require less -- less

power and, therefore, we can take a hundred percent of the value and we can make that the baseline of his comparison.

What we don't have is any evidence that today a customer would even accept a 4 high HBM product. None. So the evidence -- Mr. Kennedy says he doesn't know whether that's an acceptable alternative at all, and it's his burden. It's his burden to prove to us that, in fact, they would be acceptable to customers.

And then here's his calculation on slide 17, Your Honor. He says, well, you would have sold $771 million worth of -- of HBM2 and you would have replaced that with the -- with the 8 high HBM memory, you would replace that with HBM 4 high, which would have only netted you $380 million in revenue, and so the -- the difference is the loss -- he says, the loss that Samsung would have suffered but for improving their product and selling the latest generation is almost $400 million. He then turns around and gives a hundred percent of that to Netlist.

So he says, you're selling an improved product, you made more revenue off of the improved product, a hundred percent of that goes to Netlist. There's not a stitch of apportionment, there's not a suggestion of apportionment, and, critically, it doesn't even reflect reality. I mean, there's no -- there's no real suggestion that the old product would have sold and been a -- been a replacement at all. The world moves forward.

You just can't do that by looking back.

So he -- he does the same thing for DDR5. He says that we can use a fantasy product, this one where we're going to slow DDR5 by 20 percent to get less power utilization. He then uses a regression that he relies on Doctor Groehn for to try to come up with a number that is associated with that -- the difference in value. But then, once again, he gives a hundred percent of that to Netlist.

And here on slide 19, I show that he goes through the same kind of calculation coming up with what he says would be a $207 million revenue differential between Samsung selling the patented device under what he's -- compared to what he says is the unpatented device and awards a hundred percent of that to Netlist.

For DDR4, he does the same kind of thing. Here it's a little more apt because Doctor Groehn actually did an analysis of DDR4. He didn't do DDR5 at all. And so we have an issue with Doctor Kennedy's use of Doctor Groehn's DDR4 analysis in the 5 context. But it's the same song. He ends up awarding a hundred percent of what he says is the alternative revenue to Netlist.

The thousands of advancements that go into these features are not accounted for. He didn't do any of the -- the surveys that Your Honor hears in a lot of these cases to try to understand value. He said, I am going to pick an arbitrary

non-infringing alternative that Mr. Kennedy made up and I'm going to say that that is the baseline from which I can then create this damages model and award a hundred percent of the revenue to -- to Netlist.

We now have a hundred percent rule where the Federal Circuit has shot down the 25 percent rule and the 50 percent rule. We're now talking about the hundred percent rule. And, critically, that's revenue, not profit.

Uniloc and VirnetX, I won't go back through. Your Honor is well aware of these. There is just no authority for this. It is exposing us all to ridicule.

So what does he do? Perhaps again his hundred percent rule may be a stalking horse. They knew that I'd have to get up here and spend time talking about it, and maybe they don't want to pursue the hundred percent rule, so they give us some sort of apportionments. But the apportionment methodologies that Mr. Kennedy chooses are completely unprecedented and really just make no sense.

So he has two different ways of doing this. One is he looks at the gross margin of Samsung and says, well, you've got some profit and you've got some costs, and I'm going to assume that your costs are what you contributed yourself and the profit is then going to be -- I'm sorry. I got that backward. That the profit is for what you contributed and the costs are what the patent deserves. It makes no sense, but

that was the ratio that he utilized.

And then the other one I mentioned earlier, he looked at the R&D percentages of each company, company-wide, as some metric to tell us how the profit from a particular product should be apportioned to account for this patent. It just makes no sense.

So let's start with the -- the mid number. So he looks at DDR4 as an example, and he says the gross margin is -- is 55.7 percent. He says if you are looking perhaps at internal documents from Samsung, that that's the margin.

He then says, well, then that margin should belong to Samsung and, again for reasons he doesn't explain, he then seizes on the cost part of it and says, well, 43 percent of it is cost, and if 43 percent of it is cost, he multiplies what had been the hundred percent number by 43 percent and says that's apportionment.

Now, Your Honor, I can't follow that. I have no idea why he thinks that that's appropriate. It appears to us that he needed to find a percentage that he could use so he could say he cut down the damages and that's what he used. But it makes no sense, it's not tied to the patented invention, it's really not tied to anything. It was just a percentage that he could point to and that he could utilize, and that's exactly what he did in this case.

So his mid numbers are the result of him taking the cost

percentage, the cost of goods sold as a percentage of the product, in this case 43 percent, multiplying it by his hundred percent number, and saying, aha, I have apportioned.

The -- I already said some of this, but Samsung's costs have nothing to do with the accused features and certainly have nothing to do with what Netlist says they contributed. So the idea that they have somehow reduced power savings or reduced the capacitance in an HBM die stack, that has nothing to do with Samsung's cost of producing that.

The cost has to do with the fab that has millions of developments and patents baked into it and years of experience and technology that's -- that's beyond the pale from all but just a few companies in the whole world. Those are the costs that go into it. Again, no -- no tie whatsoever to these patents.

Mr. Kennedy says, well, I've got experience. But, Your Honor, that won't do it. We need to be -- we need to understand his methods. We need to understand his reliable and articulated methodology in order to evaluate it. And if you can't give us that, he's not supposed to take the stand.

So let me -- let me now jump to the technology, because this one really was hard for me to get my mind around. I had to read it over and over again.

So what he does is he says, let's look at R&D expenses as a share of each company's revenue. Okay? So we start out

with Netlist's revenue. Netlist net product sales in 2020 were $47.2 million. Their R&D expense was 3 million. Therefore, he calculates that they spent 6.3 percent of their revenue on R&D. Okay. That's pretty straightforward.

Samsung is much bigger, 206 -- I think that's in billions so that's $206 billion versus R&D expense of 18 billion leaves you a nine percent. And he averages that over 2020 and 2021 to come up with an average percentage of revenue committed to R&D. So what he's done is he's calculated essentially how much money each company is plowing back into research.

Okay. That's -- I can do that calculation. I can follow that.

Where I can't follow is that he takes those, he then creates a ratio between the two. So at slide 31, he creates a ratio and says that the R&D investment as a percentage of revenue between the two companies reveals that Netlist is devoting 44 percent on a percentage of revenue basis to R&D compared to Samsung.

THE COURT: Let me stop you, Mr. Cordell, and just simply ask, is this material that's adequate for the public record or does the courtroom need to be sealed?

MR. CORDELL: I think much of this is -- is public, but I would defer to Mr. Sheasby.

THE COURT: If you're not concerned, I'm not concerned. I just see a lot of numbers and percentages, and

that causes me to ask the question. But unless you're going to affirmatively ask to seal the courtroom, we'll proceed.

MR. CORDELL: I don't feel like I need to at this time.

THE COURT: Okay.

MR. CORDELL: I will tell you that we dispute a lot of what Mr. Kennedy came up here with, but much of it is pulled from public sources, his math notwithstanding.

THE COURT: Then please continue.

MR. CORDELL: So this is where we are. We have Samsung spends more of its revenue on research and development than Netlist does, although Netlist spends a relatively small amount. But then we're going to take this and make it a ratio. We're going to compare the two and then use that to apportion damages in a specific case.

So if Samsung spends -- doubles its R&D investment in washing machines, the damages in this case would go down by Mr. Kennedy's analysis, because Samsung would be devoting more of its revenue to research and development, regardless of whether it has any relation whatsoever to this case.

By the same token, the fact that Netlist in the last two years, long after all of these patents were filed for, has devoted $3 million a year as a -- as a share of its revenue, which depends primarily, I hear it, on licensing and -- and intellectual property enforcement, that metric is somehow used

to fix damages in this case.

So if they're more successful in suing people for patent infringement, the revenue will go up, their share -- the revenue share devoted to R&D will then go down. So let's say they get a $20 million windfall next year. Well, then that R&D share will fall from 6.8 percent to 2 or 3 percent, and that will then suppress damages in this case.

It makes no sense. They have no relationship to one another. It's, again, as if somebody said, Mr. Kennedy, you can't put these kinds of numbers in front of the Court, you can't go in with a hundred percent, find a way to apportion this, find a way to give us a percentage so that we can allocate damages, and that's what he did.

And I got to tell you, Your Honor, it's not the kind of thing you'll get from the first time you read Mr. Kennedy's report. You've got to read it over and over again. It was studied obfuscation, is the phrase that the boss that I had when I got hired by the Patent Office would use. It's as if he wrote this in a way to make it very, very difficult to actually ascertain his true analysis, for good reason.

I mean, these are absurd. You've got a hundred percent rule, which obviously can't pass muster; you've got this -- this fanciful idea that you're going to allocate based on -- on, you know, investment share in R&D; and then the -- the other -- the mid range method is an allocation based on the

cost of goods sold, which again has no relationship whatsoever to the patents in this case or these inventions.

So the bottom line is we -- we end up with a contribution that he calculates. I have in slides 32 and 33. But it has to come back to being tied to the contribution of these patents, and there's just no effort to do that whatsoever.

All right. So let me see if I can take you forward. Again, I already --

THE COURT: Let's go on to issue 3.

MR. CORDELL: Yes. Thank you.

So Mr. Kennedy's HBM4H, 4 high stack of -- of HBM memory he says is a non-infringing alternative. However, we know from Mr. -- Doctor Brogioli's report that, in fact, it is an accused product. So we have their technical expert saying that this 4 high stack is an accused product, and yet Mr. Kennedy uses it as his non-infringing alternative. So there's something profoundly wrong here. There's a conflict between the two.

Again, Mr. Kennedy has offered us nothing about whether these supposed non-infringing alternatives, whether the Corolla is in fact a viable or acceptable alternative to the Lexus. And there should be such evidence. If he's going to advance a non-infringing alternative, he's got to be able to do it.

He pointed to the technical experts, but they punted.

Doctor Mangione-Smith said, I don't know how to test a hypothetical system. I can't do that analysis. I don't know about the non-infringing alternatives.

So Mr. Kennedy can't hide behind the technical experts.

Critically, they not only have to exist, they have to be acceptable. You have to be able to prove that people will actually buy them and there will be some -- some commercial success that -- that will follow. And we've got none of that in this case, and Mr. Kennedy admits that he's got none.

Mr. Kennedy just assumes, I guess because people are -- you know, Samsung is a successful company, that no matter what they make, people will buy it. So he can posit these hypothetical products, he can come up with these non-infringing alternatives, and that Samsung will -- will be successful. But he needs some evidence of that, and there's just none in his report.

Again, I won't belabor this, but we have -- we have in his -- in his report, he -- he makes it very clear that he's being conservative when he says that Samsung will be able to -- to find buyers for these hypothetical devices, but he's got to be able to prove that. He's got to have some evidence of this.

Critically, Your Honor, and I've showed you this, they didn't disclose any of this in discovery. So the first time that we learned of any of this analysis is when Mr. Kennedy

served his report.

I had the *MLC/Micron* case before Judge Illston a couple of years ago in San Francisco, and she struck the whole damages case exactly for this -- this reason, that you can't wait until your expert report to launch these kinds of -- of critical analyses at first blush and expect to tell the jury about them.

These are untested, they're unproven. The factual predicates that Mr. Kennedy relies on are missing. We have no opportunity to test them, and they should have at least disclosed them earlier.

Let me -- let me move ahead. So Doctor Groehn's flawed regression analysis. We have a separate motion attacking Doctor Groehn's analysis. I won't belabor that. But one thing that we can all agree on is that Doctor Groehn was focused on DDR4. Okay? It's a different family of memory. He did an analysis -- he did a regression analysis that we don't -- we don't think is correct.

There's no evidence to support Mr. Kennedy taking that DDR4 regression and applying it to DDR5. It's a very simple formula. He just did it. And everybody we asked said they had no -- they had no ability to -- to know whether there would be a relationship between speed fallen price versus reduction in speed, which is Doctor Groehn did for DDR4 in the DDR5 context. Everybody knows they are different products,

they operate differently, the patents involved are completely different.

And so Mr. Kennedy took that regression and applied it directly to DDR4 as if it was made for DDR4 [sic]. And yet he admits this is -- this was a -- this was a -- he switched horses midstream.

Non-comparable agreements. There are a couple of agreements that could have been the basis of Mr. Kennedy's opinion but weren't. So we have the Netlist-Samsung JDLA that we heard a lot about today. That would be an agreement that perhaps some experts might have -- might have pointed to with the number in it, it involved a license to the very patents-in-suit that we're talking about. We also have the SK hynix-Netlist license, which our expert spends a lot of time focused on.

Mr. Kennedy looked at none of these. He said, in fact, that they were not comparable. He said they were -- they were not worth -- they were not worth looking at at all. What he does instead is he looks at a handful of Samsung licenses which he admits in paragraph 486 are not sufficiently comparable.

So the talisman for a damages expert to use an agreement to try to set damages in a case is that it would be comparable. That is the gating issue; that is the threshold issue. He admits none of these are, in fact, comparable, and

yet he tries to offer three--Rambus, KIP, and Nanoco.

And why? The question is pretty easy -- the answer is pretty easy: He wants to show the jury large-dollar licenses to induce them to award large dollars. The fact that Samsung has been subjected to prior high-dollar licenses in the past has no bearing on what the adequate number -- amount of damages are in this case.

The technologies were different. Your Honor is well familiar with the KIP case as is Mr. Sheasby. The Rambus case is a completely different product, completely different set of -- well, Rambus was obviously a different company at a completely different time. And then the Nanoco case is the one Your Honor would have presided over, but it settled right on the eve of trial, and it was about televisions.

So Mr. Kennedy says, well, these are not comparable. He doesn't dispute that. And yet he wants to show them to the jury because he wants to inflame passions or somehow get them on a -- on a path to awarding large damages.

Irrelevant opinions. They had Mr. Kennedy opine about a whole bunch of issues that he knows nothing about. So he talks about damages flowing from Samsung's alleged breach of the JDLA. So I warned you this morning that that was going to be a recurring theme that we heard from -- from Netlist, and we see it in Mr. Kennedy's report.

We are not here to set damages for breach of the JDLA.

They had a chance to do that in California. They came up with zero, and it's improved to a dollar. But that's where that belongs. It has no place here.

We know that Mr. Kennedy goes on and on about whether it was in good faith or bad faith for Samsung to have supplied parts under the JDLA. Well, Your Honor, that has -- that has no bearing on patent infringement. We're here to talk about patent infringement, not the supply agreement under the JDLA.

He talks about whether he believes this was ethical or unethical. He is making value judgments about the commercial relationship between Netlist and Samsung. Clearly prejudicial, clearly irrelevant.

He talks about notice and willful blindness. He's suddenly become a willfulness expert even though he confesses that he has no knowledge of these standards. So he's offering opinions that perhaps he's just mouthing the words that people are giving him because he doesn't have any expertise and doesn't even know the standard that he's opining about.

Let's talk about RAND opinions. He gives us five paragraphs of plain old this is the way it is and you're just going to have to take my opinion and my word for it on -- on Netlist's RAND behavior, that Netlist somehow does not unfairly discriminate and, therefore, his damages are RAND. He's done no work; he's done no comparison with other parties.

When Your Honor and I spent a week in December together,

we heard a lot of evidence about what a FRAND or RAND royalty rate might be and the kinds of things that go into that. Mr. Kennedy has done none of it. None of it. He talks about --

THE COURT: Are you aware of, Counsel, of any standard essential patents at issue in this case?

MR. CORDELL: I am, Your Honor. The -- the unwavering position, unwavering position of Netlist until about, I don't know, three weeks ago, maybe four now, was that every patent in this case was standard essential.

THE COURT: Well, what's Samsung's position?

MR. CORDELL: Samsung's position was that we -- we didn't think they were standard essential, but we knew that if it was a close call, they had to be submitted to JEDEC. That was the way the parties were supposed to behave. So we had taken that position all through the case that they should have disclosed these to JEDEC.

And I think that -- and, again, who knows why, but it could be that -- that perhaps Netlist felt like they were on the down side of that issue and so they've switched horses and said, no, these are no longer standard essential; we're taking all of those representations back.

THE COURT: One more time. What does Samsung's say?

MR. CORDELL: Samsung says they're not standard essential.

THE COURT: Okay.

MR. CORDELL: But it was a close call and they should be -- and the second part of that -- the second part of that phrase is -- is critical. It was a close call and they should have been disclosed to JEDEC because that would have allowed everybody to decide their essentiality.

THE COURT: But all that notwithstanding, as we sit here today, Samsung and Netlist both maintain the position that there are no patents at issue in this case that are standard essential. Correct?

MR. CORDELL: That is correct. But we're not letting them out of those representations. So they, again, have told this Court and us in discovery and the whole world that they were standards essential over and over and over again, and this jury needs to hear that. If, in fact, they're going to accuse us of infringement and they're going to deny the fact that they shouldn't be licensing these on fair, reasonable, and non-discriminatory terms, then the jury needs to understand their flip-flopping, if you want to call it that.

THE COURT: Well, we'll get to that when we get to that. You're arguing to me about Mr. Kennedy's opinion with regard to RAND. I'm trying to determine if there's any relevance to RAND opinions in a case where there are no standard essential patents at issue.

MR. CORDELL: Well, I would say it slightly

differently. He shouldn't be allowed to testify as to RAND if -- if his client, if the party who has sponsored him, is taking the position that they're not standard essential.

THE COURT: All right. Let's continue.

MR. CORDELL: All right. The future damages analysis is pretty straightforward, Your Honor. He has no basis to look into the future. That is for the Court, that is not for Mr. Kennedy. Again, they're just trying to grab at large numbers to perhaps impress the investing public. I'm not sure. But that -- that seems like it's completely out of place. We have -- we have the statute that tells us that that's improper.

And it may be that they're offering a lump sum, that they're going to offer the jury a lump sum option the way some patent Defendants say, well, we typically negotiate for a lump sum and we'll have it at that.

But Mr. Kennedy didn't do that. He didn't -- he didn't look at what Netlist normally accepts as the lump sum or what Samsung normally pays as a lump sum. He said, I'm just going to calculate running royalties out into the expiration of the patents, and then I'm going to net present value them back. That's not a lump sum analysis, and he shouldn't be allowed to do it.

There is just no question that the products in this area evolve constantly. So memory today and memory 20 years ago

are not the same memory, and we need to acknowledge that. And he has no basis to -- to look out into the future.

Critically, Your Honor, we have instituted IPRs on four of the six, and we're awaiting two more that we'll probably get in the middle of the trial, get decisions on in the middle of the trial. If they are invalidated, none of Mr. Kennedy's future analysis could ever come to pass which -- which underscores how speculative it really is.

And then, finally, the theories that -- that were undisclosed in discovery, I've touched on some of this. The initial disclosures are supposed to give us their damages case. We're supposed to have it. So in -- in the discovery order in this case, we're supposed to have their damages. Rule 26(a)(1)(3) tells us they're supposed to tell us their damages case right at the outset. We got none of that.

We asked them in interrogatories and Interrogatory 11, and I've already walked Your Honor through that. And we simply got none of these non-infringing alternatives. All of the -- all of Mr. Kennedy's analysis was nowhere in their disclosures. We cite the *MLC* case because it's one of mine so I remember it, and it -- it was -- Judge Illston used it to great effect to limit a plaintiff that does exactly what Netlist did.

So with that, Your Honor, I'll sit down.

THE COURT: All right. Let me hear opposing

argument from the Plaintiff.

MR. SHEASBY: May it please the Court, Jason Sheasby for Netlist.

THE COURT: Please proceed, Mr. Sheasby.

MR. SHEASBY: Slide 58.

Every single technical basis for these damages was disclosed during fact discovery, that our patent achieves a 30 percent savings on speed and power for DDR5, that you can only use our patent -- you can only have one DIMM per channel if you don't use our patent. And, in fact, they offered a 30(b)(6) in which we put in even more detailed information, and we expressly said, for example, more than 8 high, you would not be able to use our technology.

So I want to start with something. The detailed description of what exactly the technical benefit was to these patents was expressly disclosed during fact discovery, and they examined our technical witnesses on the subject. Precision, even late in the day, is extremely important from both parties.

Let's go to Mr. Kennedy's analysis. If we could have slide 5.

So as to the '981 and '054 Patents [sic], what Mr. Mangione -- Doctor Mangione-Smith does is the patents yield a 30 percent increase in power efficiency, and what he analyzes is, in order to make up that power efficiency in another

state, they would have to throttle the design. He then concludes that it would be acceptable or it would be -- it's -- it's available to be able to throttle the power design, and based on that, he calculates how much speed would have to be decreased if the specific patent was not used.

Doctor Mangione-Smith did exactly what the Federal Circuit asks--what is the technical world when the technology isn't used, what is the technical world when the technology is used. And then the detailed calculation was performed by Doctor Mangione-Smith and Mr. Kennedy as to exactly what the impact of that speed throttling would be in order to obtain the speed challenge.

For the '506, once again, the -- you would not be able to sell speeds above 2400 MT per second without the technology. Products less than 2400 MT per second are available. They are used and they are sold and they are used as the basis for damages that he -- that Mr. Kennedy analyzes, which is what they would lose if they could only sell the degraded low speed products as opposed to the higher speed products.

As to the '339, same thing. There are one die per channel products. They are called RDIMMs. 2 die per channel, you need an LRDIMM. So Mr. Kennedy, after we disclosed that you can't use -- you can't have multiple dies -- DIMMs per channel without our technology, he calculates what the difference in damages would be if you were not able to sell

the 2 DPC configuration.

For the HBM products, to be clear, Samsung actually sells 4H HBM products. What Doctor Brogioli and Mr. Kennedy concluded is that if they use an alternative technology, the multi-drop, they could still sell 4 high. The way the damages analysis works is even though 4 high infringes the patents, there's no damages for selling 4 high because 4 high didn't need our technology to be successful. We only are seeking damages on higher than 4 because that's what is necessary to -- because our product is necessary to sell higher than 4.

So this fiction about these products have never been sold, they don't exist, these are made up NIAs, that is -- it is -- it is inaccurate in a way that is -- that doesn't serve anyone by being inaccurate. Less than 2400 MT/s is sold today. Less than 2 DRAM per channel is sold today. And 4 high RDIMM is sold today. 4 high HBM is sold today, and speeds less than 4800 MT per second is sold today.

So I want to be very clear: This idea that these are made up or absurd -- absurd apportionment analysis, that -- that doesn't have a connection to the record.

Now, what Mr. Kennedy did is he did the following: There is no evidence in this record of a single patent or a single technological innovation that contributed to the 30 percent increase in power speed, the 2400 MT/s, the 8H HBM, there is nothing else in the record, they didn't point to any other

alternative technologies, they don't point to anything that could achieve that without our technology.

And we asked them, we asked them, are there any patents, is there any other technology that went into this other than ours, and their corporate representatives couldn't point to anything.

And so the question is, what do you do with having apportioned down to the footprint solely from the benefits of the patents, not from anything else, and they don't point to anything else and they don't point to any alternative, what is the profit split.

So I want to be very precise. What Mr. Cordell said was inaccurate in a way that is important to this case. Mr. Kennedy did not give us a hundred percent of excess sales. Mr. Kennedy did not take the -- the sale price of a product or the revenue of a product.

What he did is he calculated what would be the differential revenue if they sold the D-featured product versus the infringing product. That is black letter Federal Circuit that we are entitled to that profit. And you don't have to ask me. You can ask their own expert because their own expert does the exact same thing that Mr. Kennedy did.

The histrionics that we heard for 20 or 30 minutes about how this is unprecedented, how you're going to embarrass yourself in front of the Federal Circuit, how this Court is

incompetent, how the counsel for Netlist is leading you down a judicial error because of -- because of the prospect of awarding a hundred percent solely of the technical value of the patents-in-suit in a hypothetical negotiation, their own expert does the exact same thing. Now, his technical footprint is dramatically less, but he does the exact same thing, exact same thing.

So everything we just heard about how unprecedented and shocking this is was lawyer rhetoric that has no connection with reality.

There's another aspect to why Mr. Kennedy didn't assume that we would take the technical benefit. They still make revenue, they still make profit on the product, but the differential technical benefit we obtain, and the reason why he pointed that out is that we are not a willing licensor. And I'll point to the *Monsanto versus Ralph* case. This is 382 F.3d 274. It's a very similar case.

There's a license agreement, the license agreement was breached, the Defendant used the technology, in this case C technology, for a reason -- I'm going to take this down because it's worth just talking about without showing the slide.

In the *Monsanto* case, the -- the defendant took a seed and used it for the purpose that was outside of the contract. They breached the contract. And the Court in *Monsanto* said

very clearly two things that are very important: One, the Court in *Monsanto* talked about the fact that, quote, the jury's award reasonably reflects that Ralph misappropriated a license to use the seed without restrictions, quote.

We agree with *Monsanto*.

Federal Circuit goes on to say, the law does not require that an infringement be permitted to make a profit and, worse here, a patentee is unwilling to grant an unlimited license. The hypothetical negotiation has its limits.

And what Mr. Kennedy speaks about is exactly we've reached those limits, that the reason why we take the technical benefit for our patents, they get the benefit for everything else, and they get the revenue and profits for everything else, is because the record shows that we have never in our history ever entered into a naked patent license agreement.

I want to speak to this, about what's in the record and emphasize this. HBM 4 high, which is a non-infringing alternative, is in the record. 2400 -- less than 2400 and 2 DPC is in the record, and speed less than 4800, those are all real products. These aren't made up, absurd products. These are real products.

I'll stop there and go on to the regression that Mr. Kennedy uses. He doesn't use the regression for every design. It's not relevant to the HBM, but it is relevant to a certain

issue. And they criticize the extrapolation of DDR4 to DDR5. Once again, this is an argument only a lawyer could love.

Their own expert under oath admitted that the market for DDR5 and the market for DDR4 are the same markets. They are the same drivers. Speed, density, power, and entering into the market are the four features that drive demand and price at DDR4 and DDR5. And those are the exact same things that Mr. -- Doctor Groehn analyzed.

I want to stop for a second and I want to speak sort of about something that's I think more prudential than it is legal, which is, I think the damages number at trial is I think $400 million through trial. And that's after the Court's rulings today.

If the damages was $3 billion, if the damages was $10, the Court should apply the same rigor to it. The sort of the -- the underlying subtext of you're going to put your thumb on the scale because these numbers are large, I don't think that has any place whatsoever in a *Daubert* hearing. The *Daubert* hearing is not about the size; it's about the rigor of the methodology.

Let's go back to slide -- let's go back to the next issue. I think I've spoken enough about the regression and its rigor.

A hundred percent of the technical footprint, we are entitled to. Federal Circuit says that. If you look at the

Prism Technologies case, we're entitled -- it says we're entitled to that as well.

He does do additional apportionment measures beyond that. Those apportionment measures are rigorous. For example, Mr. Cordell spent I think all of five seconds talking about one of his cost of goods apportionment measure. And so that would be a rigorous measure as well. But to be clear, the top number where we get the technical footprint is rigorous, consistent with *Monsanto and Prism* and is in the case.

That's the end of my discussion of the apportionment. I'll now go on to the other issues unless the Court has any question.

THE COURT: Please continue.

MR. SHEASBY: All right. So the next issue is the license in the hypothetical negotiation.

Let's go to slide 39.

So, once again, it's late in the day, but precision is very important for both myself and Mr. Cordell and the Court. This idea that the Rambus agreement was not identified as comparable is a fiction. Mr. Kennedy said, None of the other agreements are comparable except the Rambus agreement. He literally said that. So that what we just heard about how Mr. Kennedy never said any of these license agreements were relevant is the exact opposite of what he said in his report.

And as to Nanoco and KIP, those were delayed in their

production until after he filed his original report. So that statement that Mr. Cordell showed you that says nothing other than Rambus is relevant, could not have included KIP and Nanoco because they weren't disclosed at the time. Those are in -- those were in subsequent reports. So any suggestion that the statements Mr. Cordell was showing you were accurate or a fair representation of this report were incorrect.

What for Rambus, what he did -- and Rambus is actually probably the most similar agreement in the case to what happened here. Rambus had a license to license Samsung, Samsung's license was terminated, and then Rambus was sued for infringement. And Mr. Kennedy talks about these exact facts.

He then speaks about the fact that the Rambus patents are technically comparable; that products that were accused of infringement were the DDR memory products that are exactly the same class of products that are at issue in this case. And what Mr. Kennedy did is he went and asked the technical expert to look at the products -- look at the patents that are at issue in Rambus, and he said that they are comparable.

By the way, he's not the only one. They actually hired an expert, Mr. McAlexander. For reasons that are unclear to me, he spends 15 pages discussing how comparable the Rambus patents are to this case. I kid you not. It's in our briefing. They have their own expert saying these patents are technically comparable.

What Mr. Kennedy does with it is what he should do. They have this view that unless you extract a final number from a license agreement, it's not relevant. That's not the law. Under that law, unless you extract a final number from any piece of evidence, it would not be admissible.

What Mr. Kennedy does is he gives the jury the tools to understand how to deal with this. He talks about the fact that this is a portfolio agreement, that there is a skew, that you shouldn't reflect this is all being valuable -- as all being the value of the entire portfolio, nor should you view this as being the value of just the five or 10 patents or whatever it is at issue.

He does exactly what you should do, which is showing that this agreement is relevant, it would be used to suggest that the royalty rates in this case should be higher because the patents are comparable, and there's nothing -- there's no Federal Circuit case that says you go to a final number.

Mr. Cordell is right. If Mr. Kennedy would have come in and said X hundred million dollars is what they paid Rambus, give us that, that would be totally inappropriate. What he does is exactly what the Federal Circuit says he should do in *Finjan*, which is, he gives the jury the tools to understand how to use this agreement and how it cannot be used. And he reaches the conclusion that this would drive royalties higher.

As to the JDLA and the Korean Tax Tribunal decision, this

is in my mind incredibly important. The JDLA agreement is one of the agreements that Samsung has presented as a comparable agreement, and Doctor -- Mr. Kennedy talks about it as well.

What he does with it is very important. He doesn't say, oh, my God, we were damaged from breach and our breach -- the damages breach would be this -- this amount of money. That would be inappropriate and he doesn't do it. What he does is the following: He says what is the true economic value of this agreement when it was entered into. The true economic value of this agreement was a supply commitment, a supply commitment that was unconstrained.

And you know how Mr. Cordell spoke about the fact how -- I think it was his colleague, actually; I apologize -- spoke about the fact it was -- Ms. Degnan, she spoke about the fact that there was a supply constraint, and they just decided not the sell us product because it was so scarce.

Our deal was they had to sell it all the time to us. And what Mr. Kennedy does is, by looking at the revenue that we were generating based on this, he projects out from the time of the hypothetical negotiation what would be the value of that supply clause. No one's going to say there was a breach, no one's going to say there's damages for a breach, but Mr. Kennedy is allowed to speak about the value of the hypothetical negotiation.

As to the KTT agreement, the KTT agreement decision is --

is crucial and relevant. Their expert admits it would be considered in the hypothetical negotiation. And the reason why it's so relevant is because the KTT decision says that the 8 million non-recurring engineering fee that was part of the agreement was not a reasonable royalty.

My brothers will tell the jury the exact opposite, that it was a reasonable royalty even though the KTT decision was adopted by Judge Scarsi and even though they admit that the KTT decision was in the case.

The next issue I want to talk about this is Mr. Kennedy didn't say that Samsung's conduct was unethical. Samsung's 30(b)(6) witness on the topic of this JDLA said the conduct was unethical. And the reason why -- the reason why that is relevant is the same reason in *Monsanto versus Ralph*, that the idea of Defendant acting ultra vires outside of the contract, totally contrary to their permission, is a factor that the jury could consider.

The idea that he should not be able to talk about JEDEC is also totally inappropriate. They have this claim that JEDEC is somehow a surrogate for value even though they have no evidence that any of the patents-at-issue were disclosed in any JEDEC contribution, but they're just saying JEDEC, JEDEC, JEDEC.

And Mr. Kennedy is telling a very accurate point, which is that JEDEC is not a tool for innovation; JEDEC is a tool

that is used to extract intellectual property from small patent holders. That is a reasonable economic view of what SSOs are.

Now, candidly, SEPs are not in the case. None of that is in the case. But I do think it's important to make clear that he's not just making scurrilous objections as to JEDEC.

As to his RAND analysis, his RAND analysis is rigorous. He specifically talks about why it's not unfair discrimination and indicates that it's not unfair discrimination because they're only paying for the use of the technology. And by the way, that's another reason why they want to keep out the calculation of the value of the supply agreement because they want to hide from the jury what the true value of the Samsung JDLA was to us.

THE COURT: Let me stop you, Mr. Sheasby. Why is there any probative value to the RAND issues if none of the asserted patents are standard essential?

MR. SHEASBY: There is none.

THE COURT: All right. And it is late in the day as you've said several times. So if you could slow down a little bit --

MR. SHEASBY: Yes, Your Honor.

THE COURT: -- I'd appreciate it.

MR. SHEASBY: I understand, Your Honor.

THE COURT: All right. Please continue. We've go

one or two more points on your list.

MR. SHEASBY: One more point. I apologize for going too fast.

The last issue is, are we allowed to accept the -- request a lump sum? Of course we are allowed to request a lump sum. I hop up and down when -- when defendants request a lump sum and I say it shouldn't be allowed and this Court allows it consistently because the Federal Circuit says it's appropriate. So the idea that there's something ultra vires about lump sum is not the case. It is not the law.

And I want to keep coming back to this. It is okay to say that we're horrible people, it's okay to say that we do bad things, but to suggest that we did not disclose the technical contribution of these patents in discovery, that has no connection to reality.

THE COURT: Well, it's not okay to say we're bad people and we do bad things and to cast aspersions toward the other side. And, quite honestly, in the argument on this motion, there's been too much of that --

MR. SHEASBY: I understand that.

THE COURT: -- in both directions.

MR. SHEASBY: I understand that.

THE COURT: And I don't want to hear any more of that.

MR. SHEASBY: Thank you, Your Honor.

So this is the last point. These have been fully disclosed, these damages points.

THE COURT: All right. Thank you.

MR. SHEASBY: Thank you.

THE COURT: Mr. Cordell, do you have a targeted and hopefully brief follow-up?

MR. CORDELL: Very targeted, Your Honor. So we can -- we can begin with the last point Mr. Sheasby made about the discovery.

Can I have slide -- I think it is slide 6.

So there we did ask them to give us their disclosed damages case. Tell us what your methods are going to be, tell us what you're going to point to.

And if I can go to slide 7.

They make they do utter the words 30 percent power savings. And if this was a case about 30 percent power savings, then I supposed the experts could argue about what that's worth. But that's not what Mr. Kennedy did. Instead, he did this comparison to a non-existent, non-infringing alternative.

I didn't hear Mr. Sheasby say anything about the product being in existence. He said the technologies could be there, and they could imagine these things.

But recall Doctor Mangione-Smith saying, I don't know how to test a hypothetical product, I don't know how to do that.

That was his technical expert. So that -- that is a problem. And then when we go to slides 8 and 9, we saw the paucity of disclosure that we've got on these things.

So for the HBM patents, it was 3 D stacked array dies with -- enabling higher capacity memory products. And, again, that may be worth something, but that isn't the focus.

Mr. Sheasby tells us, well, the 4 high stack, you don't have to use their patented invention. It still has it in there. So it's not really a non-infringing alternative. Whether you must use it or it's optional, the law of infringement doesn't consider those factors. They only ask whether the elements are there.

So he tells us that the 4 high stack has the elements, doesn't use the -- let's see how he phrased it. Doesn't use the multi-drop, but it could. Well, it uses the multi -- could use the multi-drop, but it doesn't. So he says the 4 high stack still does violate the patent, infringe the patent, but since it doesn't have to, it's a viable non-infringing alternative. Well, that simply doesn't make any sense.

Fundamentally, what I think we should focus on is the following: So Mr. Sheasby tells us it's perfectly acceptable for this Court to take a hundred percent of the value of the patented invention and give it to Netlist. 100 percent. I heard him say that; I wrote it down. He said that's -- he

showed an excerpt from our -- our expert who talks about apportioning the value back and forth.

We need to go back to the *VirnetX* case because in the *VirnetX* case what we were talking about there was the 50 percent rule. It was the Nash bargaining solution. And the 50 percent rule had to do with apportioning the value of the patented invention. It wasn't some other metric; it wasn't some other lode star. It was the value of the patented invention.

And what the Federal Circuit said in *VirnetX* is that the 50 percent rule is nonsense. Nonsense. That would be irrational. It would be irrational for someone negotiating a license to say, I'm going to give you half of whatever profits I receive from this patented invention, because the risks are too high. No one would ever do that. The same thing throughout the 25 percent rule a couple of years before in *Uniloc*.

But the reality is the law says you can't take all the value and give it to the patentee because then the manufacturer in this case has no incentive to do anything. They would never take that risk. In fact, an economist would say that the risk is higher than a hundred percent, that you're sure to take a loss because you can't account for all the bad things that might happen along the way.

So it is an absolute, rock-solid legal principle that

we're talking about here. We cannot give a hundred percent of what Netlist says is the value of the patented invention to the patentee because it won't pass legal muster. It simply will not.

With respect to the -- Doctor Groehn's analysis, I didn't hear a response or an explanation as to why it is appropriate to use a DDR4 regression analysis and apply it to DDR5. There's no linkage there. There's no technical linkage. There's no legal linkage. And it's simply inappropriate.

With respect to the JDLA damages, the scarce supply, all of this, that is a complete irrelevant distraction to this trial. And if we go into that, we're going to have the trial within the trial.

We have made a suggestion in the papers that there is a way for Your Honor to characterize the prior case if we need to, that there was a contract, it was breached -- well, it wasn't breached. There was a contract, there was a dispute about it, and it was terminated. So we can take it at that level.

But we aren't -- we shouldn't be going into why -- who shot John and why John was shot or whether John was shot. That -- that forms no part of the patent infringement issues that this jury is going to be asked to evaluate.

And then with respect to the ultra vires opinions of Mr. Kennedy, he's just not qualified to offer business advice

about what is ethical or unethical or what is prudent or imprudent or what JEDEC would do or wouldn't do. We have experts for that. We have JEDEC experts in this case who will talk to that matter. But Mr. Kennedy just doesn't have the qualifications.

And then, finally, with respect to the lump sum, the reality is there is no precedent for what Mr. Kennedy did. He extrapolated out and took the Court's role for future damages and then net presented valued it back to a -- a number today. It wasn't a lump sum analysis at all. Whether that would be appropriate coming from a plaintiff or not, we can talk about. But the fact is that's not what Mr. Kennedy did.

Fundamentally, Your Honor, he is setting us up for error, and I certainly never would suggest that this Court was incapable of handling these things. So to the extent I did, I apologize. The reality is we are all bound to give this jury something legally sound to consider, and offering them the tools to do these kinds of complex analyses only work if the tools are legally sound. And that's what we're trying to do.

THE COURT: All right. Thank you.

Mr. Sheasby, I'd like you to go back to the podium, and I want to hear one more time about Mr. Kennedy allocating a hundred percent of the incremental value to the patented technology and why you believe that should survive a *Daubert* challenge such as this.

MR. SHEASBY: Sure, Your Honor.

THE COURT: I mean, is it simply his opinion that a hundred percent of the incremental valuable is attributable to the patented technology?

MR. SHEASBY: It's the -- it's the --

THE COURT: What's the rationale here?

MR. SHEASBY: Yes.

THE COURT: Because I think you will admit that -- I can't remember another time when where I've seen a hundred percent of incremental value attributable to one side or the other. So I'd like you to focus on that and give me your very targeted argument, please.

MR. SHEASBY: Sure. So if you look at *Ericsson*, if you look at *Apple*, and you look at *Powell*, the ultimate reasonable royalty awarded must be based on the incremental value of the patented invention as the end product. So the only value that he's seeking is that he's granting the incremental value from the infringement of product. All the other value that they're generating, all the other revenue and profit that they're generating, they get a hundred percent of that. All we get is the value of the incremental benefit.

The *Apple versus Motorola* case says that estimating the value of a benefit provided by infringing features by comparing the accused products, not non-infringing alternatives, is the appropriate way of measuring damages.

And *Powell versus Home Depot* talks about the fact that the value of the benefit conferred to the infringer by the use of the patented technology is relevant to determine a use-based royalty.

So what you see from all these cases is all these cases -- in the case of *Powell*, it was a hundred percent of the cost savings that was given. In the case of *Prism Technologies*, it was a hundred percent of the cost savings that was given.

And so what Mr. Kennedy does is he has three levels of apportionment. He's got the technical apportionment, then the apportionments based on two other factors that could potentially be used. But our point is that the highest value is appropriate in this case--if you go to slide 29--because it's literally the exact same methodology that Mr. Meyer uses. They both agree that Netlist takes the technical contribution value. Mr. Meyer does the exact same thing that we do.

And so when I balance everything -- let me give you an example. The product -- let's go to slide 8.

So the product is sold for a certain amount, and they make a certain amount of revenue on it. They make less amount if they were to D-feature it. That's so that delta is the technical contribution between the -- the patented technology.

Mr. Kennedy then allows for two other additional apportionment methodologies which I assume Your Honor is not

asking about. Your Honor is just asking about the hundred percent.

And the answer is is that if you read *Monsanto*, if you read *Prism*, if you read *Powell*, we are entitled to the benefit when we're an unwilling licensor in this technical situation to the value of our technology. And I think if you go from *Prism*, the cases that make that point are *Prism*, *Ericsson*, *Powell*, and *Apple*, and *Monsanto*.

THE COURT: All right. Thank you, Counsel.

MR. CORDELL: Do you need a response, Your Honor?

THE COURT: I think not. Thank you.

All right. With regard to Document 205 and Samsung's *Daubert* motion and motion to strike regarding David Kennedy and his opinions, there are references in this lengthy report to the TI matter, Texas Instruments. The Court's previously excluded any reference to other litigation and to that particular matter, and it should be excluded from Mr. Kennedy's report and his opinions.

I'm going to grant the motion as to references to the TI matter.

With regard to the portions of Mr. Kennedy's opinion that relate to his RAND opinions, both parties have confirmed on the record that none of the patents-in-suit are standard essential patents that would be subject to a RAND contractual commitment. Therefore, their probative value, if any, is

greatly exceeded by the risk of confusion and prejudice.

Therefore, I'm going to grant the motion and strike the portions of Mr. Kennedy's report that relate to his RAND opinions.

The balance of the motion, I'm going to deny. And that will be the ruling on Document 205.

All right. Let's go next to Document 207, which is Netlist's motion to strike certain opinions of Defendants' expert, John B. Halbert.

Let me hear from Netlist on this.

MR. SHEASBY: Yes, Your Honor. Give me one moment. Your Honor, may I approach?

THE COURT: You may.

MR. SHEASBY: May it please the Court.

THE COURT: Please proceed.

MR. SHEASBY: One of the issues that is -- I think Samsung is attempting to inject into this case is sort of a suggestion in front of the jury that we acted improperly with JEDEC. And obviously since standard essentiality is not in this case, we don't believe that any sort of allegations of impropriety in front of JEDEC have a place in front of this jury.

I should note that that's never been pled, there's never been a pleading that JEDEC is -- is something that we had obligations to JEDEC that were breached from a legal

standpoint. There is equitable claims, but there's no legal claims regarding this issue. And, in fact, the only damages relationship that they talk about the JEDEC standards are -- is that we should not be seeking lost profits under them.

And so I would submit that anything about the speculative nature of Mr. -- of our behavior in front of JEDEC -- and let me just give you an example of this. Netlist was motivated by its patent litigation strategies and opportunistic behavior, that we violated in some sense the good faith of JEDEC, that we were trying to subvert the JEDEC process, that we acted dishonorably, that we acted opportunistically. None of that is appropriate subject for the jury because there's not a defense based on -- on JEDEC disclosure in front of the jury, and so that should be stricken.

I'll move on to the second issue. The second issue is Mr. Halbert speaks at length about some ITC proceedings that involve other patents in the same families. Those should all be stricken as well. There's no place for -- there's no place for discussion of extraneous ITC proceedings involving other patents in front of the jury.

And then the last point I'll speak about, and this is an example of the -- the lack of connection between the ITC proceeding they talk about in this case. I think there's some suggestion that -- that they want to indicate that we've lost a previous proceeding. I would like to do that as well if I'm

a defendant. I just don't think it's relevant in front of this jury.

The last issue is this idea of, quote, JEDEC contributions versus Netlist and others are fraud. And here's what I mean by this. So there's these DDR4, DDR5, standards and they're massive. They have tens of thousands of pages.

And what Mr. Halbert does is he calculates for the standards as a whole what percentage of them are contributed by Netlist, what percentages are contributed by Samsung, what percentage of them are contributed by third parties.

And he says something very important, and I'll quote it right here: "I have not parsed which contributions relate specifically to each individual standard." That is exactly, exactly, the point. It's no different from saying, here is 10,000 patents, how could Netlist product be any good.

These patents aren't standard essential. Discussing JEDEC contributions has tangential relevance to begin with. And this methodology of totaling up, quote, contributions without connecting them to the patents or -- not only not to the patents, but to the features accused of infringement in this case, is not a rigorous methodology.

I'll stand on that, Your Honor unless you have any questions. Thank you, Your Honor.

THE COURT: All right. Let me hear Samsung's response.

MR. MOSTELLER: May I approach?

THE COURT: You may. All right, Mr. Mosteller. You may proceed.

MR. MOSTELLER: Thank you, Your Honor. May it please the Court.

Your Honor, you heard a lot about JEDEC today --

THE COURT: Yes, I have.

MR. MOSTELLER: -- and for good reason. So all the products in this case, the DDR4, the DDR5, the HBM products, there's no dispute that those products all practice the JEDEC standards. Your Honor is aware, you're very familiar with JEDEC, you understand that it's a giant effort. There is hundreds of companies involved. As counsel on the other side just mentioned, some of the JEDEC standards are tens of thousands of pages. There is a lot that goes into these products.

Now, Your Honor, the issues that Mr. Halbert opines on and that Netlist takes issue with relate not to the standard essentiality of the patents that are at issue, but rather they relate to other issues, both that are equitable and those that are before the jury, including damages, that are still live in this case.

And so, Your Honor, if I can, Mr. Burton, go straight to slide 13, please.

Trying to move this along for Your Honor.

So there's really three categories of opinions that Netlist is seeking to strike. The first one relates to the JEDEC practices. And, Your Honor, these are relevant to the equitable issues that we've discussed briefly as well as the issues of willfulness and, in the same vein, indirect infringement that we discussed quite at length this morning.

THE COURT: Let me stop you there, Counsel.

MR. MOSTELLER: Sure.

THE COURT: Where there are equitable defenses that have been raised, which there are some in this case, those are for the Court, not for the jury. But in several earlier trials where equitable issues were raised, there has been evidence that both relate to those equitable defenses but also relate to some of the issues that are live before the jury in the trial on the merits.

MR. MOSTELLER: Correct.

THE COURT: Where there's evidence that applies to both the equitable defenses and the issues the jury are asked to decide, I have routinely allowed that evidence to be produced and put on during the trial before the jury, and when we get to the bench trial that follows the verdict, there's no need to repeat that because I've heard that evidence.

However, if the evidence only relates to the equitable issues and has no nexus with the issues the jury's asked to decide, then I have routinely instructed counsel not to

present it during the jury trial based on a risk of confusion since it doesn't relate to what the jury's going to be asked to decide, but to hold it and then present it for the first time at the bench trial that follows.

That's the approach I want to follow in this case.

MR. MOSTELLER: And, Your Honor, that's what we intend to do as well.

THE COURT: Okay. I just want to make sure we're all on the same page.

Now, with regard to these paragraphs 49, 55, and 57, Mr. Halbert is not going to disparage Netlist any more than I'm going to let any of their witnesses disparage Samsung. This is not a case about who can throw the most mud at the other side.

So language like subverted, opportunistic behavior, dishonorably, misappropriated, those are pejorative terms in the Court's view. And while I'm not going to strike them per se from the report, I'm going to apply the Court's limine orders which clearly in and of themselves prevent pejorative language and assertions made in the presence of the jury.

If there's anything this man's going to say, whether it's in his report or not, that could be considered pejorative, you need to come to the bench and you need to tell me what you intend to elicit from him before you do it and get leave of the Court. And he needs to be instructed before he goes on

the witness stand that he's not to involuntarily offer or voluntarily, I guess I should say, offer pejorative language in his testimony. If he does, he's going to be in hot water with me.

So I can go through and pick out particular words in this report and I can say this adjective is out and this one is out, but rather than do that at a high level, I'm just telling you I don't expect him to testify using pejorative terms and I don't expect the other side to launch pejorative language at him on cross-examination or at your client any more than I'm going to let him do that to the Plaintiff in the case.

Are we clear on that?

MR. MOSTELLER: Very clear, Your Honor.

THE COURT: Okay. Well, I think with that understanding, the rest of paragraphs 49, 55, and 57 can survive.

MR. MOSTELLER: Thank you, Your Honor.

If we can jump ahead then to the next section of Mr. Halbert's report at issue, which is paragraph 48. And if we can just go to slide 21, I think we can cover one paragraph in just one slide, Your Honor.

Here, as counsel mentioned, Mr. Halbert has some discussion of the prior ITC investigations. And to be clear, we understand that the Court does not permit opinions about prior investigations.

THE COURT: I don't permit testimony about prior investigations or prior litigation, and a proceeding before the ITC is litigation in the Court's view --

MR. MOSTELLER: Correct.

THE COURT: -- without leave of Court. Now, if you think the door's been opened or if you think there's a basis in real time as the case unfolds before the jury to go into something that I've told you as a general rule I don't permit, you can always approach and ask leave. And if I give you leave, you have a clear path in front of you.

But without that, you're not and he's not to talk about ITC proceedings, other litigation in Article III courts, other administrative processes such as the Patent Trial and Appeal Board. None of those things are permitted under the Court's limine practice without leave.

MR. MOSTELLER: Understood, Your Honor. And we have no need to have Mr. Halbert provide testimony or opinions about those prior investigations.

The purpose of the testimony that's given there in paragraph 48 really relates to Samsung's willfulness defenses and its defenses to indirect infringement as to the knowledge and the intent prongs of Netlist's claims, specifically that there are, amongst other obligations, that Netlist as a member of JEDEC has to the organization is to disclose patents that it is aware of, that it knows are relevant to particular

standards that have been developed by JEDEC.

THE COURT: I'm not going to try and block your attempt the defend against the willfulness allegations that are live in the case, but there are ways you can defend against those without expressly explaining, going into, and creating innumerable rabbit trails about other lawsuits and other litigations and other proceedings before other tribunals.

MR. MOSTELLER: Understood, Your Honor.

THE COURT: So that's my guidance to you on that.

MR. MOSTELLER: Thank you, Your Honor.

I believe the last issue that Mr. Sheasby raised was in regards to the rebuttal analysis that Doctor Halbert or Mr. Halbert provided in his rebuttal report related to the contributions to the various standards. And, again, as I just mentioned, Your Honor, the products at issue in this case practice the JEDEC standards.

And so the technologies and any apportionment that should be given to any contributions by the patents in this case need to take into account the JEDEC standards, the contributions of companies, including Netlist, and including everyone else, to those standards.

What Mr. Halbert did, as -- as counsel discussed, is he went through, he performed an analysis of who contributed what to the relevant standards. We have DDR4 products, DDR5

products, HBM products. And what Mr. Halbert did is he went and looked at the contributions in JEDEC to the development of the standards, both by Netlist and by others, to see effectively how many contributions went into those standards, how many of those contributions are attributable to Netlist.

There are numerous opinions in Mr. Halbert's report in terms of what he did and did not do in arriving at that analysis, and if counsel would like to cross examine Mr. Halbert on his analysis and how he conducted that, they're free to do so.

But Your Honor, the analysis that's in Mr. Halbert's report is directly responsive, both to Mr. Kennedy's analysis as well as to Mr. Meyer's analysis in which he is trying to apportion to any value to give the patents for the products that are accused.

THE COURT: Well, I have no problem with him responding to Mr. Kennedy's calculations and his responses can certainly be tested as to their reasonableness when he's on cross. So I don't see any compelling reason to strike those portions of his report.

I do want to go back and clarify one additional thing for you, though. I understand this witness is going to talk about his personal experience as a member of JEDEC, and I think as long as he stays within what he knows of his own personal knowledge is fine.

I do not expect and I will not permit him to become the voice of JEDEC and speculate about what JEDEC would do under these circumstances and what they wouldn't do. He's not -- he is not JEDEC and he's not a representative of JEDEC entitled as an agent to take a position for them.

He can talk about what he's seen as a member, he can talk about what he understands will happen in circumstances that he's observed, but he needs to cabin his testimony about the JEDEC experience with his personal experience and his personal knowledge.

I don't think it's a proper use of him to be some kind of CEO of JEDEC that's going to come in here and speak for the body, if you will, on any matter and every matter and offer opinions about what they would or wouldn't do in hypothetical scenarios. That goes too far, and I'm expecting you not to take him there on direct. And if you do, I'm expecting Plaintiff to object. And if any of that happens, I'm likely to just exactly do what I'm telling you right now.

MR. MOSTELLER: Understood, Your Honor.

One thing I would ask to clarify is that Netlist also has their own JEDEC expert in this case, Mr. Gillingham.

THE COURT: What's good for the goose is good for the gander.

MR. MOSTELLER: I appreciate that, Your Honor. I will stand down.

THE COURT: We will get to that witness when we get to that witness.

MR. MOSTELLER: Sounds good, sir.

THE COURT: All right. Thank you, sir.

Let's turn to -- by the way, for the record those are the rulings on Document 207. While the Court has not expressly struck any of the portions of Mr. Halbert's report, the Court has given clear and direct guidance to counsel as to how that report and the information and opinions of that expert are to be presented to the jury, and I expect that guidance to be followed.

Okay. So in a technical sense, the motion by the Plaintiff is denied there with the guidance and clarifications that I've clearly put into the record.

All right. Let's turn to Document 214. This is Plaintiff's motion to strike portions of the rebuttal expert report of Paul K. Meyer. This is one of those situations, Counsel, that I've referred to earlier in the day because I count 11 subparts to this one motion, and that's part of why we have such a tremendous load of material to get through during the two days I've allocated to pretrial.

But without further delay, let's get into this motion. Let me hear from Plaintiff on it.

MR. SHEASBY: May it please the Court. Jason Sheasby.

THE COURT: Please proceed, Mr. Sheasby.

MR. SHEASBY: The first issue is really a legal issue. So there's two -- there's two facts to what Mr. Meyer did. So he had Mr. McAlexander, quote, survey Netlist's portfolio and to assign values to each of the patents in those portfolios. And then he used that to divide by the value of the SK hynix agreement.

And the only legal point -- so I'm going to be able to cross-examine the heck out of Mr. McAlexander and Mr. Meyer about the fact that he actually didn't do this survey, and that's for cross-examination. That will be what it is.

The sort of the legal issue that the Court has to decide is the following.

If you go to slide 3.

What he did was he took the value of the SK hynix non-supply portion of the agreement, so just the license agreement, just the payment under the license agreement, and he said, Okay, I'm going to have Mr. McAlexander assign what Mr. McAlexander believes are the values in nine buckets relating to those patents.

The problem with that is under *Apple versus Wi-LAN,* the value to -- the value that you need to consider in the license agreement is what value did SK hynix put on each of those particular patents. And so what I don't know is, in this situation, they make no pretense about assessing which patents

in the portfolio was valuable to SK hynix, why did SK hynix make the payment they did.

So, for example, there were a number of patents asserted against SK hynix none of those passenger the patents in this case the patents in this case were never presented to SK hynix as then potentially infringing. And what Mr. Meyer does is he just sort of does this thing where he uses the agreement as a bucket and then arbitrarily assigns value within the bucket. And that, I don't think, is consistent with *Apple versus WiLAN*, but I don't need to belabor that point. That's just a legal issue that the Court needs to determine.

The second issue, and this is just a small point, is Mr. Meyer is missing a bridge from Mr. McAlexander in the sense that Mr. McAlexander assigns value to patents within a given family, but he doesn't assign value to individual patents within that family.

And so Mr. Meyer just arbitrarily assigns a certain amount of money within a patent family to one of the patents-in-suit that are in that family. And, of course, that is a sort of an arbitrary step as well. I can deal with that on cross-examination.

The more trenchant point for Your Honor to conclude is whether this is a violation of *Apple versus Wi-LAN* by the fact that he doesn't consider what the value of the patents were to SK hynix as opposed to Mr. McAlexander just arbitrarily

valuing these families. And Mr. McAlexander doesn't say these are the values for SK hynix. He just says, These are how I value the patent families. And that's a legal question that Your Honor will have to answer.

The next issue is RAND. I think Your Honor has already ruled on this, which is neither party is contending that's patents are standard essential. There has never been an allegation that we discriminated against them in the interrogatory responses based on a RAND obligation. And I think based on the Court's previous ruling that would apply here as well, that his discussions of us discriminating against them or us acting improperly under a RAND obligation has no appropriateness -- appropriate role in this court.

This is the third issue is Mr. Meyer identifies a number of non-infringing alternatives, what he suggests are non-infringing alternatives from the patents. The challenge with that is that Samsung never identified any non-infringing alternatives in discovery at all, zero. And their experts deny having ever done a non-infringing alternative analysis.

So Mr. Robins was asked and his answer is very prolix, which is the subject of another motion that -- that is not fully briefed yet, but he's never analyzed whether there are any other patents -- he's never asked whether the alternatives that he identifies would be non-infringing or would be commercially available, and neither does Mr. McAlexander.

So we're in this kooky world, and I guess I can deal with this on cross-examination, but time is pretty limited, where Mr. Meyer says these people told me these are non-infringing alternatives.

THE COURT: I gave you more time than you asked for to try the case.

MR. SHEASBY: Yeah.

THE COURT: You asked for 12 hours, and I gave each side 12 and a half. So don't tell me time's limited. If it's limited, maybe the case is too big.

MR. SHEASBY: Well, I think the case is just right in size, Your Honor.

Anyways, I don't want to pretend -- I don't want to give you a sob story. I can sort of cross them aggressively on this. But it is a point whether should you be able to put on an expert to say, oh, there's all these non-infringing alternatives where the technical people admit they didn't do that analysis and where the interrogatory response is -- is not disclosing that.

I did want to go on to the next issue which is references to other litigations and IPRs. I don't believe that has any relevance in this case, and I think it should be excluded. And I won't belabor that point any more. I think that's consistent with the other rulings. They talk about a number of litigations that -- that don't relate to any of the patents

in this case, and they shouldn't be discussed.

They talk about *inter partes* reviews that the patents have been initiated on. He has a document where he lists the number of times patents have been asserted in litigation. Once again, that would violate the Court's MIL, and I don't think has any appropriate role in front of the jury.

JEDEC contributions. This is -- we have no problem with Mr. Halbert getting up and saying, oh, there's JEDEC contributions, there's a ton of contributions. What we do have a problem with is that Mr. Meyer puts a type of rigor to it in which he uses the number of generic contributions made across these standards as a metric to assess the value of these individualized patents. And there's no nexus there because Mr. Halbert did not create a nexus between the specific areas of technology and the patents.

So Mr. Halbert can say all he wants about JEDEC. But to add a level of precision where the jury is told, you should think about the 10,000 contributions that are made to JEDEC for these entire standards when you value these patents, that adds a level of precision that is deceptive because it's not tied to either the technology at issue nor is it tied to the exact patent.

So I want to draw a distinction. What Mr. Halbert can say, he can say. For Mr. Meyer to use that as a damages tool when there's no connection between the patents or accused

products, that's a bridge too far.

The references to Micron, Intel, and SK hynix investments, I -- I don't really know how to -- how to address this. In a hypothetical negotiation, Mr. Meyer is presuming that the parties would take into account the amount of money invested by two separate third parties for which there's no evidence that they contributed any of the technology that's specifically related to these cases.

And that strikes me as -- as a really dangerous precedent to set in the hypothetical negotiation for you to start speaking about third-party products.

The next issue is confidential settlement discussions. If we write a letter to Samsung that's a settlement offer, that settlement offer has no place in front of a jury. That's the whole purpose of FRE 408 and things fall down. Samsung in their surreply brief suggests there's a well-known exception to JEDEC-RAND issues. There's no RAND issues in this case, and there was no exception in our letter.

So I don't believe it's appropriate for them to inject settlement communications into this case, and it would I think significantly damage the -- sort of the integrity of the process if we had to be aware of that.

References to the 2017 ITC action, there is an analysis that was done about a FRAND offer that we made in the *SK Hynix* case for a separate set of patents that related to LRDIMM.

FRAND is not from front of the jury because these patents are not standard essential, and because of that, I don't think this should be part of the case as well, this FRAND analysis. The jury's not going to understand why a FRAND analysis even matters because they're not going to be instructed on FRAND.

The next issue is -- is sort of a new one for me which is they're saying that our damages can't be as high as we're requesting because Samsung could just buy the company based on our market capitalization. So, in effect, that would be a tool that could always be used by a large company to sort of get out of damages from a company with a small market capitalization.

I know of no case law that has ever held that -- that the size of your market cap has any relevance to the damages that you pay. And, candidly, I think there would be a significant due process issue if that was allowed in front of the jury because it would mean that small companies could never seek reasonable damages for use from large companies.

Samsung's patent portfolio. They talk about Samsung owning 200,000 patents worldwide. I want to be very clear. There is not a single shred of evidence of a patent that is used in the accused technologies other than in the Netlist patents.

We asked their 30(b)(6) witnesses. They could not identify any. They didn't identify in their interrogatory

response anywhere. Doctor Robins could not identify any for the '060 and '160 Patents, and Mr. McAlexander could not identify any for the other four patents in this case.

So Samsung's discussion of patents has no place in front of this jury because there's no reference to them.

The last thing is Mr. Meyer talks about the invention of double data rate, DDR, in 1997, and old form factors. These are not accused products, and it would be confusing to the jury to be discussing them. And I don't think it has any place -- there's no nexus to this. In other words, he doesn't say that a contribution in 1997 related to the value of the product.

Your Honor, I will stand and rest on the papers unless you have any specific and targeted questions for me.

THE COURT: I don't.

MR. SHEASBY: Thank you, Your Honor.

THE COURT: Let me hear from Samsung, please.

MR. REGER: May I approach, Your Honor?

THE COURT: You may.

MR. REGER: Can we go to slide 18, please?

THE COURT: You may proceed, Mr. Reger.

MR. REGER: Thank you, Your Honor.

Your Honor, we've heard a couple of different motions and argument today about RAND, JEDEC, essentiality, and standards compliance. And I'd like to take those issues first because

we've -- we've heard such a large amount about it.

One of the last things Mr. Sheasby said was that there was a 408 offer, and he set that to the side. A few minutes later he said, now, Mr. Meyer is trying to hold us accountable for a $6 offer for a set of LRDIMM patents, quote, separate set of patents, was the language he used.

And if I could have slide 24, please.

This is the 408 offer that Mr. Sheasby referenced once, didn't show the context of the letter, just said 408. And then later, he says here's the $6 offer that Mr. Meyer analyzes.

It's not a separate set of patents, Your Honor. These are patents asserted in this case. This is Netlist claiming that these patents are standard essential, this is Netlist claiming that they are making a $6 offer for the entire portfolio for $6. It's not a separate set of patents, Your Honor.

And I'd like to take a step back for a moment and, again, talk a little bit about JEDEC. The issues of RAND and JEDEC are intertwined, but they're not the same, as Your Honor recognized with Mr. Halbert.

THE COURT: I see DDR4 here. I don't see DDR5.

MR. REGER: That's correct, Your Honor. This is a DDR4 LRDIMM offer.

THE COURT: Okay.

MR. REGER: And here they named it 2 DDR4 LRDIMM.

THE COURT: June of '22. Correct?

MR. REGER: Yes, sir.

THE COURT: All right. Go ahead.

MR. REGER: Okay. Thank you.

I'd like to take a step back from RAND for a moment and focus on JEDEC because really one of the things we continue hearing is that there's no reference to -- I want to get this right -- there's no evidence that anyone contributed to the patented technology. That's what we've heard a few times today. Whether it's about patents or about experts or all the different issues we've heard, we've heard that there's no evidence that anyone contributed to the technology at issue in this case.

If I could go to slide 55, please.

When it comes to the JEDEC contributions, the argument made by Netlist is that Mr. Meyer nor Mr. Halbert identifies any nexus between the thousands of contributions they claim to have contributed to JEDEC standards and specific technologies accused in this case. That's their main argument. That's their framing of the issue.

So let's look at for a brief moment what Mr. Halbert and Mr. Meyer did with respect to JEDEC. And before we do that it's important to put their analysis in context. We have, again, Netlist in their infringement contentions accusing

Samsung's DDR4 load-reduced dual in-line memory modules, LRDIMMs. This is their infringement contentions.

After identifying specific models, they say, well, wait a minute, that's not everything. Just to be clear, our DDR4 patents are asserted against any other Samsung DDR4 LRDIMM products made, sold, offered for sale, so forth, that are JEDEC standard compliant memory modules.

So they are making JEDEC standard compliance an issue in this case. And if they are able to prove infringement, it has to be conformed to the theories they disclose in their infringement contentions. Otherwise, they're outside the bounds of what they provided notice for throughout this case.

So they are stuck with this theory, Your Honor. Both in their infringement contentions and through their expert reports, they have repeatedly pointed to JEDEC compliance in their -- with respect to their accusations.

These are their infringement contentions with respect to the '506 and '339 Patent. They identify specific JEDEC standards. In fact, the very first element for the '506 1[i] says, a first data buffer, and they're citing JEDEC standards. The '339 Patent, they cite another JEDEC standard, JESD82-31A.

So Mr. Halbert took that, took them at their word, looked at their infringement contentions, and said, okay, I'm going to look at first DDR4 LRDIMM standards, I'm going to look at DDR5 standards, and I'm going to look at the specific

standards they identified in their infringement contentions. I'm going to look at the standard compliance that they talk about. I'm going to look at the standards that are complied with by the products in the case, the standards that they identify as relating to the technology at issue in this case. He actually looked at the same standards, JESD8232 and 31 and 31A, the same standards identified by Netlist in their infringement contentions for DDR4.

And Mr. Halbert didn't stop there. To set DDR4 to the side for a moment, let's look at -- set DDR4 to the side for a moment. Let's look at DDR5. There's been a little bit of discussion about the DDR5 are the PMIC patents. And Netlist's damages expert, he called them the PMIC patents control.

He says, Well, I call them the PMIC patents because really the DDR5 patents-at-issue in this case, you take the power management control and you put it on the DIMM, it's an onboard PMIC. So I call them the PMIC patents.

Mr. Halbert specifically analyzed the DDR5 PMIC standards. Again, taking Netlist at their word, they are going after standard compliance, they are talking about the PMIC, an onboard PMIC with respect to DDR5. Mr. Halbert analyzed that.

Now, if I may, Your Honor, I need to jump around for just a moment because it's important to put this into context. Mr. Meyer takes Mr. Halbert's analysis, and he responds to Mr.

Kennedy's -- that technical apportionment Mr. Cordell talked about where Mr. Kennedy said, you know, I'm going to look at refrigerators, washing machines, I'm going to look at Netlist DDR 3 products not at issue in this case, I'm going to take all that revenue, all that R&D, and that's how I'm going to come up with this technology apportionment.

Mr. Meyer said, That's ridiculous. That has no relationship to the products at issue in this case, no relationship to the technology at issue in this case, no relationship really to the time frames we're talking about. The patents were all filed and all faced priority years before 2020 and 2021, which is what Mr. Kennedy did.

So to go back and, again, I apologize for jumping around, Your Honor, but what Mr. Kennedy does is he was looking at total sales for 2020 and 2021 just like Your Honor heard, and he apportions that based on this relative value. And he includes products that aren't at issue.

So in responding to this opinion, this exact opinion by Mr. Kennedy, Mr. Meyer did two things. One was JEDEC, and that's -- and that's where he says, wait a minute, if you actually want to talk about the technical contribution of people who contributed to the issues to the products to the technology in this case, you have to look at who actually contributed to the standards of these standard-compliant products you are accusing in this case. That's what Mr. Meyer

does.

And just another point that I want to draw out from Netlist's counsel, he -- they mentioned the 30 percent power savings by moving that PMIC onboard, moving the PMIC onto the DIMM. We heard about the 30 percent with respect to Mr. Kennedy. That 30 percent was, again, moving the PMIC onto the DIMM. That is in the standard, according to Netlist. That's what they claim. They point to the DDR5 standards for that onboard PMIC.

So you have to look at the contributions of everyone involved with that standard. You have to look at the contributions of everyone involved to figure out what portion did Netlist contribute. To give you a preview of what they'll hear in trial--zero.

So ultimately, Your Honor, what Mr. Meyer is doing is trying to respond to Mr. Kennedy's, quite frankly, unfounded, almost ridiculous R&D, that technology apportionment that includes washing machines and phones and dishwashers and DDR3 and everything else. He's trying to say, hey, if you really look at the technology in issue, this is one way to do it. And we're not talking about patents here. I'm not going to do patent counting. I'm going to look at the actual contributors to the standards.

THE COURT: What else, Counsel?

MR. REGER: Okay. Your Honor, if I may, I know Your

Honor's heard plenty about RAND so I will not go through all their representations throughout the case with respect to why -- about their patents being standard essential, their infringement contentions about standards compliance. We'll address that again later I'm sure.

I would like to talk for a moment about the market comparables approach.

If I could go to slide 4, please.

In this case, as Mr. Cordell previewed, there's really only two real agreements in this case that involve Netlist patents, and that's the JDLA and the SK hynix agreement. As Your Honor knows, market comparables are a most reliable method of evaluating patents, and Mr. Kennedy agrees. Mr. Kennedy published an article, saying, looking at the portfolio and valuing the portfolio, determining the patent's value, it's the best way, really.

THE COURT: Let's don't talk about what Mr. Kennedy says. Let's talk about what Mr. Meyer says. That's the report that' before me now.

MR. REGER: Your Honor, I hear you and I will endeavor to do so. He is responding to Mr. Kennedy.

THE COURT: And I intend to allow him to respond to Mr. Kennedy.

MR. REGER: Thank you, sir.

Then I will quickly move on. So what -- we heard a

little bit in the briefing, that black box we heard from Netlist's counsel, that really they're not quite sure what happened in it.

There's over 40 pages of analysis done by Mr. Meyer analyzing only the Netlist-Sk hynix agreement, 40 pages as to one agreement. It's almost as long as the agreement itself. In that 40 pages, he identifies the royalty amounts and confirms that royalty amount based on Netlist's SEC representations.

He valued the per product royalty rates for the license. He apportioned the rate to the patents-in-suit based on Netlist evidence. He bases it based on Netlist's own testimony, actions, and documents. He groups it according to patents, according to the testimony of Mr. Tobin Hobbs, Netlist's patent counsel who testified Netlist values its patents by family.

Now, this is also another interesting point where it seems, depending on the motion, we're hearing different things. In response to Ms. Degnan's argument, Mr. Sheasby came up here and said, Your Honor, Netlist sent over a list of patents in 2016 because they wanted Samsung to get the message. They wanted to move that these are important patents and we wanted Samsung to move.

Well, that same list is what Mr. Meyer relied on as valuable to Netlist because Mr. Kennedy says it doesn't matter

when the hypothetical negotiation is, whether it's 2016, 2018, 2020, 2021. He says, the market in the negotiating position of the parties stays the same. And so based on that, Mr. Meyer looked at the 2015 grouping of the patents, the value placed on those patents by Netlist to Samsung and Micron, and then apportioned the value, given input of Mr. McAlexander and his own analysis.

Your Honor, I know there's a number of other issues in here. I know that time is short. So, Your Honor, I'm happy to go through them or rest on the papers, whichever Your Honor prefers.

THE COURT: Counsel, I think I've heard adequate argument. I don't want to cut you off. If there's something you feel a burning need to express, I don't want you not to have that opportunity. But since you asked me, I don't think I have questions or need of additional argument unless you have something you just haven't covered you think is important.

MR. REGER: You know, Your Honor, I've been excited to get up here all day.

THE COURT: You had to wait your turn.

MR. REGER: Yes, sir, I did, but I'm going to go ahead and pass the podium.

THE COURT: Okay.

MR. SHEASBY: May I respond?

THE COURT: Let's make it brief, Mr. Sheasby.

MR. SHEASBY: So there are two separate agreements. If you go to slide 35.

This is the proposal that was made in 2000 -- while this lawsuit is ongoing that is subject to FRE 408 that did have a $6 offer in it. That is ultra vires. It's FRE 408 and shouldn't be allowed. There's just a confusion, and I don't think it was intentional.

There is another $6 FRAND offer that was made on LRDIMMs back in 2017 that was before the patents at issue had even issued that was made to SK hynix as part of a FRAND offer and because those patents were being asserted as standard essential patents. I don't attribute any ill will to it.

There's two separate $6 offers. One was an FRE 408 in 2022, and one of them was this FRAND negotiation that was done for standard essential patents, none of which were the patents at issue in that case. And I just want to clarify that for the Court.

THE COURT: All right. Thank you.

All right. With regard to Document 214 and the numerous subparts in Plaintiff's motion to strike portions of the expert's rebuttal report tendered by Paul K. Meyer, I'll try to address these in order which they were set forth in the briefing.

With regard to the Meyer market comparables approach,

which is identified as paragraphs 289 through 305 of Mr. Meyer's rebuttal report, I'm going to grant the motion and strike those paragraphs. I think they do fall within the purview of *Apple versus Wi-Lan* and they are probably not reliable in light of that and other precedent.

With regard to the references to RAND, I'm going to grant that portion of the motion as we've established beyond doubt today there are no standard essential patents at issue in this case. Therefore, any mention of a RAND contractual obligation would violate Rule 403 and be highly confusing and prejudicial with limited, if any, probative value.

With reference to the NIAs, I'm going to deny that portion of the motion.

With regard to the references to other litigation and IPRs, consistent with the Court's earlier discussed practices, I'm going to grant that. Again, I think that's mandated by a fair application of Rule 403 where there's considerable risk of prejudice and confusion and limited probative value.

With regard to the references to the alleged JEDEC contributions, I'm going to deny that portion of the motion. I believe this witness is entitled to respond to Mr. Kennedy's opinions. Particularly in that regard, I don't find anything in this section of his report that cannot fairly be addressed by Plaintiff through cross-examination and rises to the level of needing to be excluded.

That's going to be the Court's same ruling on the references to Micron, SK hynix, and Intel R&D investments. That portion is denied, can be fairly addressed by cross-examination.

With regard to the settlement discussions governed by Rule 408 and the particular letter that's been referenced in the argument today, I'm going to grant that. I do believe it falls under Rule 408. It's being used to disprove the amount of a disputed claim, and I think that's what the rule stands to protect against.

With regard to the expert analysis from the 2017 ITC action, I've made it clear earlier today that the ITC action is an adjudicatory process that has limited probative value and a high risk of confusion. I'm going to grant that portion of the motion, primarily on a 403 analysis.

With regard to the comparison of the patent damages to Netlist's market cap, I'm going to grant that. That's highly prejudicial and it's innately pejorative in the way it's presented.

With regard to the references to Samsung's patent portfolio, that's relevant and probative, and I'm going to deny the motion in that regard.

With regard to references to unaccused products, I think that can be fairly dealt with by cross-examination and I'm not going to exclude it. That's not going to preclude relevance

objections if they're well-founded when it's actually put forward before the jury.

And I believe those are all the many subparts of this motion.

Have I missed anything, Counsel?

MR. SHEASBY: No, Your Honor.

THE COURT: All right. Let's move on to Document 206. This is Plaintiff's motion to strike portions of the rebuttal report of Phillip Kline.

MR. ROSEN: May I approach, Your Honor?

THE COURT: You may.

MR. ROSEN: Good afternoon, Your Honor. May it please the Court, Michael Rosen on behalf of Netlist.

THE COURT: Please proceed, Mr. Rosen.

MR. ROSEN: Your Honor, this -- this motion is related in some ways to Docket No. 203, which relates to the testimony -- I mean, the expert opinion of Doctor Groehn. So there may be some common issues that will come up, but for this particular motion it's actually very straightforward.

THE COURT: That's why we're taking it up. It's very short, it's very straightforward, it's a good candidate to end the day on.

MR. ROSEN: Well, Your Honor, this -- this motion centers on two words, one adjective and one noun, and those are inaccurate data. And this is a phrase that Netflix -- I'm

sorry, Netlist respectfully submits that Samsung should be barred from uttering at trial.

Now, if I can get slide 2, please, Mr. Huynh.

So in our opinion, Your Honor, Samsung and its witnesses can feel free to take issue with how Doctor Groehn analyzed and processed his data, they can try to poke holes in how we presented the data, but they should not be allowed to say that the data he used is inaccurate. It's simply not true, it's prejudicial, and it's pejorative. In fact, the data that Doctor Groehn did use is identical to the data that Samsung's expert, Mr. Kline, used.

So this is the opinion, this is the particular phrasing that Mr. Kline uses in his report that he claims Doctor Groehn appears to have relied upon inaccurate data, and citing a particular spreadsheet, which I'll just read into the record so that we have it, and this is an important spreadsheet, it's Sam-net 013333862, and this is the same exact spreadsheet, the same data, that has been used by both experts, and as well by Mr. Kennedy. It's the same spreadsheet and it's acknowledged, in fact, by Mr. Kline that this was in slide No. 6, that this was the spreadsheet that he did use.

In addition to that, perhaps more importantly, Mr. Kline testified that the numbers that he had used were accurate and hadn't been filled with any sorts of errors. Here he testified in particular at his deposition that the product

numbers were accurate. That's in slide No. 8.

In slide No. 9, that he coded -- that he, meaning Doctor Groehn, coded the numbers properly and there were no coding errors in his SQL, S-Q-L for the record, coding of the data; and that there were no errors in the mathematical code that he used, either; and that, finally, in the correlation of the data that Doctor Groehn embarked upon, that his correlation didn't contain any errors, either.

So for all of these key elements, Mr. Kline conceded that not just the source of the data was correct, but then any of the manipulation and processing and analysis of the data was correct.

And then, finally, the point here, Your Honor, is that, in addition, both Mr. Kline and Samsung didn't actually fully analyze the work papers that Doctor Groehn had produced as part of his report.

So -- so here the supporting work papers that -- that he had relied upon included a pretty wide range of different -- different materials here. There were spreadsheets, there were source code, there were notebooks, there were a variety of different statistical programs that are commonly used by economists and by statisticians.

And in its opposition, what -- what Samsung tries to argue is that there are other areas where -- and you're going to hear this from them as well -- where Doctor Groehn

supposedly erred by filtering and processing the data in a different sequence than they could have. And that's an argument that Samsung should feel free to present at trial. It's classic grist for the cross-examination mill.

But what we respectfully suggest is that, in any event, Samsung should not be permitted to use the words inaccurate data to describe what Doctor Groehn did.

And with that, I'll stand unless Your Honor has any questions.

THE COURT: I think I understand your position.

Let me hear from Samsung briefly, please.

MR. REGER: Your Honor, may I approach?

THE COURT: You certainly may.

MR. REGER: Thank you.

THE COURT: Go ahead, Mr. Reger.

MR. REGER: Thank you, Your Honor.

Your Honor, this will be brief. There's the common phrase garbage in, garbage out, and what Doctor Groehn does is he manipulated the data before inputting it into the regression analysis that's really the subject of his report. We'll talk a lot more about it tomorrow, but it is a 10-page report with only four paragraphs dedicated to his so-called analysis.

What is undisputed is that he changed the data before it was input. We agree, I'll stipulate, that Doctor Groehn said

he started with the raw data that we even saw that in the slides that Netlist's counsel showed us. The raw data. Agreed.

The problem is that the Hedonic regression doesn't occur until after Doctor Groehn manipulates it. He could have run regression on the actual sales data that was stipulated to by the parties. He did not. He chose not to.

Instead, he manipulated the data by removing sales that he didn't think were appropriate. He manipulated the data if they weren't big enough for his tests. He manipulated the data by removing whole chunks of sales if they were less than 10,000 units. He relied on inaccurate data to actually perform the Hedonic regression.

So Netlist is trying to do nothing more than shield his expert from criticism by Mr. Kline.

THE COURT: Does Mr. Kline call out this manipulation in his report expressly? Or does he simply just say Doctor Groehn used inaccurate data?

MR. REGER: No, Your Honor. He actually explains the inaccuracies over and over again, and then Mr. Kline actually explains the impact that that has on the damages numbers.

THE COURT: All right. Well, it's clear to me that both these witnesses start with the same raw data, and Mr. Kline is not going to be permitted to convey to the jury that

the data he started with is something different and he's not going to be permitted to call it inaccurate.

Now, if his report otherwise opines that, starting with the same accurate data, he later manipulates it, contaminates it, confuses it, whatever he says was done to it after he acknowledges they started on the same starting line when the gun sounded, I'm not going to strike that. But he's not going to convey that one side had different data than the other side when they started because I don't think it's disputed they started with the same raw data.

MR. REGER: I agree, Your Honor. I apologize.

THE COURT: Now, let me finish.

MR. REGER: Yes, sir.

THE COURT: To the extent he avoids conveying or testifying that the other side relied on initially inaccurate data, I'm not going to grant any of the remaining portion of the Plaintiff's motion. I think it is responsive to the other expert, and they're entitled to respond.

But it's a material misleading or misrepresentation to say that Doctor Groehn relied upon inaccurate data without clarifying that it's the same data that Kline relies on. Groehn may do different things to it once he starts with it, but he starts with the same data that Kline starts with, and that's got to be clear. And if Kline is going to indicate otherwise, then -- let me say it this way.

To the extent his report supports an indication from him to the jury that that's not the case, that they didn't start with the same data, then I'm going to grant the motion in that regard. I don't want this witness conveying to the jury they didn't start out with the same data. What each of them did with it once they started with the same data is fair game. All right?

MR. REGER: Yes, sir.

THE COURT: Beyond that, I'm going to deny the balance of this motion. Okay?

MR. REGER: Thank you, Your Honor.

THE COURT: All right.

Now, we still have a long way to go, but we have been at it since 9:00 this morning. It's now after 5:00 p.m. We're going to recess for the day. I will see you back here at 9:00 in the morning.

I would like to see the same four lead and local counsel in chambers for just a minute before you leave today. Well, I'll tell you what. I'll just talk with you now about what I would have talked with lead and local in chambers.

It is common in the scenario in which we find ourselves--that is, in the middle of pretrial with a plethora of substantive motions--that once parties on both sides of the docket get the guidance provided by the Court's rulings on these dispositive motions, it is common, if not almost

universal, that each party adjusts to those rulings. There's nothing wrong with that.

And if the Plaintiff decides it needs to narrow some of its case, if the Defendant decides it needs to jettison some of its defenses, if there are adjustments that are reasonable in light of the guidance that both sides are getting from the Court's pretrial rulings, there's nothing wrong with that.

What is wrong with that is hiding the ball until it's too late for the other side to respond or hiding the ball until the other side, not knowing the intended change of position, has expended time and resources preparing to present or defend the same.

So if there's going to be narrowing, either on the affirmative Plaintiff's side or the defensive Defendants' side, I understand we're not through the pretrial yet. We're going to get there tomorrow by hook or crook. But once we've completed the pretrial process, if there are going to be adjustments on either or both sides of the case, they need to happen. They don't need to be delayed.

I do not want motions for sanctions that the other side waited to drop this defense or to drop this claim or to dispose of this theory and we've spent time, resources, and money wasted that we can't get back, time we couldn't put on the remaining parts of the case, I don't want that kind of an argument presented to me.

And though there's been, by and large, professional comportment today, I'm well aware of the challenges that the parties have experienced with each other getting this case to where it is. I won't go so far as to say rancor, but there have been some confrontations. There is a late-breaking motion for sanctions before me that is pretty good evidence of that.

So in an effort to avoid more of the same, I'm telling both sides, feel free to make the adjustments you think need to be made to streamline and focus your respective cases in light of the guidance you're getting from the Court, but you're going to do it promptly, you're going to fully disclose it to the other side, and I'm not going to tolerate allegations of ambush or delay or gamesmanship in that regard.

Everybody understand where I'm coming from? Any questions from either Plaintiff or Defendant?

MR. SHEASBY: Nothing from Plaintiff, Your Honor.

MR. CORDELL: Your Honor, I do have a question. The Court sent an order out last week, maybe last weekend, setting the dates for final designation of claims for trial and prior art from the Defense.

THE COURT: And that might need to be expedited.

MR. CORDELL: Given that we moved the trial up a few days, it would be helpful if we could have those --

THE COURT: Well, Monday to Friday's not a huge

jump, but what are you proposing, Mr. Cordell?

MR. CORDELL: I would propose that we get the claims, you know, Thursday or Friday of this week, and we'll give them our defense narrowing on Monday.

MR. SHEASBY: Your Honor, that's an impossibility. I have a contempt hearing on enforcing an injunction on Monday and Tuesday that I'm going to be a hundred percent prepared for.

We can go -- if the Court's order is self-executing, it's 10 days before the jury trial, which means that they'll get them three days earlier now, no matter what. And so I would ask that the Court stick with the narrowing process that has already begun.

THE COURT: I think that's fair. If it's 10 days before the 17th, now it's 10 days before the 14th.

MR. SHEASBY: Thank you, Your Honor.

MR. CORDELL: I may be misremembering the order, but that's --

THE COURT: Well, let's just put a finer point on it. The date we're talking about as it currently stands is, if my memory serves me right, April the 4th. Is that right? Or is it a different date?

MR. CORDELL: It was the 10th -- the 7th was claim narrowing, and the 10th was defense narrowing.

THE COURT: All right. Well, the trial was

previously set on the 17th. Go back 10 days and you get to the 7th. We're now on the 14th. Go back 10 days, and you get to the 4th. So the first step in that process should be on the 4th and not the 7th. The second one should follow, what is it, three days later?

MR. CORDELL: Yes, sir.

THE COURT: That would put it on the 7th. So 4th and the 7th instead of 7th and the 10th.

All right?

MR. CORDELL: Thank you.

MR. SHEASBY: Thank you, Your Honor.

THE COURT: All right, Counsel. We stand in recess until tomorrow morning.

MR. SHEASBY: Thank you, Your Honor.

(The proceedings were concluded at 5:15 p.m.)

I HEREBY CERTIFY THAT THE FOREGOING IS A CORRECT TRANSCRIPT FROM THE RECORD OF PROCEEDINGS IN THE ABOVE-ENTITLED MATTER. I FURTHER CERTIFY THAT THE TRANSCRIPT FEES FORMAT COMPLY WITH THOSE PRESCRIBED BY THE COURT AND THE JUDICIAL CONFERENCE OF THE UNITED STATES.

S/Shawn McRoberts 04/02/2023

_____DATE_____
SHAWN McROBERTS, RMR, CRR
FEDERAL OFFICIAL COURT REPORTER

IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION

NETLIST, INC.,	(	CAUSE NO. 2:21-CV-463-JRG
Plaintiff,	)	
vs.	)	
SAMSUNG ELECTRONICS CO., LTD., et al.,	)	MARSHALL, TEXAS
Defendants.	)	MARCH 29, 2023
	)	9:00 A.M.

VOLUME 2

PRETRIAL CONFERENCE

BEFORE THE HONORABLE RODNEY GILSTRAP
UNITED STATES CHIEF DISTRICT JUDGE

SHAWN McROBERTS, RMR, CRR
100 E. HOUSTON STREET
MARSHALL, TEXAS 75670
(903) 923-8546
shawn_mcroberts@txed.uscourts.gov

A P P E A R A N C E S

FOR THE PLAINTIFF:

IRELL & MANELLA, LLP - LOS ANGELES
1800 AVENUE OF THE STARS
SUITE 900
LOS ANGELES, CA 90067-4276
(310) 203-7096
BY: MR. JASON SHEASBY
MS. YANAN ZHAO

McKOOL SMITH, P.C. - MARSHALL
104 E. HOUSTON ST., SUITE 300
MARSHALL, TEXAS 75670
(903) 923-9000
BY: MR. SAMUEL BAXTER
MS. JENNIFER TRUELOVE

FOR THE DEFENDANTS:

FISH & RICHARDSON PC - WASHINGTON DC
1000 MAINE AVE., SW
SUITE 1000
WASHINGTON, DC 20024
(202) 783-5070
BY: MR. RUFFIN CORDELL
MS. LAUREN DEGNAN
MR. BRIAN LIVEDALEN
MR. MATTHEW MOSTELLER
MR. MICHAEL McKEON

FISH & RICHARDSON, PC - SAN DIEGO
12860 EL CAMINO REAL
SUITE 400
SAN DIEGO, CA 92130
(858) 678-5070
BY: DR. FRANCIS ALBERT

FISH & RICHARDSON, PC - DALLAS
1717 MAIN STREET, SUITE 5000
DALLAS, TEXAS 75201
(214) 747-5070
BY: MR. MATTHEW COLVIN
MR. THOMAS REGER

FISH & RICHARDSON, PC -
NEW YORK
7 TIMES SQUARE, 20TH FLOOR
NEW YORK, NEW YORK 10036
(404) 724-2764
BY: MS. KATHERINE REARDON

GILLAM & SMITH, LLP
303 SOUTH WASHINGTON AVENUE
MARSHALL, TEXAS 75670
(903) 934-8450
BY: MS. MELISSA SMITH

GILLAM & SMITH, LLP
102 N. COLLEGE, SUITE 800
TYLER, TEXAS 75702
(903) 934-8450
BY: MR. TRAVIS UNDERWOOD
MR. TOM GORHAM

OFFICIAL REPORTER: SHAWN M. McROBERTS, RMR, CRR
100 E. HOUSTON STREET
MARSHALL, TEXAS 75670
(903) 923-8546

THE COURT: Be seated, please.

All right. We'll continue the pretrial hearing in the Netlist versus Samsung matter, which has been in recess since yesterday late afternoon. And we'll return to the unaddressed dispositive motions that are before the Court.

Before I do that, is there anything from either side that is either late-breaking or I need to know about? Are both sides ready to go forward with what's remaining, or is --

MR. SHEASBY: Your Honor, we are ready to proceed.

There's some disputes that have been removed from two of Defendant's motions to strike based on the Court's rulings yesterday. I communicated those to the arguing attorney for the Defendants, and my understanding is he will be apprising the Court of those narrow disputes when he introduces --

THE COURT: Okay. Well we'll address them when we get to those motions.

I take it that none of the motions themselves have been resolved completely; only portions based on yesterday's rulings?

MR. SHEASBY: That's correct, Your Honor, in our view.

MR. CORDELL: I did have one additional issue, Your Honor.

THE COURT: Okay. What's that?

MR. CORDELL: So just now we got an email from the

Plaintiff saying that they wanted to substitute exhibits on the exhibit list. Given some of the Court's rulings, some of the exhibits that they had put on the list are no longer relevant, so they want to swap in their prior lists. We don't have a problem with that per se, as long as we're allowed to do it as well. But we thought we needed to get the Court's guidance on that before we began changing the exhibit list.

THE COURT: Well, let's put it this way. Once we get through the dispositive motions, my plan is to take up the disputed limine issues. Once we get through the disputed limine issues, I will next turn to the disputed matters to be considered for pre-admission as exhibits. That's when I'm going to lower the boom on any changes. So if you have substitutions to make, make them before the time the Court turns to the issue of disputed exhibits for pre-admission.

MR. CORDELL: Excellent. Thank you.

THE COURT: And do it on both sides.

And don't be surprised if I order you to meet and confer about narrowing that list of disputed exhibits. I'm probably not doing myself any favors. If I made you stick with the list you have, then apparently a good bit of it would drop out based on these rulings, and now I'm just letting you reload between now and the time we get there, which means the likelihood of the list of exhibits being narrowed any is going out the window.

MR. CORDELL: We will ensure the exhibits we have that are relevant, which I think is important.

THE COURT: All right. I'll stick with what I've just told you. If you want to substitute between now and the time that we get to taking up exhibit disputes, both sides will have that latitude.

MR. CORDELL: Thank you.

THE COURT: But I'm still going to probably ask you to see what can be narrowed when we get there. Because again, counsel, we have today and that is all the time we have, and whatever we don't reach doesn't get reached, at least for argument in this context.

Okay. Let's pick up, then, with the remaining disputed motions. We ended yesterday with argument on Document 206. Let's pick up a somewhat related motion, which is Document 203, and that is Samsung's motion to strike under *Daubert*, testimony of Doctor Groehn.

Let me hear argument from the moving Defendants on this.

MR. REGER: Your Honor, may I approach?

THE COURT: You may.

MR. REGER: Thank you, Your Honor.

THE COURT: All right, Mr. Reger. You may proceed with argument for the Defendants.

MR. REGER: Thank you, Your Honor.

Your Honor, yesterday -- we heard a little bit about this

yesterday when we talked about Mr. Kline, and I previewed the insufficient report of Doctor Groehn. We talked about how both Mr. Kline and Doctor Groehn started with the same raw data. That was accurate and stipulated to. But Doctor Groehn then manipulated that data before putting it into the regression software to get one number out of that software.

Now, Your Honor, we're going to dig a little deeper into Doctor Groehn's report this morning. The good news, it will not take very long, and it will not take very long because it is only 10 pages. The entire report is 10 pages, with the substantive analysis only four paragraphs. So I expect this motion to be resolved relatively quickly.

Now, while there are a number of key issues I want to talk about with respect to Doctor Groehn, the real focus is on the utter lack of disclosure by Doctor Groehn in disclosing his process, his methodology that he followed. Again, his report is only 10 pages. His failure to adequately explain his methodology is enough to strike his limited report in and of itself. And it's meant to shield the jury from confusing unreliable conclusions that are used by Netlist damages expert without any check, without any verification. And Doctor Groehn's unsupported and undisclosed conclusions impact tens of millions of dollars.

So, Your Honor, when we focus on his barely disclosed methodology, we learn that not only did he manipulate the

data, he did not account for known issues in the data, and he did not check his results. Ultimately, Your Honor, Doctor Groehn's proffered testimony is fundamentally flawed from beginning to end.

Now, hedonic regression is a tool that's used in economics quite often. It's often used in real estate, in cars, and things of that sort. Really the key to building a reliable regression is you want to include all the variables that a customer values.

When you look at real estate, you do not just look at square footage and the number of bedrooms; you look at a number of different features with respect to the house. You look at all the specifications of the house, you look at the location of the house, is it close to schools, is it close to a bad neighborhood. You look at all these things to try to determine the impact of each of those variables on the price of the house. You look how big is it, how many bedrooms, the location, the number of bathrooms, the type of roof, does it have a pool or not.

THE COURT: All right. I get the picture.

MR. REGER: Yes, Your Honor.

THE COURT: Let's stop before we get to what color the wall paint is. Okay?

MR. REGER: My wife might find that important, Your Honor, but understand.

At the end of the day, Your Honor, the key to really understanding the impact of a particular feature of a product, you have to know what other features are available and their impact on the price.

Another key factor to building a reliable regression is using the largest data set available. If I run a regression on the sale price of one home, that doesn't tell me anything. Two does not tell me anything. If I go back 10 years and look at all the home sales over the last -- in a particular neighborhood, then I might learn something because I have a larger data set and I can actually start determining the impact of square footage in a particular neighborhood.

Not only that, but you have to start looking at the relationship between the variables that you're looking at. It's called -- one of the things is called multi-collinearity. So many homes have square -- a set square footage and a set number of bedrooms. Oftentimes those features are related. Very few homes have a large square footage and few bedrooms, but some do. So you have to consider that in your analysis.

And ultimately what you have to do is you have to test your results. You say, I've run my regression, this is the result I get, let's try a couple of different -- let's run it a few different ways, a few different times to see if my results are valid and are reliable.

Doctor Groehn did none of this. He has 10 pages, and his

analysis is limited to four paragraphs where he talks about the actual regression that he runs. This is what his -- what we were able to glean from his 10 pages that he disclosed to us via his report.

And then ultimately in his late-disclosed limited work papers and through his deposition, what we were able to learn is he did start with Samsung's accurate and stipulated-to sales data. Then he only chose three variables. There are many variables he could have chosen to look at when it comes to Samsung's memory products. He chose three--time, speed, and capacity. And he's trying to understand their relationships when it comes to price.

Not only did he choose three tightly related variables; he then manipulated Samsung's sales data when it comes to those three variables. He starts consolidating them. He aggregates them. He removes data. Once he does all this, he then removes almost a quarter of the total data set. And again, that's taking -- you look at the entire sales in a neighborhood, and he says, I'm not going to look at these homes because I don't like their data. He then inputs this manipulated data into different software and he outputs one result and one result only.

Now, it's important to note that when it comes to the actual software Doctor Groehn used and relied on, he did not disclose that either. So he didn't disclose his methodology,

he didn't disclose the software he used, and ultimately, as we will see, he did not disclose other critical things.

His ultimate result -- his ultimate conclusion is that he suggests that a one percent increase in speed leads to a 1.445 increase in price. And he has a similar conclusion with respect to the decrease in speed and decrease in price. That's important to note that his confidence interval goes everywhere from 0.5 to 2.3. That doesn't sound very big, but when we're talking tens of millions, hundreds of millions of dollars, this wide range really has an impact and it will be confusing to the jury.

Now, one of the original focuses of our motion with respect to Doctor Groehn is, it appeared to us based upon his limited report, that he was asked to analyze DDR4, he said he analyzed DDR4, and then he concluded with respect to damages generally. And our understanding of that was confirmed, it seemed, because Mr. Kennedy, Netlist's damages expert, uses the DDR4 analysis by Doctor Groehn to analyze--excuse me--not to analyze, he does not analyze it, but he uses that number in the DDR5 calculation. So it appeared to us that they talked and said, Let's analyze DDR4, and we're going to use this in DDR5.

Now, Doctor Groehn pulled back from that and said, No, no, no; I only meant DDR4. Netlist also backtracked and said, No, no, no; we only mean DDR4. Doctor Groehn did not analyze

DDR5 And he did not mean his results to apply to DDR5.

THE COURT: We've already dealt with Mr. Kennedy's report yesterday.

MR. REGER: We have, Your Honor.

THE COURT: We're not going to reopen that.

MR. REGER: Understood, Your Honor.

THE COURT: If you want to complain about Mr. Kennedy and how he applied Doctor Groehn, that's appropriate for when the Kennedy report is being reviewed and tested by the Court.

MR. REGER: Understood, Your Honor.

THE COURT: So if Groehn doesn't take DDR4 and morph it into DDR5, then we need to move on.

MR. REGER: Understood, Your Honor. I was trying to provide some background, but I understand your instructions, sir.

So one of the key factors here is what Doctor Kennedy does not do is he does not analyze does not verify, does not check. And again, I'm not trying to revisit that issue; I just want to explain the context of why Doctor Groehn's lack of methodology, his lack of double checking his results, Doctor Groehn's failures, the impact he can have on the damages case because, again, if Doctor Groehn had said 1.2 or 1.0, 0.6, if Doctor Groehn had come up with different numbers it would have a profound impact on the damages numbers, as we

see here.

One of the key factors, as I mentioned, Your Honor, was actually selecting the correct inputs. And we see different cases where this Court and other courts have stricken damages experts, not for conducting hedonic regression but for improperly conducting hedonic regression, as Doctor Groehn did.

We have cited literature that Doctor Groehn reports as, you know, the theoretical foundation. They admit that selecting the correct variables is important. So Doctor Groehn chose, like I said, three variables with respect to price--speed, capacity, and time. He does not explain in his report why he chose those three variables. He does not explain why he chose only those three variables. He ignores testimony from Samsung witnesses explaining different factors that go into the price and what drives demand in Samsung's memory products. All of these factors are completely ignored by Doctor Groehn. He does not even address this testimony. He doesn't mention this testimony in his four paragraphs describing his analysis.

Other things he ignored in Samsung's sales data are different component compositions, rank and organizations. Do these different factors, do these different features impact the ultimate regression analysis? We don't know because Doctor Groehn doesn't tell us. He doesn't tell us why he

excluded this. He doesn't say he ran this and found no meaningful results. We don't know.

Doctor Groehn admitted his team went to Samsung's website to get some of the information about our products, things he ignored. Again, that could have had an impact on price. He did not explain why he ignored these variables.

THE COURT: This all sounds like good cross examination, but why does it rise to the level of actually being excluded?

MR. REGER: Your Honor, ultimately it is an unreliable -- it cannot be tested because he does not explain -- I understand he's limited to the four corners of his report, but ultimately he is going to be presenting a number on unreliable methodology. That should -- that would prejudice the jury, Your Honor, and confuse the jury.

THE COURT: What else do you have for me?

MR. REGER: That is all, Your Honor.

THE COURT: Let me hear ha response from Netlist.

MR. ROSEN: Good morning, Your Honor.

May I approach?

THE COURT: You may approach, Mr. Rosen.

Please proceed.

MR. ROSEN: Thank you, Your Honor. Good morning. May it please the Court. Michael Rosen for Netlist.

So contrary to counsel for Samsung's representations,

Doctor Groehn's hedonic regression analysis uses commonly employed reliable statistical techniques, all of which Doctor Groehn does explain in detail in his report, in the accompanying materials to his report, and at his deposition. And perhaps most importantly, as Your Honor indicated, all of -- any disputes that Samsung has with his methodologies or his results are classic material for cross examination and not for exclusion.

And I'll just note that toward the end of our session yesterday when we touched on this briefly, I believe counsel referred to Doctor Groehn's work as something like garbage in garbage out, and in addition to being impolite, that's simply an inaccurate characterization of what Doctor Groehn did here, and we'll take a look at that here in detail.

All right. So, first of all, Your Honor, as counsel for Samsung explained, hedonic regression analyses are commonly accepted, widely used techniques in the statistical and, in particular, the economics field. There's nothing here that Doctor Groehn did that hasn't been used in numerous cases that have been validated, and the case law, in fact, itself states that these types of techniques are widely used.

And, in particular, the main issue that Samsung has raised are the variables that Doctor Groehn used or, more importantly, the variables that they claim he did not use. And what we would submit, Your Honor, and is in the papers as

well, is that Doctor Groehn picked exactly the right variables, because variables, if they don't actually vary they're not variables. And in this case every other variable that Samsung claims that Doctor Groehn should have used that he did not use, apart from speed and capacity and price and time, all of these variables either didn't change, don't change at all, or were closely correlated with the variables that did change, and that's how Doctor Groehn explained it.

So overall his finding was, as we saw, that there was a statistically significant relationship between speed and price. And indeed, Mr. Kline, their expert, agreed that if you were to run the same specification you would get that coefficient. You would get that statistically significant result. And that was something he agreed to as well.

Now, so the main variables that he looked at, as we said, were speed, capacity, time, and price. In fact, Samsung's own experts, their technical expert Mr. McAlexander, who we'll be hearing other motions about today, he himself testified at his deposition the primary drivers of pricing are speed, density, which is also known as capacity, and power. We'll come back to power in a minute, but even Samsung's experts agree these are the main variables that matter here.

In addition, the article by Doctor Triplett, the renowned economist that Samsung is fond of quoting and referring to in its papers and at the hearing today, one of the main arguments

or assertions that Doctor Triplett makes is one of the main ways to reduce the multi-collinearity problem we'll come to in a second is to reduce the number of variables. The more variables you put in, the worse your results are going to be.

So we saw a version of this chart earlier that counsel put up, and as you can see, of the many, quote unquote, variables on this page, many of them are identical. Right? The DDR type is the same. The DIMM type is the same here, the DDR type. The speed is what varies, obviously. Voltage or power, which is what Mr. McAlexander referred to, is identical throughout the time. And the rank order, rank by org, and the component composition columns, those do differ, but Doctor Groehn was asked about that at his deposition. Samsung had the chance to ask him about it then, just as they'll have the chance to ask him about it at trial, and he testified that they're very highly correlated with the other two variables.

So essentially voltage, also here he said there was no variability in voltage. There's no reason to include variables that don't vary. And many courts have found on this very specific issue, Your Honor, that the issue of which variables are chosen and are not chosen, that's classic, classic grounds for cross and not for exclusion.

I'll just come quickly to multi-collinearity. Just to be clear, Your Honor, this is the potential for having too many variables confounding the result, and you have to account for

that. Doctor Groehn did just that. He explained in his report how this all works in context of all of these statistical checks on the data. He produced what's called a correlation heat-map in his software. This was something that was produced with his report.

Samsung had every opportunity to review it, understand it, it was coded in the both -- the Mathematica coding language that is very typical in this field, and they had every opportunity to ask him about it at deposition, as they will at trial. He explained in detail at his deposition how the confidence interval relates to multi-collinearity and, in fact, Mr. Kline, their expert, also got to review that.

Now, just one very brief point on -- and this is all very technical, but Doctor Triplett says it is true multi-collinearity can be a problem, but there is little empirical evidence that it's a greater problem in hedonic analysis than it is in any other sort of context. And really the point -- the part, the most important part is to try to reduce the variables as Doctor Groehn did. And here as well, the coefficients, the results were so striking or so dramatic that any collinearity issues are really going to be discounted.

Once again, here specifically courts have found that the issue whether or not multi-collinearity was accounted for properly, this is an area for cross and not for exclusion.

Finally, he also -- Doctor Groehn conducted additional confidence interval statistical checks. He used all of these terms that we won't bore the Court with, but just for the record, the R squared, the P value, the confidence interval, these are all, by the, way the same types of checks that Mr. Kline put in his report that a typical statistician would do, and that's exactly what Doctor Groehn is, although I would say he's even more accomplished than a typical statistician and has specific expertise in the field of hedonic regressions.

We'll just come to one last issue that they raise, which is the issue of the sales data, Your Honor, that Samsung argues that Doctor Groehn inexplicably excluded a certain number of Samsung sales. And the reason he did that, he explained very carefully what he did and using the same exact sales data spreadsheet that everyone else has used. The reason he did this is that there were certain products, certain Samsung products that were either at the end of their life cycle or selling at very small volumes, and the pricing on those objects, on those products, as you can imagine, toward the end of their cycle gets reduced dramatically. In his view those were outlier products. He excluded them from his analysis and from his processing.

Once again, that is a choice that statisticians make in analyzing data. Samsung may disagree with that choice, and

they're free to dispute it, but he explained the nature of that choice, he explained why he did it, and why the outlier should be excluded.

So with that, Your Honor, I'm happy to entertain any questions the Court may have on this issue, but I think it's very straight forward.

THE COURT: All right. Thank you, counsel. I don't have any additional questions at this point.

Anything further from the Defendant before we move on?

MR. REGER: Your Honor, very briefly.

There was a brief mention of a heat-map. And just so the record's clear, I believe the indication was that this was disclosed with his report. That's not accurate. When we reviewed the scarcity of Doctor Groehn's report, we asked for any work papers and notebooks that would come with that, and only then was this given to us.

There is an indication that the heat-map actually addresses the issues that we raise. The fact is it does not. All it does is create a bar chart--nothing more, nothing less. And what's important there, Your Honor, that bar chart is not referenced in his report. They still have not produced this heat-map that supposedly shows collinearity or multi-collinearity. It doesn't change the input, does not change the regression, does not change the output. This heat-map does not fix anything; does not make Doctor Groehn's analysis

anymore reliable or any less prejudicial to the jury because it is actively not part of his report.

Thank you, Your Honor.

THE COURT: All right.

You know, Mr. Reger you said several times in that argument this was an extremely short, brief, and limited in length report, the implication being it can't be very important because it's so short. I'm reminded that the Gettysburg Address is 271 words long, so I'm not sure that length of the document translates automatically into its weight or importance.

All of that observation notwithstanding, I've considered the arguments on Document 203, and I find that these arguments substantially go to the weight of the issues. They're appropriate for vigorous cross examination, but they don't rise -- the complaints don't rise to the level of warranting outright exclusions, and I'm going to deny Defendant's motion to strike the report and testimony of Dr. Andreas Groehn.

All right. Let's go next to Document 204, which is Defendant's *Daubert* motion and motion to strike the testimony of William Henry Mangione-Smith.

And let me hear from the moving Defendants on this, please.

MS. TRUELOVE: Your Honor, could I interrupt before we start this motion and just make a request to the Court?

Mr. Rosen has airplane tickets to travel to Europe to meet his family this afternoon, and we would just respectfully request that he be permitted to leave the hearing in order to make that flight. He is not slated to argue anymore motions before the Court.

THE COURT: All right. To make your flight you need to leave at this time, Mr. Rosen?

MR. ROSEN: Yes, Your Honor.

THE COURT: Okay. We obviously have plenty of lawyers on both sides of this case. I see no problem with that, so you're excused.

MR. ROSEN: Thank you, Your Honor.

THE COURT: Travel safely.

MR. ROSEN: Thank you, sir.

THE COURT: All right. Let's go ahead.

DR. ALBERT: May it please the Court. Frank Albert for Samsung.

THE COURT: Please proceed, Mr. Albert.

DR. ALBERT: May I approach, Your Honor?

THE COURT: You may.

DR. ALBERT: We have six issues in this motion, Your Honor, and the good news is I believe a good bit of it has been resolved.

The first issue is whether Dr. Mangione-Smith can rely on materials the Court has already struck. We believe, after the

redactions provided by opposing counsel and discussions this morning, this is largely resolved. I think there's one small issue that I'll get to right away.

THE COURT: All right.

DR. ALBERT: As a very brief reminder, Your Honor, we were in this court just a couple of months ago about some documents that were disclosed in November that the Court struck from the case that should have been disclosed under Rule 3-2, particularly documents regarding conception, reduction to practice, development of the alleged inventions.

THE COURT: I've reviewed the Court's order from last fall in preparation for this argument.

DR. ALBERT: Thank you, Your Honor.

THE COURT: I've got a hard copy in my hand.

DR. ALBERT: Thank you, Your Honor. I will go ahead and skip right to the meat of the argument, Your Honor.

So we have paragraph 77 of Mangione-Smith's rebuttal report. We believe this is now resolved after conversations with counsel this morning. That's slide 14 of what I just handed up.

Slide 15 of what I just handed up refers to paragraph 313, Mangione-Smith's rebuttal court report. Again, we understand that has also been resolved. Plaintiff will not be relying on either of those paragraphs.

The issue that remains, Your Honor, we think could use

some of the Court's guidance. There are presentations that Netlist claims discloses the invention under Rule 3-2. I have it highlighted here on slide 16. They were required to disclose documents sufficient to evidence each discussion with disclosure to other people or, you know, others. And this is for the purposes of § 102. § 102 is an invalidity issue. It doesn't matter who you disclose it to. If you disclose it to -- the inventions to anybody, it impacts invalidity under § 102.

The issue that we're dealing with, here, Your Honor, is they have largely redacted from their reports the disclosures to the claim or disclosures of the invention to parties other than Samsung, but are maintaining their use as to presentations to Samsung before the application of the patents.

The second issue is the dates of the documents. There are portions of 3-2 that deal with the priority date as the key date for documents that need to be disclosed. You need to have -- but for the documents regarding disclosures to others, the date right there in the local rule is prior to the date of application for the patents-in-suit. And here what we have are documents. Here's a reference to paragraph 19 of Mangione-Smith's opening report. We see there at the bottom this is the Bates number 54708. That's right smack in the middle of that range of documents we were discussing just last

month. And what they're saying right there, you heard it yesterday, this distributed buffer design, that's what Netlist is claiming is the invention. We disagree, but that's what they're claiming is the invention. And this was a presentation that they're using that was before the date of the particular patents that are asserted in this case and was disclosed to Samsung.

We have, again, next slide 19, a reference to another presentation to Samsung, and then listed on page 20. So the issue here is whether the order on the motion to strike materials from just a couple of months ago relates to also disclosures to Samsung before the date of the application of the patents or whether it only refers to people -- parties other than Samsung.

The second issue is Netlist's new infringement theory, and we believe this is resolved but we just want to confirm that it is on the record. So the issue here is the Netlist -- in their report accused an SPD hub as a claimed logic element. That was not disclosed in the infringement contentions. From Mr. Sheasby's representation this morning, we understand Netlist is not going to claim at trial that the SPD hub is a logic element under the claims. If I have that right, then I believe this portion of the motion is resolved.

Moving on to issue No. 3, Your Honor, this is the issue of untimely doctrine of equivalents opinions. As Your Honor

well knows, we have a very active case. We had four contentions -- we were on the fourth amended contentions by the end. Even through the end there were no doctrine of equivalents theories. We saw these in the opening reports -- expert reports from Mangione-Smith. Here they disclose a new theory on this converter circuit element and a new theory on the non-volatile memory elements. And you look back to their contentions and it's just a blanket statement, a blanket reservation of rights, which is not something that this Court has allowed before. You can't just blankly reserve a DoE for any possible issue. You have to put in your DoE arguments into your contentions. We didn't get it in the original contentions, not in the first amended contentions, not in the second, third, fourth. It came in the opening report. We don't have a motion to amend that argues that these DoE contentions are timely or there was good cause. We just don't have any of that.

So with that, I'll move on, unless the Court has any questions.

THE COURT: Let me stop you and ask for a little point of clarification.

On Issue 2, you mentioned the SPD hub and you indicated that that was resolved as being moot. That portion of the briefing also seems to address the mapping of vin_bulk_ovlo, or vin_bulk. Is that also included in the representation that

you made that that's now somehow worked out or resolved, or is that still somehow a live matter we need to focus on?

DR. ALBERT: That's not an issue the Court needs to resolve anymore, Your Honor. That has been resolved.

THE COURT: Okay.

DR. ALBERT: Not by issues this morning, but through the briefing.

THE COURT: I don't care how it got resolved; just so it is. Thank you.

MR. SHEASBY: May I just clarify that, which is --

THE COURT: Let's wait until you get to the podium, Mr. Sheasby.

All right, Mr. Albert. Go ahead.

DR. ALBERT: And to clarify, we are not pursuing that issue for Court resolution.

THE COURT: Understood.

DR. ALBERT: The fourth issue relates to whether a -- their infringement expert can opine on the state of mind issues. The law is clear that an infringement expert is no better than the jury for concluding state of mind. He can talk about the facts, he can talk about what was exchanged between the parties, what was said between the parties, but the conclusion as to what was in Samsung's mind, their intentions, that is not an issue that is appropriate for their infringement expert.

So let me just go through some of these here. Starting with slide 34, down at the bottom we have a statement Samsung would have to act intentionally to avoid discovering its products. That's the expert putting himself in the mind of Samsung, of all Samsung engineers. It's not appropriate.

Slide 35, referencing Mangione-Smith's statement that Samsung did not actually know they are infringing, Samsung was intentionally avoiding learning such information--again, putting himself in the mind of Samsung. It's not appropriate for him to make such statements about Samsung's state of mind.

Paragraph 465, Samsung gained knowledge regarding participation in JEDEC. That's a disputed issue, Your Honor. He's referring to -- there were some disclosures in JEDEC that related to other things, and they're making conclusions that there would have been, you know, knowledge gained about the issues in this case that is a leap for him to do. So again, gaining knowledge is another issue. It's a state of mind issue that's not appropriate for the expert.

And with that I think I can move onto the next issue, unless the Court has any questions.

THE COURT: No. Let's move on.

DR. ALBERT: Great.

Issue No. 5, Your Honor, is fairly simple. There was a deposition and then there was an errata that was issued. The expert is relying on the deposition without any reference to

the errata. We believe that is improper. We -- and in the briefing Netlist says, you know, the deposition and the errata comprise the complete record, and we agree with that. We agree with that. We're not saying that they can't rely on the deposition with the errata. What we're saying is they can't rely on the deposition alone. They have to rely on the deposition and the errata, which is the complete record. They didn't do that. Doctor Mangione-Smith didn't do that and, therefore, his opinions relating to his deposition should be struck.

THE COURT: You know, I have seen this come to the Court's attention in two different ways in the past, counsel. I've seen it where the errata was issued timely and the expert just chose to ignore it. I've also seen it where the errata was not issued timely, the expert went ahead and rendered an opinion without the benefit of the errata, and then when the other side saw the opinions of the experts, suddenly the errata appeared out of the blue to effectively try to fix what the expert had opined on.

The timing here is an issue that is of some interest to the Court. How does the issuance of the errata comport with the rendering of the opinions? Which came first here--the chicken or the egg, in other words?

DR. ALBERT: The errata came first, Your Honor.

THE COURT: Okay. What else?

DR. ALBERT: And last up, admittedly, this may take a little bit more time, but we don't think a whole lot more, and this is relating to the technical valuation opinions. We heard reference to that yesterday. Here is a -- sorry. It's a busy slide. It gives you an overview of the three issues that we're going to deal with. I have a slide for each one that will break it out.

So with regards to the '918 and '054, this is the set of patents that we heard yesterday from Netlist. They claim it's all about the on-board power management. We disagree with that, but that's their claim. And their claim is that, you know, to find some measure of damages, what they do is they calculate the power efficiency of these memory devices with on-board power management compared to not having on-board power management. And they say, Well, based on that power difference, if you wanted the same power, what you would do is you reduce the speed, and that alleged reduction of speed is what is used by Kennedy in order to calculate a kind of dollar figure for this benefit--alleged benefit. And we pointed out in our briefing this conclusion as to this power benefit, you know, what you would actually get, is not supported. It's not there.

And so here I have a -- the issue here is -- the statement is the same DIMM at the same speed would draw 12.3 times 1.43 watts, i.e., 17.6 watts. There's a lot there to

unpack. You heard the 30 percent power efficiency from Mr. Sheasby yesterday. That's what that 1.4 refers to. It's actually 1 divided by, you know, I think 0.7--that 30 percent. And that's how they get the 17.6 watts.

But the key here, the key here, Your Honor, is what Doctor Mangione-Smith says is that what you would want to do for these products, what -- to get the benefit is you would keep the power at 12.3 watts. That's that 12.3 number. That's a number that the products operate at today. And so he would say, Well, if you wanted to keep that number, then you would have to reduce the speed and that would be your technical benefit. And Your Honor, that is a big if. That is a hypothetical that is unsupported. Why would you keep it at 12.3 watts to reduce the speed? Why would you do that? What is the basis for that? This notion that Samsung would automatically do that is not supported. We invited them to explain that in the briefing. We didn't see it.

THE COURT: All right.

DR. ALBERT: What we did see is a claim that power management was somehow the linchpin of this invention--that if you put on this power management chip, that that is the key, that's what they invented. And we asked their expert Doctor Mangione-Smith, you know, Did you actually consider that? Did Netlist invent the idea of on-board power management? And he says, No, that's not something I considered. It's not

something he considered.

So this idea that on-board power management is the invention and the technical benefit and you would do this power savings because of it, it's just not supported in the record.

So we have two key issues here on the '918 and '054 Patent that are not supported. And obviously we can't cross examine Mangione-Smith on this basis for why he says you have to have 12.3 watts when it's not there.

THE COURT: All right.

DR. ALBERT: On the '339, Your Honor, it's actually a little bit simpler. And we heard this claim yesterday as well, that the '339 -- that without it you cannot have two DIMMs per channel. That is a bold statement that was not supported in the record.

Now, we pointed this out, and the response what we got was some of the tersest response I've ever seen in an opposition to a *Daubert*. It was one line pointing to a whole section of a report without pointing to anywhere in that section that matters for this claim. Nowhere in their opposition did they say this 2 DIMMs per channel, that bold claim is supported by this line of this report or this statement in the report. So I don't know what they're going to come up and say, but it's not in the briefing.

We also heard that this concept of a distributed buffer

architecture is the basis of the alleged technical benefits of the '339 Patent. Now, we heard that over and over again yesterday. And if that's the case, then we have a problem because Doctor Mangione-Smith didn't do anything to confirm that was the case. He's not provided an opinion that Netlist invent distributed buffer architecture, but yet claims that the technical benefits is because of the distributed buffer architecture. We have fundamental disconnect there.

Moving on to the '506 Patent, Your Honor, we have a similar issue. The claim from Netlist, the bold claim is that without this accused feature you would not be able to go above 2400 MT/s, mega transfers per second, and that's a bold assertion. We would like to know where that came from. We pointed this out in the -- in our motion, and we got the same terse response--no identification of where particularly this is supported in Doctor Mangione-Smith's report; really just a cite to 10 paragraphs and say, Go look for yourself. That's -- again, that's -- I don't know what they're going to come in and say right now, but I haven't seen where they think that certain number of paragraphs this statement is supported.

And with that, Your Honor, I will pass the podium, unless Your Honor has any questions.

THE COURT: No. Thank you, counsel.

Let me hear a response from Netlist, please.

MR. SHEASBY: May I approach, Your Honor?

THE COURT: You may.

Please proceed when you're ready, Mr. Sheasby.

MR. SHEASBY: Thank you, Your Honor.

So I'll start with the first section, which there was materials subject to the Court's order. I understand that a number of arguments have been withdrawn related to that. In particular, I think there was in the opening motion an attempt to strike these documents, which were either produced by Samsung or were provided during the early parts of litigation that don't discuss our product, don't discuss conception and reduction to practice, and are purely simply disclosures of the patents-in-suit. And I didn't hear any argument -- I heard that this was all resolved except for one presentation, so if I'm wrong I'll just focus on that one presentation to save -- to save time with your permission, Your Honor.

DR. ALBERT: Your Honor, I think there's a number of paragraphs in there. I can direct you to the slide, but there's one issue remaining as to whether --

MR. SHEASBY: I just want to confirm I don't need to fight about this anymore. This is not from the paragraph that you care about.

DR. ALBERT: Okay.

MR. SHEASBY: Okay. So lets's dig into what they care about.

The issue here is, and let me go to it directly. The

issue here is -- and I don't think -- just for the record, I don't think this is being picky that these sorts of materials which are long after the filing date and are disclosing patents I don't think are fairly covered by the Court's order. We took the Court's order seriously. I think there is sometimes when clever lawyers get an order and find a way to sort of parse and cut the mustard so thin that you're not really doing justice to the exact issue. That's not this case. I think we in good faith engaged and cut accordingly.

The one presentation that they're concerned about is this one. So this is a presentation from Samsung's records, and every single one of the paragraphs they cite that's still an issue is this presentation. This presentation was given to them in October 2008. This presentation is not being used for conception or reduction to practice. This presentation is not being used to square behind any reference. This presentation was not given to a third party. It was literally in Samsung's records as well, and it's solely being used for one purpose and one purpose only--it's an advertisement. This is not the invention, but if you wanted to use this, you would be highly motivated to use our invention.

And so the reason why this paragraph is -- why this document is relevant is not for any purpose relating to anything that was the subject of the Court's order. No one's going to say it relates to reduction to practice. It's

solely, solely going to be used to -- as part of a willfulness case to understand why Samsung was motivated to eventually use our technology. And if I'm wrong on that and I'm misrepresenting to you, I know there will be a serious sanction associated with that. I say it again. This is not conception, this is not reduction to practice, this is not disclosure of any element of the invention before the priority date. It has nothing to do with that.

If Your Honor --

THE COURT: And how would you describe with particularity this, quote, presentation so that at a later date if this arises we'll be able to track what's being said today?

MR. SHEASBY: Sure. What I'll say exactly is the following: Samsung was shown a design for load reduction. In order to commercialize that decision for load reduction, they had to adopt the technologies in the '506 and '339 Patent.

THE COURT: All right.

MR. SHEASBY: Infringement theories, I think that has been resolved, and I appreciate Doctor Albert's confirmation that one of them we've withdrawn, which is the SPD hub; the other one they conceded I think is correctly in the disclosures.

The doctrine of equivalents issue, this is sort of a paralysis issue, and we were sensitive to file yet another

infringement -- amendment to the PICs because we think we had properly disclosed it.

Let's go to 18.69.

So we fairly disclosed to them that there would be a doctrine of equivalents that we are maintaining on elements in the claims, to the extent they disputed anything. We then explained to them in our original PICs that we were pointing to LDOs as converter circuits. Samsung remained utterly silent. They didn't dispute that an LDO could be a converter circuit. Through claim construction, no dispute that an LDO could be a converter circuit, so there was no opportunity to seek construction of it or anything about that. And suddenly on November 10th, close to the end of discovery, after claim construction, they say, Oh, you don't include a converter circuit. This was in a late amendment to their supplemental interrogatory responses. And that was very soon -- the expert reports were being due in December -- were due in December.

There's no claim that there's a prosecution history estoppel on the phrase 'converter circuit'. There's no claim that there was any prejudice in the sense that they would have done something differently. They examined, and I think this was -- this was tactical, which is they didn't disclose the sort of the tactic of saying an LDO is not a converter until after they took our 30(b)(6) witness' deposition so they could sort of catch him in some sort of contradiction as to whether

an LDO is a converter or not. Nothing wrong with that. That's good litigation tactics. But the point is, is that they finally said, We don't think an LDO is a converter. We said, Of course an LDO is a converter in our opening, and Mr. McAlexander fully joined it.

So there can be a theoretical debate as to whether the disclosure was sufficient enough in the contentions, but we can't be mindreaders. If they don't disclose the non-infringement contention until the end of fact discovery, if we fairly disclose -- and we specifically told them it's the LDOs; in our PICs from the beginning we say it's the LDOs. And then at the end of fact discovery saying an LDO can't be a converter, there's really no other choice than to engage it in the expert report. It was fully disclosed in the first round and they had the right to respond.

THE COURT: Let me ask you this question. You just said it was fairly disclosed in the first round. How many rounds were there? How many iterations of these contentions were put out?

MR. SHEASBY: Four.

THE COURT: And was it addressed anywhere in the subsequent iterations beyond the first one?

MR. SHEASBY: That it was -- an LDO was a converter circuit? Every single one said the LDO was the converter circuit. We were consistent from the beginning. We never

changed it.

THE COURT: So you carried what was in the preliminary infringement contentions forward into subsequent iterations.

MR. SHEASBY: All the way. We've been saying LDO is the converter circuit from the beginning, and then at the end of fact discovery they say LDO can't be a converter circuit, and then we're in a situation to have to scramble.

THE COURT: All right. What else, Mr. Sheasby?

MR. SHEASBY: The alleged state of mind, this is sort of a difficult thing. There are highlights in the paragraph where Mr. -- Doctor Mangione-Smith hypothesizes about what Samsung's state of mind would be, but there are vast elements of that paragraph that just discuss the factual record about when Samsung was on notice, objectively on notice of certain things. And so based on the Court's order yesterday, I think the idea that an expert -- I think you said that for Mr. Halbert that experts should not be hypothesizing on the state of mind of --

THE COURT: That would apply here as well.

MR. SHEASBY: That would apply here as well. And that would be a blanket thing, but the paragraphs contain lots of perfectly appropriate things about notice and state of mind. So on those paragraphs I would just ask that the Court apply the same standard.

If the Court has any questions, I'll move on to the next one.

THE COURT: No. I mean, my intention here from the briefing and hearing Defendant's argument until you addressed it was to basically say Doctor Mangione-Smith is not going to be permitted to opine on Samsung's state of mind. Now, that should be adequate guidance without going back and taking paragraphs and saying this sentence is in and this sentence is out. And there's no reason why counsel for both sides can't do that with the kind of high-level guidance I'm giving. But that is the Court's guidance--the witness is not going to opine on Samsung's state of mind.

MR. SHEASBY: Okay. And then in terms of the Sung Mo Jung's deposition testimony, I think -- so first off, Doctor Mangione-Smith will testify that he read the entire deposition testimony, I believe. As to the --

THE COURT: This is the one with the errata?

MR. SHEASBY: Yes, this is the one with the errata. I don't think he has the memory to say, I read the pages of the errata or not. They were given to him. As to whether he physically read the errata, we don't know the answer to that because he doesn't remember it anymore. But we know that the original is part of the record, so you can't get away from the original. The way this works is you have the original and then they can also put in the corrected version. And Doctor

Mangione-Smith did quote from the original version and they're fighting about things from 'implements' to 'supports', that they change these individual words.

And there was no idea to be evasive or to sort of do anything underhanded. If Doctor Mangione-Smith publishes this testimony to the jury, depending on what testimony's admitted, we can publish the corrections as well. But there was no intention to play games. These are either minor amendments or improper amendments, and the courts have said that the testimony -- the original testimony is part of the record that he can consider. And they can certainly cross examine him on the fact there were these changes.

I can tell you we do not accept these changes as proper in many instances, and we'll actually be moving to strike the changes as outside the appropriate federal rules because they were substantive.

THE COURT: You will be moving to strike them?

MR. SHEASBY: Yes.

THE COURT: When is that going to be? The day we select the jury? This is pretrial.

MR. SHEASBY: So the way I think it would work is we designate the testimony, they counterdesignate the corrections, and then we'd object to the corrections.

THE COURT: Let me ask you this. Do you dispute Doctor Albert's representation that the errata was available

to Doctor Mangione-Smith at the time he issued his report?

MR. SHEASBY: It was absolutely available to him. It was given to him.

THE COURT: And your argument is he had it, he didn't use it or implement it because, in his view, it was improper?

MR. SHEASBY: No. I'm not going to put any words --

THE COURT: What's your explanation if he had it and it was available to him why he apparently didn't factor it into his report?

MR. SHEASBY: Because he was told he could cite whatever he wanted from the deposition transcripts, and so this is what he picked to cite. He was just given the generic instruction.

THE COURT: All right.

MR. SHEASBY: So he didn't know that there was some magic where he had to put in the footnote with the changing testimony. And we didn't catch it.

THE COURT: Let's go onto the next argument.

MR. SHEASBY: The last issue is technical valuation. We disclosed -- I think there was some suggestion that we didn't disclose the concept of voltage regulation on DIMM and the fact that it is lower power and the fact that our technology allows for multichannels. Scott Milton, our 30(b)(6) witness, discussed that at length, the technical

reason for that.

The reason why the watts are what they are is something I think relatively prosaic. And he talks about this in his report, which is Samsung doesn't get to control the watts that its modules use. The watts are standardized based on DDR5 specification. So you have to fit within the watt envelope of the DDR5. And so the reason why he uses the watts he does is because he talks about in his report that that's the specification for the watts, Your Honor.

THE COURT: All right.

MR. SHEASBY: Okay. And then the next issue, which is the -- if you go to slide 94.

For the '506 Patent, he talks about the fact that it creates a more open data eye, and it's the more open data eye that allows you to get above a 2400 MT/s, so I think that is fairly in his report.

And then as to the two dies per channel, he talks about the fact that the reason why it works is because the two dies per channel allows you to selectively -- the patent allows you to selectively turn off and on data paths between different channels. And he speaks about that in the report, and that's what does 2DPC.

So I'll say he's going to be limited to his report, but his explanation for DDR5 watt envelope is the standard, his explanation for '509 is the width of the data eye that's

created, and then his explanation for the '339 2DPC in paragraph 304 is the -- is the fact that you can address different channels at the same time by turning on and off the data buffer.

Thank you, Your Honor.

THE COURT: All right. Anything further from Samsung, Doctor Albert?

DR. ALBERT: Very briefly, Your Honor.

The -- regarding this presentation you have before you, the notion that they're not going to use this to claim that there is a disclosure of the alleged invention when we heard all -- you know, multiple times yesterday the alleged invention was somehow about this distributed buffer architecture, and this exactly how they want to use this presentation.

With regards to the technical benefits issue, I invited counsel to point you to portions of the report, the '918 and '054 report regarding this requirement for power, and I don't see where that comes from. We have a general cite to a footnote in the slides, but this notion that there is a standardized power is -- I haven't seen them point to that either in the briefing or here today in their report as support for this portion.

So this notion of the requirements, that's absolutely necessary for their opinion, and they seem to concede that if

they can't show that then they can't show technical benefits. I haven't seen it in the briefing, I haven't seen it in the -- you know, in anything they showed you today.

With regards to their current reference to the standard also raises other issues that, if it was ripe, then that would probably have been another motion that Your Honor would have had to deal with today. To claim a technical benefit not only the actual benefit of, you know, speed or power or whatever, but a standardized requirement that's different than the patent, that is inappropriate under the case law, I believe it's -- *Ericsson v. D-Link* would have created that problem. That's not before the Court because that's not something they raised in the briefing because all they did, again, was just point to a whole section of the report without explaining where this came from.

Regarding the '506, Your Honor, here we have the thinnest of references to a portion of the report in I believe it's slide 94 of Plaintiff's slides, a more open data eye contributes to better data signal integrity. That's not the statement. The statement is you cannot do this at all above 2400 MT/s without the claims. How do they go from, Oh, we got something that contributes to a better signal, to, You absolutely cannot do this. And that link, that link, we invited them to explain it and I still haven't heard what the explanation is.

So with that I will be done. Thank you, Your Honor.

THE COURT: Thank you.

All right. With regard to Document 204, Samsung's *Daubert* motion and motion to strike regarding Dr. William Henry Mangione-Smith, with regard to what's been described as these untimely theories based on the corrected and redacted report as submitted, I think that's adequate and I'm going to deny anything beyond what's been corrected and redacted in that regard.

All right. With regard to the alleged late-disclosed conception documents for rebuttal purposes, and Defendant's request to strike paragraph 77, I'm going to deny that motion. Same as to paragraph 313, paragraph 88, and paragraph 692.

With regard to the section of this motion that's been characterized as new infringement theories for the '918 and the '054, I'll accept what appears to be a representation from both parties that these issues are no longer live and are moot and have otherwise been disposed of.

With regard to the assertion of an untimely DoE opinion for the '918 and the'054 Patents, having heard the argument and considered the briefing, the Court's going to deny that portion of the motion.

With regard to allegations regarding Samsung's alleged state of mind, which are addressed by the witness, the Court's going to grant that. As I earlier indicated, Doctor

Mangione-Smith is not going to opine as to Samsung's state of mind, and I'm going to leave the Court's guidance at this higher level rather than delve into paragraph numbers, sentence numbers, parts of sentences. I leave it to the parties to implement accurately my guidance here.

With regard to the issue of the errata and the errata from Mr. Jung's testimony, to me that goes to the credibility and weight of the expert's report and can be dealt with on cross examination. I don't necessarily welcome a number of deposition disputes and counterdesignation disputes, but if that's how it presents itself, I'll take those up when they're presented. At this point I'm not persuaded that there should be an outright exclusion of the differences between the original deposition and the subsequent errata.

And as to what's been designated as the unsupported technical valuation opinions not disclosed in fact discovery, having heard the argument, having considered the briefing, I think this goes more to weight than admissibility, can be fairly dealt with on cross examination, and I'm going to deny the motion as to this portion thereof.

And that appears to me to cover all of what's been disputed in the multi-subparts of this single motion under *Daubert* concerning Doctor Mangione-Smith.

Do you have a question?

DR. ALBERT: May I ask for a clarification, Your

Honor?

THE COURT: You may.

DR. ALBERT: On the first part of the motion regarding the untimely documents, there were paragraphs 77 and 313 that we presented during the argument as moot based on representations that Netlist would not carry those forward. I believe you denied the motion with regards to those paragraphs. Were you referring to the parties' agreement?

THE COURT: I think it all comes out to the same thing.

DR. ALBERT: Okay.

THE COURT: I'm not going to try to stand in the way of the parties' agreement, and to the extent the parties' agreement moots it, fine. My denying those particular paragraphs doesn't indicate an attempt by the Court to circumvent the parties' agreement, if that clarifies your question.

DR. ALBERT: Thank you, Your Honor.

And our understanding of the agreement was that those paragraphs were not going to be presented or relied upon.

Thank you, Your Honor.

THE COURT: Okay. All right, sir. Any other questions about the Court's ruling on this motion?

MR. SHEASBY: Nothing from the Plaintiff, Your Honor.

THE COURT: All right, then. Let's turn next to Document 208, which is the Plaintiff's motion to strike portions of the rebuttal report of Mr. Joseph McAlexander.

Let me hear from Netlist on this, please. Here, again, we've got a multi-facetted motion. One of the side's briefing showed there were seven issues, and by the time I got through reading the other side's briefing they had become 10 issues, so let's see if we can wade through them.

MR. SHEASBY: May I approach, Your Honor?

THE COURT: All right, Mr. Sheasby.

MR. SHEASBY: Can we have the first -- let's go to the table of contents, which is slide 2, Mr. Huynh.

So the first issue is --

No, that's the wrong deck. It's deck 6. It's deck 11.

THE COURT: If I could dangle incentive to counsel, we are going to recess after this motion is heard, so the quicker we get through it, the quicker you can have a recess.

MR. SHEASBY: The first issue, new claim constructions. Let me give you an example, because I know Your Honor has read the briefs and let me sort of speak precisely about what the concern is.

If you'll go to 11.6.

So there's this term called 'converter circuit'. They didn't ask for constructions of 'converter circuit', and then what they do in their rebuttal non-infringement report is they

craft a construction of 'converter circuit'. And the way they did that is not by saying this is what a converter circuit means to a person of ordinary skill in the art and it's plain and ordinary meaning; the way they do to it is exactly the way you would conduct claim construction--they cite specifications, they cite extrinsic evidence and they use that to come to a construction of 'converter circuit'. And so what they've done, in effect, is imported into the claim construction into the expert's brief a claim -- expert's report a claim construction brief.

Their response to this is, We can't seek construction of every term in the patent, and that is a perfectly reasonable and fair position. But if you think there's a term for construction that requires reliance on intrinsic evidence to explain it, requires the importation of textbooks and documents to explain your meaning, that strikes me as something that is going beyond what an expert can do and what is appropriately a claim construction process.

So I don't want to be insensitive to their concerns that this was a very large case with lots of patents, but these are not even situations -- so sometimes people sort of whistle through the graveyard in the sense that someone proposes a term and says, We should construe it; and then we say, Well, let's really not construe it and let's just -- we'll deal with it at the back and we'll fudge it. That's not the situation.

These were terms that there was no fair engagement that they ever thought intrinsic evidence and extrinsic evidence was necessary to inform the term.

Another example of that is non-volatile memory. Their construction of non-volatile memory excludes registers, and the way it excludes registers is not by Mr. McAlexander testifying that non-volatile memory is not present in registers; it's by saying, Well, they discuss registers in the specification, and they discuss those registers, and then sometimes they discuss non-volatile memory, and so non-volatile memory can't be a register.

So this is another example in which -- we can have a fight whether the data, the sort of the solid state material or cells within a register that we're accusing of infringement is non-volatile memory, as that term is understood and it's plain and ordinary meaning, but to say non-volatile memory can't be registers because of the specification is a bridge too far.

Let me give you another example. There's a concept of transferring of data, and Mr. McAlexander is taking the position that a transfer of data cannot be the -- cannot be the transfer of control data to a register.

Now, it's not because he's applying some plain and ordinary meaning of the word 'transfer'. What he does, if you look at his specification, is he literally -- it was a claim

construction declaration put into his expert report. He goes through pages and pages and pages of the specification and discusses how this informs his interpretation of 'transfer'.

Let me give you another example.

If you'll go to 12.14.

This is a bit of a different animal in the sense that we actually had an agreed-upon construction state in which the memory model is operated. We thought we were okay, but then he goes on to say, Well, the state in which the power to DRAM has been removed is not a second operable state. Now, he can talk about why from a plain and ordinary meaning to a person of skill in the art when the DRAM is not operable it's not a -- DRAM is depowered, it's not an operation, but that's not what he does. He goes through page and page and pages of speaking about the transfer of data as it's, quote, used in the specification, and operable state as it's used in the specification, and uses that to tease out a meaning. This is claim construction. This was a claim construction declaration. It should not have been in the specification. It should not be something that's used for expert reports.

I'll give you another example. 'Determine' and 'determining' in the sense that a predetermined amount must be determined in the memory module. This is not just a new claim construction; this is an evasion of what Judge Payne previously construed. So Judge Payne previously said that the

pre-determined amount could be from a source that it was off the module, and the way they get around that is they recite the same portions of the specification to argue that determining of a pre-determined amount has to be on the data buffer, but the way they do that is by citing to the specification, not by a plain and ordinary meaning understanding.

So in our mind -- and this is not -- I think they say, Well, you -- experts have to be able to discuss some amount of the specification. We agree with that. It is a judgment call. Has this gone too far? Has this gone beyond plain and ordinary meaning and has this now entered surreptitious claim construction? And Your Honor, I think it has.

Your Honor, I'll move on to bucket 2, unless you have any questions.

THE COURT: Well, by my calibration, that was bucket 2, but --

MR. SHEASBY: New claim constructions, yes.

THE COURT: Right. What about the non-infringement theory regarding pre-regulated input voltage? Do you have any additional argument on that?

MR. SHEASBY: I do, yes, Your Honor.

THE COURT: Why don't we follow my order and then you-all can jump around to I think my order rather than me spending my time flipping back and forth trying to catch up

with your order. Let's talk about pre-regulated input voltage.

MR. SHEASBY: Sure.

So Doctor Mangione-Smith says separate voltages have to exist for -- that 'pre-regulated input voltage' and 'input voltage' must be two separate entities. We've always said in our original PICs that the input voltage is the pre-regulated input voltage; they are one entity. They talk about how we don't have a pre-regulated input voltage in their interrogatory response, but they never say that the pre-regulated voltage and the input voltage must be two distinct entities. And if they would have said that, we would have fully engaged it.

The reason why I say that is, in Samsung's claim construction brief they say that 'pre-regulated input voltage' and 'input voltage' are presumed to have different meanings, but they didn't say they require different structures in claim construction. The reason why they didn't say that in claim construction is because the Court expressly rejected it. The Court argued that -- the Court noted that there's nothing that required the pre-regulated voltage to be on-module, and so, therefore, the pre-regulated voltage regulation imposes no additional limitation on the construction -- the rate -- the phrase 'pre-regulated voltage' has no additional limitation on it as to how it's constructed, where it's constructed, or what

constructs it.

And so what they've done, quite candidly, is taken what should be a claim construction argument is an input voltage separate -- requires separate structure from a pre-regulated input voltage and convert it into a non-infringement argument in which Mr. McAlexander engages the issue in the way he does with the claim construction declaration. He talks about antecedent basis. He talks about how pre-regulated voltages and regulated voltages, input voltages, are created in embodiments in the specification. None of this is an argument that, one, should have been in claim -- should be in an expert report, because if it was anywhere it should be in claim construction; and two, none of it should have been here at all because they never disclosed this argument that the input voltage and the pre-regulated voltage must be separate.

THE COURT: All right.

MR. SHEASBY: Where should I go next, Your Honor?

THE COURT: Let's next talk about other previously undisclosed non-infringement theories, particularly paragraphs 419 through 30 and paragraphs 572 through 578--excuse me--572 through 581.

MR. SHEASBY: Yes, Your Honor.

THE COURT: That's the conversion circuit that includes at least three --

MR. SHEASBY: Yes, Your Honor.

THE COURT: -- buck converters?

MR. SHEASBY: So if you go to slide 32.

So there is an argument that they disclosed late, in late November, that an LDO is not a converter. There's another argument in which they disclose that not only is it not a converter, it's not a buck converter. In other words, in addition to not being a converter, it doesn't require -- it doesn't satisfy the elements of buck. This is not an argument that was never made, it was not -- the only argument they made in late November was that the LDO was not a voltage conversion circuit. They never made the argument that the LDO was not a buck converter.

So the reason why this is pernicious is we would have spent more time in our expert report discussing what a buck does if they were going to make this argument.

THE COURT: All right. Why don't we talk about what's been characterized as previously undisclosed alleged non-infringing alternatives.

MR. SHEASBY: There's another important one on non-infringement theories, Your Honor. It's relating to configuring information.

THE COURT: Okay.

MR. SHEASBY: And this is another example. They are saying that configuration event -- that error information can't be a configuration information, and that's what they say

in the report. They never in their interrogatory response disclose that configuration information can't be error information. All they disclose is the distinction between non-volatile memory and registers. So this is another example where it was nowhere disclosed in fact discovery that error information could not be configuration information. It appears nowhere.

Undisclosed non-infringing alternatives. This is a legal issue in its purest form. There was an interrogatory in which we asked them to identify any non-infringing alternatives. They took the strategic strategy of -- that's redundant. They took the strategy of identifying -- saying there are no non-infringing alternatives, and then they did something bizarre, which is in their expert report they describe what they quote as non-infringing alternatives, but when we pressed them as to whether they are non-infringing alternatives, Mr. McAlexander never says that he did any analysis as to whether they are available or non-infringing or anything like that; it's just, Let me talk to you about some things that I think would work differently. And we never had the ability to test that information during fact discovery because they said there are no non-infringing alternatives other than our own products.

So this is an issue of prudence. There are situations in which Plaintiffs remain tight-lipped on their technical

damages theories, and in that situation it is impossible for a defendant to engage with the issue of NIA until the expert report. We totally get that. That's not the situation here, and in our interrogatory response in our 30(b)(6) witnesses we gave detailed information of the benefit that we thought achieved. We gave them the roadmap.

And what they did is, having seen this, as opposed to disclosing it -- their theories during fact discovery so we could test their witnesses, they wait until their reply brief to talk about other ways of achieving lower power. But the problem is, is that that creates an absurdity, because the people who know the best, which are the Samsung engineers, were immunized from ever discovering it.

So there is -- there should be no relief given to a plaintiff who sits on their hands and hides behind the log as to the technical benefits, but when it's fairly disclosed in 30(b)(6) so they know, for them to not disclose anything and immunize their witnesses from examination on that subject is, in my mind, a bridge too far, Your Honor.

With your permission, I'll move on.

THE COURT: All right. Let's go to non-infringement arguments based on the specification.

MR. SHEASBY: Yes, Your Honor.

So there are a couple here. So, first off, they talk about 'determine' or 'determining' in a claim 'pre-determined

amount', and the problem is, is this is not a passing reference to the specification to give context for the jury; he's literally, Mr. McAlexander, just conducting claim construction by citing portions of the specification and explaining why that should limit the term of 'non-volatile memory'--sorry--of 'determine' and 'determining'. He does that for 'determine' and 'determining', he does that for 'non-volatile memory', and he does it for -- those are the two primary ones.

So let me take a step back there. It is absolutely the case that the specification can elucidate embodiments that are within the scope of the claim, and that is reasonable. And the use of the specification for that purpose by Mr. McAlexander is fine. What is not appropriate is for Mr. McAlexander to say, And this claim excludes X because X is not in the specification. N, this claim limitation excludes Y, because this term, I know it has this meaning because I looked at the specification. Using the specification to constrain the scope of a claim, using the specification to constrain in an exclusionary manner the meaning of a term is claim construction and it's not appropriate for an expert.

Your Honor, I'll move onto the next one, with your permission.

THE COURT: All right. This is -- the next one relates to what constitutes invention. It's been referred to

as other improper infringement opinions.

MR. SHEASBY: Yes, Your Honor.

THE COURT: It's paragraphs 326 and 27.

MR. SHEASBY: Yes, Your Honor.

So this is just some -- a point of sensitivity, which is that anything -- one of the things that Mr. McAlexander does aggressively, he says, Well, the '981 [sic] Patent and the '054 Patent are all about hybrid memory modules, and we don't have a hybrid memory module. That -- this is in his report on infringement--not in his validity report; in his report on infringement. Talking about the gist of the patent when dealing with infringement is dangerous and it invites error. There's no gist analysis as to whether it's infringed. You don't look at the title of the patent whether it's infringed; you look slowly at the specification. And these paragraphs I think, in my mind, go too far by talking about the, quote, alleged invention. That's not appropriate for a jury to hear.

In addition, it's not appropriate for the jury to hear prosecution history estoppel arguments, which Mr. McAlexander is presenting in paragraph 118. That's really for the judge.

Your Honor, may I move on? I think I may have jumped one by going to prosecution history estoppel.

THE COURT: All right. Let's talk next about paragraphs 103 to 110 and 336 and 337 about prior proceedings, prior Netlist litigation, IPRs et cetera.

MR. SHEASBY: Right. So this has no -- I'll go there.

If I can have slide 59.

So there's a series of references to the IPRs that have been filed in this case. There is a series of references to a previous lawsuit with SK hynix regarding the '185 Patent. It's not a patent in this case. There is a serious reference to ITC proceedings; once again, not a patent in this case. And there is I guess repeated references to the PTAB proceedings.

Lawsuits against SK hynix, ITC proceedings against SK hynix, PTAB proceedings, none of those have any relevance for the purposes of infringement. None of those involve the patents at issue, and so any amount of -- any probativeness would be of marginal value as opposed to its prejudice, Your Honor.

THE COURT: All right. Next let me hear your argument on paragraphs 146 through 148 about prosecution history estoppel.

MR. SHEASBY: Yes, Your Honor.

So there are narrow instances in which the prosecution history can be relevant to put in front of the jury, and that is something that we will engage because the prosecutions will be in the record, but the use of them must be carefully cabined. And in this situation to talk about Netlist

distinguishing prior art on a particular ground, using that as estoppel to make an argument is classically information that is for the judge, not for the jury.

THE COURT: All right. Let's talk about paragraphs 122 through 145 regarding opinions based on documents that Samsung seeks to exclude.

MR. SHEASBY: Yeah. That's moot. Samsung has withdrawn those. Yes, that is moot. Samsung has withdrawn those.

THE COURT: Then let's talk next about the relative technical value issue as being unreliable. It's paragraph 655 through 659.

MR. SHEASBY: Yes.

So this is 11.78. And just to talk to you what the issue is, is Mr. McAlexander purported to do an analysis in which he analyzed the value of patents within families, and used that to assign values with those families. But the problem is, is on cross examination at his deposition he admitted not doing it. And so this is a situation where I don't actually know how to handle this. There's something in his report that purports to analyze the value of the patents, but he admits that he never analyzed the value of the patents in his deposition, and so I don't know how to navigate that. It strikes me as it could be cross examination, but if he literally didn't do the analysis that's in his report, it

should not even be cross examination; it should just be stricken from his report.

Thank you, Your Honor.

THE COURT: All right. Let me hear from Samsung. We'll take these up in the same order.

MR. LIVEDALEN: Good morning, Your Honor.

THE COURT: Just a minute, counsel. Let's get everybody in their places.

MR. LIVEDALEN: I have slides for you. Just one second.

May I approach?

THE COURT: You may.

MR. LIVEDALEN: Thank you.

Good morning. Brian Livedalen for the Defendant.

THE COURT: Please proceed, counsel.

MR. LIVEDALEN: Thank you, Your Honor.

THE COURT: Let's start with the pre-regulated input voltage issue that's been discussed.

MR. LIVEDALEN: Sure.

If I can go to slide 16, please.

So this issue, on the one hand there is an issue of whether or not this theory was disclosed, but there's a second issue of whether or not Mr. McAlexander is making improper claim construction arguments.

THE COURT: And it's the second one I'm the most

interested in.

MR. LIVEDALEN: Okay. So this issue persists across many of the claimed limitations that Mr. Sheasby discussed, and really the issue is that every time Netlist disagrees with a non-infringement argument, it asserts that Mr. McAlexander is applying an improper claim construction. But Mr. McAlexander applies the plain meaning throughout his report. They had an opportunity to ask him at his deposition if he applied some other meaning. They did not.

And so what Netlist is asking to do is that Mr. McAlexander remove any analysis, whether it's reliance on extrinsic evidence or the specification itself to confirm the plain meaning. And so what Netlist seeks is to strike any rationale such that Mr. McAlexander is essentially arguing in a vacuum. He certainly needs to be able to look at the patents, he needs to look at textbooks and things to confirm what a person of skill in the art would understand.

So with respect to the pre-regulated input voltage issue, the first question is whether or not it's properly disclosed. As you see here in our interrogatory response, we clearly indicated that there was not a pre-regulated input voltage. Mr. Sheasby is asking that there be a greater level of clarity, but as we've seen in their disclosures, for example, the doctrine of equivalents, they made broad disclosures and relying on those for specific arguments. Mr. McAlexander is

clearly just providing additional detail that was not in the interrogatory response. But as we see here, the non-infringement position was clearly disclosed in discovery.

There's a separate argument that this non-infringement position is inconsistent with what Judge Payne ruled on. That is also incorrect. Judge Payne never addressed whether the pre-regulated input voltage and the input voltage are separate. He clearly -- he only addressed that the pre-regulated input voltage must be generated on the memory module from the input model, so there's an issue as to where the pre-regulated input voltage would be created. He disagreed with Samsung's position on that and gave it his plain meaning, but it is not an issue where he separately addressed the constructions of pre-regulated input voltage and input voltage and their relationship.

And so when you get to Mr. McAlexander's opinions, he looks at the accused products, he looks at the structure of the claim, and he asserts as a matter of plain meaning that simply that the voltage that Doctor Mangione-Smith points to being in bulk does not satisfy both the pre-regulated input voltage and the input voltage, according that their plain meaning.

THE COURT: All right. Let's go onto this category that's been entitled 'new claim constructions'.

MR. LIVEDALEN: Yep.

THE COURT: And there are several subparts to the subparts, so let's just walk through them as quickly as we can.

MR. LIVEDALEN: Yes, Your Honor.

Let's go to --

THE COURT: Do you want to talk about converter circuit?

MR. LIVEDALEN: Sure. We can go there.

We can go to slide 8.

So as you'll see here, there's essentially a battle of the experts as to what is the plain meaning of 'converter circuit'. Doctor Mangione-Smith in his deposition said a converter circuit is anything that changes the voltage, and, in fact, he said a wire can be a converter circuit. So he's taking a super broad definition of what a converter circuit means to one of skill in the art.

Mr. McAlexander, on the other hand, says as a factual matter that a load drop-out regulator, which is called an LDO, is not a converter circuit. That's his opinion. And he defines what a person of skill in the art would understand the converter circuit to mean.

As we see here in the *Gree v. Supercell* case, it says, Doctor Claypool is giving an opinion as to what a POSITA would consider the plain meaning. This constitutes a non-infringement position; it does not amount to claim

construction. So we see here Mr. McAlexander is doing the exact same thing. He's looking at the term 'converter circuit', he's instructing what that means to a person of skill in the art, and he's rendering a factual opinion that the LDO does not constitute a converter circuit.

And we see that here at slide 9 he identifies what a person of skill in the art would understand is a converter circuit, and then he actually identifies textbooks hand other evidence that confirms his opinion. And he says that unlike the converter circuit that converts power from one form to another, a linear drop-out regulator dissipates heat essentially by throwing away the power. And we see here in paragraph 404 he cites to textbooks that explicitly distinguishes an LDO from a converter circuit. So this is a purely a factual issue; it's not a matter of claim construction.

THE COURT: All right.

MR. LIVEDALEN: Is there another term -- I guess one more issue, Your Honor.

Doctor Mangione-Smith, Netlist's own expert, as we see here for purposes of validity says, if you look at the bottom, that a person of skill in the art would have used an LDO, and he makes this opinion to say that the converter circuit limitation would not have been rendered obvious. So this is another issue where even Doctor Mangione-Smith in

instances--not for infringement, to be clear, but for validity purposes--has taken the position that an LDO is different. So there's no disagreement there.

Which term would you like to go to next, Your Honor?

THE COURT: Quite honestly, I think -- I mean, we can go through these one by one, but it's effectively the same back and forth on all of them, or the majority of them.

I would like to hear your argument on the determine/determining.

MR. LIVEDALEN: Yes, Your Honor.

So we see here -- this is claim 1 of the '506 Patent, and the issue is that Mr. McAlexander is looking at the logical structure of the claim. The entire claim is directed to a memory module comprising. Right? We see that here--the memory module comprising. And his opinion is that there is a wherein limitation that talks about a determining step and the logic of reading a claim directed to a memory with the wherein limitation saying a requirement about the determining step is done by the memory module. This is not a claim construction argument. He's looking at the logical relationship of the limitation at issue to the rest of the claim.

With respect to the other issues that Mr. Sheasby raised, he's arguing that any time that Mr. McAlexander goes to the specification to confirm his understanding is inappropriate. But at no time has Mr. Sheasby identified an area where

Mr. McAlexander's arguing for a disclaimer. Nowhere had Mr. McAlexander said there's a lexicographical definition. Nowhere has he said, I'm going to limit it to an embodiment. He identifies the plain meaning, and in instances he looks to a textbook or he'll look to the specification to confirm that his plain meaning is correct. And the reason why he does that in many instances is because Mr. Doctor Mangione-Smith is interpreting claims beyond their plain meaning. He's saying look off the module; it can be done anywhere. But that doesn't make any sense. That's not what the claim says and that's not what the patent says.

THE COURT: All right. Let's talk about previously undisclosed non-infringement theories, Issue 3 in my agenda. That's paragraphs 419 through 30 and 572 through 581.

MR. LIVEDALEN: Thank you, Your Honor.

Let's go to slide 22, please.

So this issue's a bit strange. So Mr. Sheasby argued that Samsung only said that an LDO is not a converter but never said that it's not a buck converter. NOw, does that make any sense? If I say that my boat is not a car, do I have to also say that it's not a Lexus, as we talked about yesterday? A buck converter is a type of converter.

THE COURT: Yesterday we heard about Corollas versus Lexus.

MR. LIVEDALEN: We're going to get more into that in

a minute.

THE COURT: A Corolla is a model; Toyota is a make. Lexus is a make. Lexus has various models like an ES300, but we heard makes compared to models yesterday--Corollas compared to Lexus. It would have been more accurate to have Corollas compared to ES300s, but we'll overlook that inaccuracy and move on.

MR. LIVEDALEN: The point is--right?--

THE COURT: I've owned several of them, so I know.

MR. LIVEDALEN: So we said an LDO is not a converter, so that means it's not any type of converter. It's not a buck converter. It's not a boost converter. So the fact that Mr. Sheasby is saying they weren't put on notice of this argument is just inconsistent.

The reason why they are trying to get this non-infringement argument out of the case is because Mr. Mangione -- Doctor Mangione-Smith argues in one instance that an LDO is a buck converter, and then for validity purposes he says an LDO would not render obvious a buck converter. So he can't have it both ways. They are trying to take this issue out of the case so Doctor Mangione-Smith can skirt over this issue.

And we see here he says, Doctor Mangione-Smith during his deposition, a buck converter is a particular form of voltage converter. So if we say an LDO is not a converter, it's not a

buck converter.

THE COURT: All right. Let's talk about undisclosed alleged NIAs.

MR. LIVEDALEN: Slide 27, please.

So we start out with Mr. Sheasby saying that Netlist disclosed all these non-infringing alternatives. That's just flat-out wrong. We see here -- this is what he pointed to. He says that -- we said, Hey -- you know, we said that VR on-DIMM is good because it gives lower power, or that the timing adjustment is good because it gives you higher speed operation. That is not a disclosure of what the actual non-infringing alternative is. They never disclosed during fact discovery what the actual non-infringing alternative was that they are going to rely on for the incremental benefit of the patent. Had they done that, then we could have responded and provided narrower non-infringing alternatives.

And the reason why this is an issue is because when you look at Doctor Mangione-Smith's report that -- what a non-infringing alternative is, he goes way too far. So to come back to our analogy, he says Lexus is the accused product, and he says the Model T is a non-infringing alternative. Right?

And so Mr. McAlexander needs to say, Wait a minute; there's lots of prior art alternatives that are between the Lexus on the one hand and the Model T from many, many years

ago. And so that's what Mr. McAlexander does--he critiques Doctor Mangione-Smith's non-infringing alternative analysis. He says, Wait a minute, you picked a non-infringing alternative way too far. There are many different alternatives you didn't consider between the accused product and what you're identifying as the non-infringing alternative.

THE COURT: I'd like some specific argument on paragraph 685. 685 is raised by Netlist in their reply saying that Mr. Meyer relies on Mr. McAlexander's rebuttal report to conclude that Samsung could readily implement numerous other viable alternatives, and then Samsung in its surreply says that it does not affirmatively rely on NIAs to compute damages. Can I get some specific focus on that one paragraph, please?

MR. LIVEDALEN: Sure.

Right. So the issue is, is that, as I said, Doctor Mangione-Smith originally asserted that the only non-infringing alternative is to move the power management IC from the module back into the host. And so what Mr. McAlexander does in response is he says, Well, the patents themselves are very narrow. It recites the specific types of buck converters, there has to be a certain number, and they have to perform certain functions. And Doctor Mangione-Smith didn't recognize the simple fact that you could replace an LDO with a buck converter. And for validity purposes Doctor

Mangione-Smith says an LDO is not a buck converter.

So that's a simple alternative that Doctor Mangione-Smith did not consider. And that's what Mr. McAlexander identifies in rebuttal as one example of a non-infringing alternative that he did not consider, and that is something that Samsung could have implemented to design around the patent in terms of the non-infringing alternative.

THE COURT: Do you have paragraph 685 you can put on the screen for me? Either --

MR. LIVEDALEN: Yes. I can put it on the elmo, if we can switch over.

THE COURT: All right. Is that the conclusion of the paragraph, or does it carry over to the next page?

MR. LIVEDALEN: That's the conclusion, Your Honor.

THE COURT: Okay.

MR. LIVEDALEN: And so preceding this is he discusses what's in the prior art, he discusses as the example of simply replacing one of these buck converters with another element that's outside the scope of the claim, and provides all those as his counterexamples to Doctor Mangione-Smith's opinion about his one non-infringing alternative.

THE COURT: Okay. Let's move on to the next issue in this motion regarding non-infringement arguments based on the specification.

MR. LIVEDALEN: Go to slide 32, please.

So I think this issue spills over to what Mr. Sheasby's calling improper claim construction arguments. Essentially he's trying to take away Mr. McAlexander's ability to look at the specification to confirm his non-infringement position. And we see here that when you're applying the plain and ordinary meaning, you need to consult the specification. We cited the case from *Apple v. Samsung* that says it's appropriate for an expert to look at the specification to confirm the plain meaning.

So this is not, again, an instance where Mr. McAlexander is using the specification to opine on a disclaimer, to opine on there's a specific definition provided in the specification; he's simply using it to confirm that his plain meaning opinion is appropriate and correct and that Doctor Mangione-Smith's interpretation is correct.

Do you have any specific questions on any particular term?

THE COURT: Well, we are revisiting the determine/determining term again --

MR. LIVEDALEN: Right.

THE COURT: -- that's come up earlier.

MR. LIVEDALEN: Yep.

So if you look at slide 34.

Mr. McAlexander -- we talked earlier about the claim structure of claim 1 of the '506 Patent, that the structure

indicates as a matter of common sense that determining is done on the module because the claim is directed to a module, and Mr. McAlexander here is only pointing to instances in the specification to say that's what the whole invention is about; it's about putting the determining step on the module.

THE COURT: Let's talk about paragraphs 326 and 27 and what constitutes invention with regard to the '918 and the '054.

From what I've seen in the briefing, counsel, it looks to me like Samsung's position is a refrain I've heard many times, which is, This is just background, Your Honor. Is that what you're saying here, or is there more to this?

MR. LIVEDALEN: You said 315, Your Honor? I want to make sure we're on the same page.

THE COURT: 326 through 27. It's been characterized as other improper infringement opinions.

MR. LIVEDALEN: Yep. That's right, Your Honor. So my slides have the other set of patents. But that's exactly right.

So essentially what's going on here is Doctor Mangione-Smith in his opening report, which is obviously about infringement, but he does other things like characterize what the patents are about, he characterizes the contribution, the technical benefit, in other words, of the asserted patents, and so Mr. McAlexander here is only providing rebuttal

testimony on those points. He's not going to argue that you should look at the prior art to determine infringement. He's not going to argue that you should look at the patent title to determine infringement. It's just clearly a different part of his report.

THE COURT: This is where I think I've seen some of his verbiage where he uses the phrase 'the alleged invention'. I think there's a somewhat pejorative flavor to that. I'm going to order him to modify that. And I'm not going to strike it, but he's going to say 'asserted invention', not 'alleged invention'. And to the extent it is -- you know, It's just background, Your Honor, is something that's covered a multitude of sins in the past. If that's truly what it is, then it's okay, but that usually is an indication that it might be something much more than that.

I'm not going to strike this particular provision as long as it is limited to truly background. If it becomes a sword instead of a shield, then I'm going to expect an objection to be raised and I'm probably going to grant that objection.

MR. LIVEDALEN: Yes, Your Honor. Understood.

THE COURT: Okay. Let's talk about the references to the prior Netlist litigation, the IPRs, et cetera.

MR. LIVEDALEN: Yes, Your Honor.

Slide 42, please.

So I think there are two issues here. One is the issue

of a motion in limine. That doesn't really resolve the whole issue--right?--because Mr. McAlexander's opinions go to both equitable and issues that are going to be heard by the jury. So to the extent that his opinions go to the equitable issues, I don't think that the motion in limine would apply. There wouldn't be any cause for confusion or prejudice. So I think -- for that reason alone I think that the opinions should not be struck.

With respect to issues that would be heard by the jury, I think we talked about this yesterday that if Netlist were to open the door, for example, for willful infringement, to say that there is an issue of standards compliance and that they should have known that these patents were seen as complying and that Samsung would infringe, in that instance we believe that we should be able to identify instances in other courts that have found the exact same patents that Netlist said were standard compliant were actually were not infringed by standard compliant products. So we think that's a fair rebuttal.

But until Netlist opens the door, you know, we have no intent to introduce prior proceedings, per the Court's standard MIL No. 13.

THE COURT: Talk to me about the prosecution history issue as estoppel. That's paragraphs 146 through 148.

MR. LIVEDALEN: Yes, Your Honor.

So I think we may be in agreement here. Mr. McAlexander is not going to argue as a legal matter that there's been an estoppel. We do think, though, he should be able to talk about prosecution history. And to the extent that Netlist is on record of saying that its patent does not cover a particular technology, that we should be able to provide that information before the jury. We see that here in slide 45, another ruling--right?--where the Court has granted the --

THE COURT: Did you pick that case because I was counsel for C.C. Forbes in that case, or is that just a coincidence.

MR. LIVEDALEN: I think that's a coincidence, Your Honor.

So we don't -- Mr. McAlexander is not going to offer any legal conclusions, but we do think the prosecution history is an important fact that he should be able to consider.

THE COURT: All right. Mr. Sheasby represented that the next issue, opinions based on documents Samsung seeks to exclude, which targets paragraphs 122 through 145 of Mr. McAlexander's report has been withdrawn. Can you confirm that for me, or do we still have an issue.

MR. LIVEDALEN: That is correct, Your Honor.

THE COURT: All right. Then I'll note that it is withdrawn.

Lastly, talk to me about this relative technical value

analysis that's at paragraphs 655 through 59.

MR. LIVEDALEN: Yes, Your Honor.

So what Mr. McAlexander did is he was asked by Mr. Meyer, the damages expert, to evaluate the relative technical contribution of certain patent families compared to the patent portfolio as a whole, and so to do that he was asked to consider nine patent families. And there -- to be clear, there was other patents that Netlist has that were not considered, and the reason why he did that, it was to be conservative. Right? So the ones he did not consider, he assigned a zero value, which means the patents-in-suit would have automatically been given a greater value.

And so we see here what Mr. McAlexander did to do this analysis. He considered the area of technology of the patent family, the technical scope of the patent family claims, the potential applicability or use of the claimed technology in the field, and the number of other family members within a family.

And so Netlist does not take issue with this analysis at all, did not say that he forgot an important issue, didn't say that he considered something he shouldn't have considered, so they don't take issue with his methodology here. What they do take issue with is that they think that his analysis was not sufficient. They criticize his analysis for not considering or being detailed enough and going into details about the

specific patents. But what he does is consider patent families. As we see here, family 2 has six patents. And so if he looked at them, he considered the applicability to the actual products that are being sold today and their scope that they would claim.

And so Netlist's criticism here, even if true, would go to the weight of his testimony and not the admissibility.

We see here in the *i4i* case from the Federal Circuit that Mr. Sheasby cited previously yesterday. The one thing I want to clarify is we saw Mr. Sheasby get up here and he said, I asked Mr. McAlexander about this analysis, and he said, I didn't even do the analysis. That's incorrect. What Mr. McAlexander said is that he didn't do the same analysis that he did for the KIP and the Rambus portfolios, and that's because for that analysis he did a comparability analysis between specific patents in those portfolios to the Netlist patents. That's a different thing than he did here with respect to this technical contribution.

And Mr. Sheasby said, Well, if he said he didn't do the analysis, I can't even ask him about his analysis. Well, Mr. Sheasby did not ask a single question about any of the paragraphs in his actual report. So we see here a paragraph from his report, paragraph 659, Mr. Sheasby took Mr. McAlexander's deposition, had plenty of time to ask him about whether or not he wrote this paragraph, whether or not

he had any information about the paragraph, whether he had any other opinions that were not in here about the paragraph. So the fact that Mr. Sheasby says, Aha, Mr. McAlexander didn't do the analysis so we should just strike the whole thing is just not true.

THE COURT: Okay. What else, counsel?

MR. LIVEDALEN: That's all that I have, unless you have any other questions, Your Honor.

THE COURT: I think I've asked my questions as we've gone along. Thank you for the argument.

MR. LIVEDALEN: Thank you, Your Honor.

THE COURT: Mr. Sheasby, is there anything you feel compelled to revisit?

MR. SHEASBY: One thing.

If I can have paragraph 366 and 367.

So in this paragraph --

THE COURT: Let me save you some time, Mr. Sheasby. I'm going to strike 366 and 367.

MR. SHEASBY: Thank you, Your Honor.

THE COURT: All right. With regard to Document 208 and the Plaintiff's motion to strike portions of the rebuttal expert report of Joseph McAlexander and the multiple subparts that have been set forth and argued in this motion, as I noted, I am going to strike paragraphs 366 and 367. I don't think it's accurate that Judge Payne considered the

Defendant's arguments and found the claim term to not be so limiting that 'pre-regulated input voltage' need not be the same as 'input voltage'. Well, he did consider the arguments and he found the claim term not limiting as presented at claim construction, and not limiting in the way that 'pre-regulated input voltage' is not necessarily the same and need not be the same as 'input voltage' referenced at claim 16. And I think this expert report treads on that and should be properly struck, and I'm striking those two paragraphs.

With regard to the portion of the motion that deals with what's alleged to be new claim constructions, I'll take these in turn.

Paragraphs 402, 406 through 408, that's denied. That's the issue with regard to converter circuit. I think that's short of new claim construction and is the expert talking about plain and ordinary meaning. And it's clear we're going to have a battle of the experts in various places throughout this trial. This is one of them.

In paragraphs 529 through 2534 and 537 through 538, same ruling.

In paragraphs 530, 533 and 540, that's on updating the bits of registers, same ruling.

On paragraphs 492 through 495 and 504 dealing with the interpretation of 'operable state', that's the same ruling. That's denied.

However, when we get to paragraphs 167 through 172, the Court's persuaded that the expert does cross the line when it comes to determine/determining and does engage in what's effectively claim construction, and I'm going to strike paragraphs 167 through 172 of this report.

With regard to what's been called previously undisclosed non-infringement theories, with regard to paragraphs 419 through 430 and the buck converter issue, as it's been described, and with regard to paragraphs 572 through 581 where what is and isn't configure information is disputed, those objections are denied as well, and I'm not going to strike those portions of the report.

Moving on to what's called previously undisclosed alleged non-infringing alternatives, this one is very close to the line, in my view. I'm not going to strike these provisions, but I am going to preclude the use of these paragraphs about technical valuation and benefits from being used beyond rebuttal to Doctor Mangione-Smith's opinions about technical valuations and benefits, and I'm going to limit them to that and I'm going to preclude their use as support for ultimate damages opinions. The back and forth between McAlexander and Mangione-Smith is fair, but going beyond that and taking these opinions as an entree into the damages area beyond these experts field of knowledge is not proper. That's the best way I know how to deal with that issue.

With regard to the asserted non-infringement arguments based on the specification, that's paragraph 169, paragraph 235 through 36, and 251 through 262, I'm going to strike those provisions. I think they do improperly constrain based on the use of the specification.

With regard to what's called other improper infringement opinions, paragraphs 326 and 27 and 88 through 99, 115 through 121, I've previously addressed this. I'm not going to strike these. I am going to require that they be limited to truly background information. And I am going to direct the witness not to use the phrase 'alleged invention', but alternatively to use the phrase 'asserted invention'.

With regard to what's called reference to prior proceedings, this includes references to the Netlist litigation elsewhere, Samsung -- parties other than Samsung in litigation with Netlist, I should say, and IPRs, et cetera, there's no justification for the use of this. The prejudicial effect, the confusion of the jury, far outweighs any probative value.

As far as the jury's concerned, I'm going to strike these. Now, when we get past the verdict and when we have a bench trial on equitable defenses, certainly I don't have a problem with this, but I'm not going to permit it for use before the jury.

Now, that being said, no door that's been closed can't be

opened, and as Defendants argued, if the Plaintiff should open the door to these matters, then they certainly have the right to approach and ask for leave to counter them. But without a door opening recognized by the Court in advance, these are not going to be gone into by the expert. And barring something unforeseen, they are excluded from his testimony.

With regard to what's called prosecution history estoppel as a question of law, this deals primarily with paragraphs 146 through 148. To the extent these paragraphs include legal -- or assert legal conclusions--this is prosecution history estoppel, for example--that's granted. To the extent they don't assert actual legal conclusions but discuss the concept, the factual story concerning what prosecution history is and what it was in this case, the factual portions of these paragraphs are not excluded, but to the extent the expert goes beyond the facts and opines about legal matters and says this is estoppel based on prosecution history estoppel, that's excluded and those opinions are struck.

Okay. The issues regarding paragraphs 122 through 145 are no longer before the Court and, as the parties have confirmed on the record, that's been withdrawn.

With regard to the relative technical value issue as being unreliable, the Court believes this goes to the weight and not the admissibility, and I am going to deny the motion to strike this portion of the report. That's paragraph 655

through 659.

And those are the Court's rulings with regard to Document 208.

All right, counsel. It's 11:25. Let me give you a choice. We're close enough to break for lunch instead of taking a recess, which my earlier incentive of having short argument to get us to a recess obviously didn't work. Or we can take a short recess, come back, and work through the noon hour. Do you have a strong preference?

MR. SHEASBY: Plaintiffs would request lunch now because we think ultimately that will save more time at the end.

THE COURT: All right. Well, it's 11:26 now by the clock here on the bench. We're going to make this short. We'll reconvene about 10 minutes after 12:00.

The Court stands in recess.

(Lunch recess.)

THE COURT: Be seated, please.

All right, counsel. We'll continue the pretrial hearing in regard to Netlist versus Samsung.

Let's next turn to Document 202, which is Samsung's motion under *Daubert* and to strike certain testimony of Dr. Michael Brogioli.

Let me hear from the moving party Samsung.

DR. ALBERT: May I approach, Your Honor?

THE COURT: You may.

DR. ALBERT: May it please the Court. Frank Albert for Samsung. May I proceed?

THE COURT: You may proceed, counsel.

DR. ALBERT: Thank you, Your Honor.

We have four issues in our motion to strike materials from Doctor Brogioli's reports. The good news is we believe two of these issues have already been resolved.

So the first issue regarding reliance on excluded materials, we believe that the redactions provided by Netlist in response to the motion resolve just about everything. We think that there is one issue that counsel for Netlist and I discussed this morning that may also be resolved, and let me just skip right to that, Your Honor.

THE COURT: If there's anything left here that's not covered by the previously tendered redactions? Let's get it knocked out.

DR. ALBERT: Okay.

THE COURT: Preferably if there's not, we'll move on, but you need to tell me.

DR. ALBERT: So slide 27.

Counsel for Netlist and I discussed paragraphs -- excuse me. I believe this is the wrong one. I apologies, Your Honor.

Slide 39.

Your Honor, counsel and I discussed paragraph 479, the -- that part of this was not redacted in the original redactions that we received. The second part by late 2007 presenters have described this line of register device and data buffers and internal technical specifications, and then the next sentence as well. My understanding is that Netlist is no longer going to be offering this paragraph. If that's correct, then this portion of the motion is entirely resolved.

THE COURT: Mr. Sheasby, what's Plaintiff's position with regard to paragraph 479?

MR. SHEASBY: That's correct, Your Honor.

THE COURT: All right. Then this portion in the motion appears to have been addressed.

Let's move onto the second motion regarding claims 15 and 22 of the '060.

DR. ALBERT: And regarding claims 15 and 22, these are not elected claims, so we believe this portion of the motion is also resolved.

THE COURT: All right. Any dispute from Plaintiff on that?

MR. SHEASBY: No dispute, Your Honor.

THE COURT: All right. Then there will be no action taken on this section.

DR. ALBERT: The third issue, Your Honor, that unfortunately will take a little more time is the array die

issue. And we made a long discussion about that yesterday. We understand Your Honor's ruling that there are fact issues as to whether the accused HBM products contain an array die under the Court's construction. This issue is different. This issue is the role that Doctor Brogioli may play in framing that fact dispute at trial.

And we are requesting that he be precluded from providing any testimony that the accused products have an array die under the Court's construction because he failed to follow the construction, he applied something else, and what's he's applying is entirely inconsistent with the plain meaning of the term 'DRAM circuit'.

So what do we have in this construction? Array day was construed as array die that is different from a DRAM circuit. So in order to meet their burden, they need to show that what they're accusing as array die is different than a DRAM circuit. There's negative limitation there--it's different from. So in order to prove that it's an array die, they need to prove what is a DRAM circuit, and that we don't have that.

Doctor Brogioli has done none of that analysis. What he has done -- you heard that over and over again yesterday. What he has done is he has delved back into the prosecution history to reargue what the scope of the disclaimer was. He did not take this construction that the Court applied, array that is a different DRAM circuit, and then determine, Okay,

what is the plain meaning of DRAM circuit; how do I apply that.

And we, of course, understand the Court's guidance in claim construction process. The Court construes a term. It is always up to the parties at that point to take the Court's construction and apply the plain meaning.

So the scope of what was excluded from this construction, the DRAM circuit, of course is up to the experts. And we understood the Court's guidance to provide guidance on that that you take this construction, you apply the plain meaning, and you tell that to the jury. Instead, what we have here--we heard that over and over again yesterday--is that they took this construction and then they went back to the well to reargue what was the scope of the material that was disclaimed in prosecution.

Now, if we had a situation here where they said different than a DRAM circuit and this is what a DRAM circuit means, and I come in and say, You don't have that, Samsung, and, therefore, you don't have any DRAM circuits, therefore, it's an array die, that would be a different situation. But we don't have that, and you'll see that as we go through this.

So what are they actually doing? We heard over and over again yesterday that the issue is not being framed as whether we have DRAM circuits; the issue is whether we have DRAM circuits that are similar to Rajan's DRAM circuits. And here

I have Netlist's objections to the Court's claim construction order, which it ultimately withdrew, but here they are contrasting. They are saying that the prosecution statement argument at issue is limited to Rajan's DRAM circuits, 206A-D, and then--this is very important--contrasting that broadly and not to any DRAM circuits. So they are expressly telling the Court, to you, that Rajan's DRAM circuits are different than the plain and ordinary meaning of any DRAM circuits. And what they haven't done in this case is take that plain and ordinary meaning of DRAM circuits and apply it to the products at issue in this case.

How do we know that? Well, we asked the expert directly, "What is the plain and ordinary meaning of the term 'DRAM circuit'?"

"I don't recall discussing the plain and ordinary meaning of that phrase."

He does not recall discussing it because he doesn't discuss it. What he is discussing is Rajan and, specifically, his interpretation of Rajan. And you heard that interpretation over and over again yesterday. Rajan supposedly discussed -- it disclosed only wire-bonded DRAM circuits. So you have these DRAM circuits that are connected together via wire bonding, and that's what they say Rajan was all about. And so what they've done here in this case is not provide an opinion -- he's gone out of his way, out of his way

not to provide an opinion that there are no DRAM circuits; instead, it's there are no DRAM circuits like those in Rajan.

And he goes further. In his report he is going right back, right back to the prosecution history to argue that the intention of these statements that were made that ultimately resulted in disclaimer. As I understand the Court's practice, that is completely irrelevant at this point. When the Court has decided the scope of the disclaimer based on claim construction, we are to apply the Court's construction, the plain meaning. Going back and saying, Well, the Court was wrong, we actually discussed something different in the prosecution history, that to us seems out of bounds for an expert.

He goes on. Samsung argues that HBM products do not include array dies because the inventor supposedly argued to the examiner that DRAM circuits that -- are different from DRAM dies.

So, to be clear, that is not our contention. Our contention is these HBM products do not include DRAM -- array dies because they don't satisfy the Court's construction. In the Court's construction, array die is different than DRAM circuits. Ours have DRAM circuits. They are all about DRAM. That's a misstatement of what we're arguing. But then they go on to argue--you know, again delving back into this prosecution history--to argue the scope of what the disclaimer

was. That's improper. Samsung misquoted the inventor's statement to the office which was Rajan merely stacks DRAM circuits 206A-D which are different from array dies. So they're delving back into this claim construction process.

We had a lengthy claim construction exchange and a hearing. They had opportunity to present all of this. Judge Payne decided against them. They had opportunity to present objections. They did, and then they withdrew them. Now they're presenting those same arguments in the guise of expert testimony, which is completely improper.

I won't belabor the point, but I heard yesterday, Your Honor, that there was a dispute, that somehow Doctor Brogioli actually did provide a plain meaning analysis of DRAM circuits and array dies in his report for his infringement analysis. And the devil is in the details. The devil is absolutely in the details here. And so I would like to flip to that section of his report, one of these sections.

So here we have starting on paragraph 174, this limitation, "a plurality of stacked array dies, including a first group of array dies and a second group of at least one array die, each array die having data ports." You see here, first and foremost, array dies is a separate limitation from how array dies are connected, whether there's TSVs or not. Array dies stand alone--the plurality of stacked array dies. Okay? So what does Doctor Brogioli do?

THE COURT: What are you trying to do, counsel.

DR. ALBERT: So Doctor Brogioli goes through this on paragraph 174. No discussion of DRAM circuits; no discussion. No discussion of what DRAM circuits mean. No discussion that DRAM circuits are absent.

He gets all the way through 174, 177, all the way up to 183 before he starts even discussing DRAM circuits. And what you'll see is, if you go paragraph by paragraph, it's all the same. Every paragraph in this section, it's all the same. It's all, There is not DRAM circuits like those in Rajan. There's not a statement in here that says, Well, Samsung in these DRAM dies, which I agree are there, have, you know, DRAM cells and, you know, input/outputs and peripheral circuitry, he doesn't do any of the analysis to say, Well, I analyzed it and those aren't DRAM circuits. No. What he does is he says, here 183, "HBM products do not have DRAM circuits like those in Rajan." Like those in Rajan; again, referencing Rajan's DRAM circuits, the very next line.

And he does this -- he does this paragraph by paragraph where he carefully avoids discussing whether or not there is any DRAM circuits in these products. Instead, he always ties it to, Do you have the DRAM circuits of Rajan or the types of DRAM circuits that he's saying are in Rajan. Remember we heard yesterday Rajan they're saying discloses a type of DDR -- a particular type of DDR DRAM and a particular type of

connection scheme, the wire bonding.

So here, again, he says the stacked DRAM dies with long wire bonds--that's what he says is Rajan--are somehow different than Samsung's products. Samsung's has TSVs. Again, hasn't gone through DRAM circuits, hasn't mentioned what they are, hasn't said that there's no DRAM circuits, and certainly hasn't addressed their own admission that Rajan's DRAM circuits are different than any DRAM circuits. He does this over and over again. I've highlighted paragraph by paragraph.

185, next paragraph, he's talking about wire-bonded DRAM circuits. 186, instead of applying the DRAM circuits generally, he's then saying, Well, HBM doesn't have the very exact type of DRAM circuits that are in a different DRAM product. Again, he has not applied the Court's construction of DRAM circuits, negative limitation that he has to do, and he said, I did not discuss that in my report. He did not discuss that.

So we understand there are fact disputes, we understand the Court has ruled that there are fact disputes, but there are no factual issues that are relevant for the jury that should come from Doctor Brogioli because he does not, he does not apply the Court's construction.

And we heard -- again on 187, the DRAM dies in HBM are not interchangeable with another type. So he is --

Mr. Sheasby mentioned yesterday the Rajan is a wire-bonded DRAM -- this DDR type of DRAM. What they're doing is over and over again they're going back, going back to this Rajan analysis, this DDR1, this DDR2, this DDR3, the long wire bonding. It all relates to this Rajan analysis rather than an analysis under the Court's construction plain and ordinary meaning of DRAM circuits.

So what they're doing here -- doing here, Your Honor, is plainly what the Court prohibits. They're trying to redo claim construction, they're trying to redefine this term, instead of what the Court normally requires, which is you take the Court's construction and you apply the plain meaning of that term. Didn't apply it. Certainly doesn't do it in any of the paragraphs of his report. As we go on, 189, again referring to the --

THE COURT: Let me ask you this, Doctor Albert. You just said for the second time what you do is take the Court's construction and apply the plain meaning of that term. The Court didn't construe 'DRAM circuit', so what's your response to the argument that what this witness is doing is simply providing his version of the plain and ordinary meaning of 'DRAM circuit' that the Court didn't otherwise construe?

DR. ALBERT: And -- very good question, Your Honor. We struggled long and hard with this because that's what we have endeavored to do in all of our analysis. But that's not

what's happening here. He has gone out of his way, out of his way to avoid providing an opinion that there are no DRAM circuits in these devices.

THE COURT: And I assume you have an expert who's going to opine that DRAM circuits are not limited to what this witness says they are and the jury's going to decide what the plain and ordinary meaning of 'DRAM circuits' is. Is that not the case?

DR. ALBERT: It's not the case, Your Honor, because they haven't provided an opinion on what the plain and ordinary meaning is. So if they had come in and they had said --

THE COURT: Wait a minute. Are you telling me you don't have an expert that's going to take the opposite side of this issue as to what constitutes the plain and ordinary meaning of a DRAM circuit?

DR. ALBERT: We certainly have an expert that will take the position on what the plain and ordinary meaning of it.

THE COURT: Okay.

DR. ALBERT: I wouldn't say that it is the opposite issue, because they don't have the opposite issue. They don't have any testimony that this is the plain and ordinary meaning of 'DRAM circuits'. They -- all they do is they delve right back into the prosecution history to say, not what the plain

and ordinary meaning is, but what the disclaimer is. And if there is -- there are no statements in here that say there are no DRAM circuits anywhere. In fact, he goes in to tremendous detail about the prosecution history regarding the -- what the scope of the disclaimer was.

So if it was the case where he said -- Doctor Brogioli said there is a plain meaning of this term to one of ordinary skill in the art, that's something we can cross examine on, but he doesn't have that. He doesn't have an opinion on what -- the plain and ordinary meaning of this term, and, therefore, he doesn't have an opinion on whether the HBM products satisfy this limitation under the plain and ordinary meaning. In fact, he says in his report -- or says in his depo, I have not provided an analysis of the plain and ordinary meaning of this term.

So if he's allowed to testify about this issue, he's going to come in, he's going to say, you know, whatever he's going to say based on his report, but it's not going to be plain meaning. Those words do not come from his report. Those words do not come from his deposition. What he's going to say is this bold conclusory statement that array dies are in the products without doing any of the analysis that there are no DRAM circuits under the Court's construction. And that's inherently prejudicial that the Court would allow an expert to come in and make those bold assertions with no

opinion, with no opinion.

And we gave them an opportunity. It's not like he just, you know, forgot to say the words. We gave them an opportunity in his deposition, and he flatly said that he doesn't have it.

THE COURT: All right. What else do we need to hear from you on in this motion before I hear from the other side?

DR. ALBERT: Moving on to unsupported technical benefits opinion, Your Honor, this is a different issue than what you heard with Mangione-Smith. Mangione-Smith, he provided testimony that there were certain technical benefits, which we obviously disagree with, and based on the Court's ruling we'll cross examine him on them. Doctor Brogioli, he provides a next step, and what he says is that without the asserted patents, that Samsung would not be technically competitive in the marketplace for certain products. And Mr. Kennedy relies on that and he says those products with stack height of eight or more would not have been possible without the asserted claims.

But doctor Brogioli's testimony, his statement is Samsung would not have been technically competitive. Does he say that we can't make a different product at eight stack that would perform slower, different, more -- you know, more power? You know, certainly he'd say stuff that like, but the conclusion that Samsung would not be technically competitive is

unsupported, and he's not an expert in that area.

So what would that mean? Well, you know, first off, he's not an economist. He's not going out and doing any market studies. He says he's not an economist. He did not perform any market analysis. So if you're going to make a bold statement that, Hey, you couldn't sell this product in the market, generally you would want some sort of data on that. You would want, you know, to survey some customers, survey the industry, what -- and you would want to ask them, Would this be something that you could sell.

And that gets into the second issue, is he doesn't do a -- you know, what we call a kind of product level analysis. And so what I mean by that, Your Honor, is he analyzes very specific features and says, This would be some effect on the product. And he says that, you know, without the claims that you could do something else, but it's going to have some effect.

What he hasn't done is he hasn't said, Okay, this other hypothetical product that you could make would have these sorts of behaviors. You know, this would be the speed, combined with the capacity, combined with the power, combined with the temperature load, combined with the cost, combined with the price that a company in this field could charge. Didn't do any of that. And so because he didn't do any of that, he couldn't input any of that into a kind of market

survey.

To say that you wouldn't be able to sell these, period, without some alternative begs the question what are the -- what is the complete set of performance parameters for that alternative, which he didn't do. And then it begs the question, well, would a customer buy it for 95 percent of the price, 90 percent of the price, 80 percent of the price? None of that is in his report. He didn't do any of that analysis. But what he says is, Well, it absolutely wouldn't be technically competitive. And then Kennedy says, Well, because it wouldn't be technically competitive, I get all of your revenue. Not 95, not 80, not 70, not whatever a customer would have said they would have paid for this alternative, but he gets all of the differential revenue between that product and what he says is non-infringing.

And because he hasn't done any of the market analysis, because he hasn't actually shown what this other product would actually do, that type of conclusion is completely inappropriate and unsupported and should be struck as unscientific.

THE COURT: All right.

DR. ALBERT: Thank you, Your Honor.

THE COURT: Thank you, counsel.

Let me hear from Netlist in response, please.

MR. SHEASBY: May I approach, Your Honor?

THE COURT: You may.

MR. SHEASBY: Mr. Huynh, we are on deck 16.

THE COURT: All right, Mr. Sheasby. Please proceed.

MR. SHEASBY: In paragraph 80, doctor Brogioli speaks about the fact that he interprets the plain and ordinary meaning of 'DRAM circuit' as meaning a wire-bonded die with edge binding. That is an issue that will be joined at trial by Doctor Robins and is appropriately joined at trial by Doctor Robins. He literally says the words plain and ordinary meaning in his report in paragraph 80.

If we go to slide 29.

He in these paragraphs speaks about the fact that the core dies that Samsung uses meets his understanding of an array die. They are memory cells organized in an array format. He goes on to particularly point out that they are different from DRAM circuits. And the way he knows they're different from DRAM circuits is because Samsung's documents describe them as different from circuits and Samsung witnesses describe them as different from circuits. An example of this --

THE COURT: Let's go back just a minute, Mr. Sheasby. You said that Doctor Brogioli says in paragraph 80 of his report that this is the plain and ordinary meaning of DRAM?

MR. SHEASBY: DRAM circuit.

THE COURT: Show me that, please.

All right.

MR. SHEASBY: He also says it in paragraph 190, Your Honor. He says in paragraph 190, "Unlike the 3D stacked HBM cores, a person of ordinary skill in the art would recognize that the wire-bonded dies, for example, Rajan, are adhered together using insulating adhesives." He goes on to discuss that at lengths. And so there is an issue that's fairly joined which is what is the meaning of DRAM circuit.

He doesn't just stop there. In these paragraphs he's saying DRAM dies in HBMs are not DRAM circuits. And if this -- you can't refer to Rajan was actually the methodology or the standard that would be used, Your Honor; it would be not -- it would still -- his methodology would still be sound because there's no reference to DRAM or Rajan in any of these paragraphs. He's discussing a DRAM circuit and applying it in determining that the HBM products are not DRAM circuits.

He does it here again in paragraph 191. He talks about the fact that HBM stacks do not exist in traditional DRAMs, which he refers to as a DRAM circuit. He points to the testimony of the experts at Samsung that they are not DRAM -- that what HBMs have is not DRAM circuits.

And to be clear, it's not just I define Rajan as -- I define it as limited to Rajan 260A-D, and, therefore, it's not a DRAM circuit; 260A-D of Rajan depict a stricture. That

structure he defines as edge connected DRAM circuits. His point is, is that those circuits do not exist in the products at issue in this case. He cites to Jihwan Kim without any reference to -- he cites to Doctor Kim without any reference to Rajan. He cites to Doctor McAlexander who also admits -- Doctor McAlexander also admits that there is a distinction between DRAM circuits and which on HBM.

THE COURT: Let me stop you for a minute. At a high level, what it seems like to me the essence of your response is, and will be, is that Defendant is seeking to hamstring your expert in opining on the plain and ordinary meaning of 'DRAM circuits' because the Court didn't expressly construe that term at claim construction, and they want to basically tie your hands on how you can opine to the jury through this witness what the plain and ordinary meaning of 'DRAM circuits' is. And I'm very suspicious, having looked at the briefing, that when we get to Doctor Robins that we're going to hear the same story in reverse, where they're going to tell me that you're trying to hamstring their expert as to his opinions of the plain and ordinary meaning of 'DRAM circuits' based on whatever arguments you presented in that motion.

I mean, is this not a tit-for-tat situation where these two experts are simply going to tell the jury what they believe the plain and ordinary meaning of what a term is that was not construed at claim construction, and the jury is going

to decide which of these battle of the experts they're going to be persuaded by? Because if that's what this is all about, we can all save ourselves a whole lot of time without going through the weeds over and over again about all these esoteric reasons why this is improper claim construction when, at the end of the day, each side's trying to keep the other guy's expert from saying something that hurts them so that their expert can say what helps them. I mean, I've seen this movie before.

MR. SHEASBY: This is the movie. With one exception this is the movie. And so there is one thing that I -- that we do think is out of bounds.

THE COURT: And in this movie there are no good guys; both of the actors are bad guys trying to do the same thing to each other.

MR. SHEASBY: I think it's -- yes.

THE COURT: Well, you know, I've got the rest of the day and we're going to get done everything we can get done in the rest of the day, but this seems like a colossal waste of time to me.

MR. SHEASBY: I'll move on, Your Honor.

THE COURT: And I'm not speaking solely to Mr. Sheasby; I'm speaking to everybody in the room.

MR. SHEASBY: I understand, Your Honor.

So the last issue is technical valuation opinions.

If we go to slide 53.

Doctor Brogioli does say that he's -- that they would not be able to achieve eight stack or higher with the required speed thermal budget -- and the required speed in thermal budget that you needed. And he didn't just make that up. If you go to slide 56, he bases that on Samsung's modeling as to how much additional capacitive load you add every time you go above four. He specifically discusses it at 340.

Again, he cites to simulation data that Samsung created itself about what would -- the thermal budget would be if you went above four using the prior art design. He discusses -- discussion on what the power budget would be if you went above four with this decision. And then he discusses marking data and how they have to stay within a thermal and power envelope for them to compete with each of the -- each of their competitors in the HBM space.

So I agree with Doctor Albert that if he would have said they can't go higher than four high but I'm not going to tell you why, that would be inappropriate; but I think, fair to say, there are some pretty rigorous analysis here, and I will stand on that, Your Honor.

THE COURT: All right. Anything further?

DR. ALBERT: Briefly, Your Honor.

THE COURT: I've heard that before.

DR. ALBERT: We take your Court's guidance, and this

is exactly why we asked Doctor Brogioli the question, Did you discuss plain and ordinary meaning of the term 'DRAM circuit'? He said he did not do that.

So this is not a case where one expert is saying plain and ordinary meaning and the other expert is saying, Well, different plain and ordinary meaning. We have one expert that says nothing, and our expert that says under the plain and ordinary meaning this term is not satisfied.

THE COURT: All right.

DR. ALBERT: Your Honor asked to see paragraph 80. I don't see plain and ordinary meaning at all referenced in paragraph 80. To say that this paragraph 80 shows that Doctor Brogioli is discussing the plain and ordinary meaning of DRAM circuits is -- we just don't see it. He says a person of ordinary skill in the art with regard to stacked DRAM circuits with edge wiring bonding such as those found in Rajan. He doesn't even say DRAM circuits. It's stacked DRAM circuits. And he says that those are different from array dies stacked via TSVs; nothing about the plain and ordinary meaning of the term DRAM circuits.

And we also were directed to paragraph 190 after Mr. Sheasby was forced to show paragraph 80 to the Court, and it's the same, Your Honor. In fact, it's one of the paragraphs that I showed Your Honor. There's nothing in here about plain and ordinary meaning. Again, it's actually

referencing the DRAM circuits of Rajan.

So with that, Your Honor, we do not believe this is a case of one expert saying plain meaning and the other expert saying not plain meaning; this is a case of their expert saying he didn't apply the plain meaning but their attorneys saying that he did.

THE COURT: Thank you, counsel.

I'm going to carry ruling on this motion regarding Doctor Brogioli, Document 202, and we're next going to take up Document 210, Plaintiff's motion to strike portions of the rebuttal expert report of Gabriel Robins. And then after I've heard both motions, I'll give you a ruling.

MR. SHEASBY: Thank you, Your Honor.

I'll go through, based on the index in our original motion, section by section.

So the first issue is improper claim construction arguments, and the first issue there is 'electrical communication'. Doctor Robins says that 'electrical communication' means a lack of electrical connection. There's nothing wrong with that. What is wrong is what he does with that information is the way he gets there is by going through the specification and finding examples of how the specification supports his interpretation of 'electrical connection', and in our mind, that's *ultra vires*. That's not appropriate use of the specification. If they wanted a

construction of 'electrical communication', they should have sought it from the Court. They shouldn't be advancing the specification as setting forth the limits of it.

So this part is fine. This part talking about what the goal of the patent and referring to incorporating by reference a broad understanding of 'electrical communication' is also suggested by the specifications referenced to load-on die, all of that is *ultra vires*, it's not plain and ordinary meaning, and it's not using his knowledge. If they wanted that construction, it should have been at the time.

Die interconnects. It's the same thing. He's using extrinsic evidence to argue that die interconnects are broader than TSVs. And he speaks about how there was an intentional lack of definition in certain documents. That's, once again, using the extrinsic evidence to construe the claims. The first and second driver -- this is a pure matter of claim construction, Your Honor, which is the claim requires two drivers, a first driver of a first size, a second driver of a second size, and Doctor Robins says that both of the drivers have to be different sizes. That is a pure issue of claim construction. That is not any issue that would be appropriately before an expert.

There's no -- the Federal Circuit law is quite clear that the recitation to a first and second must be distinct structures, but there's no requirement that the recitation to

first and second be different sizes or different in any ways. It's just requiring two separate things--a first driver that has a driver size, a second driver that has a driver size. And dependent claims, leaning off of that, there could be additional limitations about those driver sizes, but there's none in claim 15.

The next issue is his discussion of 'selection in consideration of a load'. The problem with that is, once again, that he's focusing on the plain meaning specification -- he's referring to the specification as opposed to giving his plain and ordinary meaning.

The last issue is 'chip select signal' where he's actually giving a contradictory interpretation of 'chip select signal' than the Court gave. The Court gave a plain and ordinary meaning and rejected certain portions of that, and it rejected a limiting term, and then Doctor Robins is attempting to put the limiting term back in.

The next -- Your Honor, I'll stand on that, if you have anymore questions.

THE COURT: I don't.

MR. SHEASBY: Okay. The next issue is his interpretation of why the first driver size and the second driver size are not met, and he talks about the fact that the specification doesn't mention the word 'test' or the fact that the specification's embodiments don't use a driver size for

testing, which is what Samsung does; that they don't infringe. And that is improper to use the specification to narrow the terms in the -- to narrow the meaning of the terms. You can't say you don't infringe because your embodiment is different from the specification embodiment. It's not just once; it's repeated again and again discussing certain embodiments. And so his comparison of the embodiments to the claims for purposes of infringement is improper.

There's an odd thing that is happening here, and I think it's the same thing that we talked about briefly, is Doctor Robins speaks extensively in his infringement report about prior art--in effect, comparing the asserted claims--sorry--comparing the accused products to prior art. This is paragraphs 52 to 63, 67, paragraphs 113 to 130, talking about it was well-known, well-developed, well-studied; suggesting that because Netlist -- Samsung may practice elements of the prior art, that that doesn't infringe. Whatever relevance that have for the purpose of invalidity, it is no relevance for the purposes of infringement.

Next issue, previously undisclosed prior art. This is the same issue that was talked about with Doctor McAlexander. Samsung admits that Doctor Robins does not offer any non-infringing alternatives, so Doctor Meyer should not be able to rely on Doctor -- Mr. Meyer should not be able to rely on Doctor Robins' claim construction.

The next issue is the one we discussed previously, which is how do we handle this issue of the -- I'm sorry. I think I'm skipping one. Oh, no. How do we handle this issue of the Rajan circuits. There is one issue that I think stands out differently and is difference in kind, Your Honor, from all the others, and this has to do with paragraphs beginning at 155 of Doctor Robins' report, and it's the only part that I'll focus on.

What Doctor Robins does in this section of his report is he talks about how Rajan can be read in concert with another reference called Elpida, and if you combine Elpida and Rajan, you could then interpret Rajan to be disclosing TSVs and interconnects. And it's the one thing -- this is the only thing that I'm going to spend time on in Rajan on the die interconnect, which is that it is not proper for him to do a hypothetical scenario in which he said, Well, if Rajan was combined with this Elpida patent, this would allow a person of ordinary skill that it was obvious from Rajan to do X. That's the only issue I'll focus on is this analysis on page 155 where he uses this Matsui reference and combines it with Rajan to sort of present a non-infringement position. Other than that, the issues are just he said, she said; they're right back at each other. But this is a different in kind. This is dangerous and it's actually violating the Federal Circuit case law.

And if I can, let's go to deck 6, please.

So the danger of what they are proposing to do is talked about in *Genuine Enabling*. So in *Genuine Enabling* there was a claim limitation, there was disclaimer, and it says, "signals containing audio or higher frequencies are excluded." And then what the district court did is the district court said, Well -- which is actually, oddly enough, the opposite of what Judge Payne did. If Judge Payne would have construed 'DRAM circuit', it probably would have been reversible error. But the district court in this sense says, Well, I'm going to construe slow varying as anything less than 500 megahertz, and there was nothing in the intrinsic record that says it was less than 500 megahertz.

What they did is they put in a declaration from a Doctor Chizeck who cited another reference, and that other reference said, Oh, well, slow varying would be 500. And so then he said -- they said, Well, that's our basis for limiting it to 500. So what they did is they didn't take the disclaimer as is, they took the disclaimer, then they combined it with another reference to limit the claims.

That's the only issue that I think is different in kind from what we've been dealing with on both Robins and --

Let's go back to deck 13, Mr. Huynh.

That's the only issue that is different in kind. This is a dangerous paragraph, and this does have very significant

issues with Federal Circuit precedent, in particular the *Genuine Enabling* case that I just read.

And if you go to slide 41.

"To the extent Mr. Nguyen's statements may implicate other scope, the record did not rise to the level establishing a clear and unmistakable disclaimer." So what Robins does is he says, I'm going to take Rajan, but if you combine Rajan with Matsui, it would be obvious to do X. That's a bridge too far and not appropriate.

The last issue I'll discuss, Your Honor, has to do with references to Netlist's pleadings. And I take your -- the Court's admonition seriously. Us wasting time on taking completely contradictory positions doesn't serve anyone. But this is different. This Matsui tactic is -- rises to a different level.

And the last issue is he -- Doctor Robins cites to a number of our pleadings. Our pleadings are not appropriately in evidence. No one's going to dispute that the HBM products have DRAM dies on them, and so there would be no reason to put our pleadings or infringement contentions in front of a jury.

If there are no other questions, I'll stand down, Your Honor.

THE COURT: All right already, counsel. Thank you.

Let me hear a response from Samsung, please.

MR. COLVIN: May I approach, Your Honor?

THE COURT: You may.

MR. COLVIN: Matt Colvin for Samsung.

THE COURT: Please proceed.

MR. COLVIN: Your Honor, I will start with the discussion of die interconnects not an electrical communication with array dies. I think it's important to remember, and I'm not going to belabor the point, but a person of ordinary skill in the art understands the claim terms in light of the specification and in light of all of the other knowledge that a person of ordinary skill in the art has.

And so much of what is being complained about by Netlist here is that Doctor Robins shows his work. He has an extensive background section where he shows what a person of ordinary skill in the art would know and why he would know it. Here are all the references that show X, Y, and Z, and, therefore, a person of ordinary skill in the art would understand those things.

Then he looks at the claim term and he says, Well, what does the specification have to say about this, because that's what the Federal Circuit tells us a person of ordinary skill in the art is supposed to do.

So looking at paragraph 416, which is the first paragraph -- well, it is one of the paragraphs that Netlist complains of, we see doc Robins is doing just that. Doctor Robins says, as I explained in that section, which is his background

section, "a person of ordinary skill in the art would understand electrical communication is a broad concept that involves the transmission of electricity." Now, Netlist may say that statement alone is okay, but they have a problem that he shows his work and says why a person of ordinary skill in the art would do that. That's the prior sections that are in the background that Netlist would have this Court strike. But Doctor Robins is doing exactly what we want our experts to do, and that's show his work, show why a person of ordinary skill in the art would understand 'electrical communication' to be just what he says it is here.

Paragraphs 417 and 418 and 419 tell a similar story. Here's 418. This is where he Doctor Robins confirms that his understanding is consistent with the specification. Again, this is what we want our experts to do.

I have some slides here about what Doctor Brogioli does, which I know the Court has already heard, but it's important to note that Doctor Brogioli goes even further. So to the extent what Doctor Brogioli does is okay, Doctor Robins is certainly okay, because Doctor Brogioli takes the specification and he says, Well, based on the specification, which I have highlighted, each array die has receivers and drivers that enable electrical communication.

And we asked him about that at his deposition. And just to clarify, we asked him if it was his opinion that in order

to infringe, the die interconnects must be an electrical communication with the receivers for the drivers on the core die. He says yes. The problem with that is that that's not in the claim. And so Doctor Brogioli takes what was in the specification and narrows it, narrows the claim scope by adding what I have in red on this slide, which is far more egregious than anything that Doctor Robins does.

There were several different arguments about driver size, but the arguments about driver size are not claim construction. The arguments about driver size come right out of the claims themselves. Claim 1 of the '160 Patent requires "the second driver size is different from the first driver size." This isn't Doctor Robins looking in the specification to pull out some non-infringement argument; this argument is directly in the claim itself.

With respect to the '060 Patent, here's claim 20, "the first driver and the second driver drive the data signal." Again, this isn't Doctor Robins looking into the specification and determining that the signal being sent is actually a data signal and not a test signal; no, it says right in the claim itself that the signal that's being sent is a data signal.

One of the claims that Mr. Sheasby showed you was claim 15 of the '060 Patent, and his argument in claim 15 is not an asserted claim.

I think the technical benefits section deserves a bit of

a closer look. So you saw this paragraph from Doctor Brogioli's report, Your Honor, where Doctor Brogioli makes this large claim that Samsung would not be technically competitive in applications that require high capacity memory, i.e., those that utilize eight high stacks and higher.

And to do that, Doctor Brogioli -- well, Doctor Brogioli's technical analysis does something that's relevant to this discussion of Robins, so I need to spend a moment on it.

Doctor Brogioli takes a particular feature of the accused product, which is a twist or spiral connection between the die, and he says, That feature is what going to use to assess the technical benefits of the '060 and the '160 Patents. What's super interesting about that is that the twist and spiral connection is not something that's found in the '060 and '160 Patents. That feature is found in the accused products. But, nevertheless, Doctor Brogioli says, I'm going to take that feature and I'm going to compare the attributes of that feature to a feature called multi-drop. And whatever the benefits are from twist and spiral, Doctor Brogioli says, Well, those are the benefits of the asserted patents.

And so Doctor Brogioli bases his entire technical benefits analysis on this fact that the '060 and the '160 Patents, the benefits can be shown by twist and spiral. So that's the baseline we're starting from.

What Doctor Robins does is says, I disagree. First of all, you can do multi-drop with eight high or higher, and he shows his work. For example, in paragraph 477, he says, Various multi-drop connection schemes can indeed work, and he cites to this Rio and Elpida art. These things teach stacking above the 32 DRAM dies using TSVs arranged in multi-drop. He shows his work.

Second, Doctor Brogioli states that the industry requires eight high capability. Doctor Robins disagrees with that. He says, You can do four highs next to each other. He cites this Cane patent, and again he shows his work.

To rebut Doctor Brogioli's assertion that twist and spiral can stand in as the technical benefit solution of the '060 and '160 Patents, he cites several other patents that shows that twist and spiral came before the '060 and the '160. So it can't be the case that the '060 and '160 Patents' technical benefits can be shown by looking at twist and spiral because twist and spiral came before.

Here's an example--the Shibata patent. Doctor Robins discusses another set of patents that disclose the same thing--Elpida from 2009 and 2011. These all show that point to point and twist and spiral came before, and so this is a direct rebuttal to what Doctor Brogioli has to say of the technical benefits of the '060 and '160.

With respect to array die, just a few things I want to

point out there, Your Honor. Netlist in its briefing says this statement: "Doctor Robins does not apply the Court's ruling that the disclaimer was a structural disclaimer that was limited to Rajan's DRAM circuits." Your Honor, I think they slipped up when they said that. I think what they did there was reveal what they're really thinking that this -- what they're really trying to do here, because the Court's ruling was not limited to Rajan's DRAM circuits; it was just DRAM circuits. And they take issue with Doctor Robins not looking at DRAM circuits as though they're Rajan's DRAM circuits, but there's nothing to take issue with there because that wasn't the Court's construction.

Doctor Robins analyzes DRAM circuits -- array dies, and, therefore, DRAM circuits, in a few different ways. The first way he does it is he says, The core dies used in the HBM products are not different from DRAM circuits because they are DRAM memory dies with DRAM circuits. Netlist criticizes this and says that Doctor Robins doesn't have any support for that, but paragraphs 378 through 409 support it.

And when you go and look at footnote 10 where he says DRAM circuits -- DRAM memory are known to have DRAM circuits, and you go and look at paragraphs 80 and 89 of his report, just for example, Doctor Robins has a huge discussion of DRAM and DRAM circuits.

On the left is a figure, and Doctor Robins says that

this shows a DRAM circuit. On the right we have a DRAM cell structure. Doctor Robins goes on. He cites JEDEC, a JEDEC spec entitled "Terms, definitions and letter symbols for microcomputers, microprocessors, and memory integrated circuits." And it talks about RAM and it talks about DRAM. And that is the type of analysis that's missing from Doctor Brogioli's report.

Doctor Brogioli starts from the premise of prosecution history estoppel, and then says, Based on this prosecution history estoppel, he construes it to be broad -- I mean, he construes it to be very narrow. And he looks at Rajan and he says, Well, based on Rajan, DRAM circuits must be this special thing. But that's the wrong thing to look at it, Your Honor. And it's not apples to apples.

THE COURT: We've had the motion on Doctor Brogioli. I've heard those arguments. If this is all about, If you let them do theirs you have to let us do ours, I understand the subtext here, counsel. So let's talk about the motion that's before the Court.

MR. COLVIN: My point would be simply that it's different; that you can deny what they did with Doctor Brogioli and allow what Doctor Robins did because they're different.

THE COURT: I understand I have that option.

MR. COLVIN: Finally, Your Honor, with respect to

the -- Doctor Robins' reliance on infringement contentions and statements made in the complaint, those are party admissions. And so it's going to be an issue in this case whether or not the HBM products had DRAM circuits, and Netlist has admitted that the die in the HBM products are DRAM. Those are one of the things that Doctor Robins looks at. And because they're party admissions, we think he should be allowed to do that.

THE COURT: Anything further?

MR. COLVIN: That's all, Your Honor.

THE COURT: Anything further from Netlist on this motion?

MR. SHEASBY: Two things.

Slide 33.

Counsel does not dispute that they conceded that Doctor Robins does not offer non-infringing alternatives. His discussion of random parts of the Samsung product that look like the prior art has absolutely no relevance to anything he would be opining because they admit it is not a non-infringing alternative. It would be a violation of Federal Circuit law to get him to stand up and talk about all these random prior art references in his NIAs when they admitted he did not do an NIA analysis. It's plain as day in front of Your Honor, and I didn't hear anything from counsel in response.

And the second issue is the only point that I focused on on the DRAM circuit issue was the point that counsel ignored,

and counsel is very able and I don't believe that was unintentional, which is this combination of Matsui, the use of Matsui to inform Rajan is not proper.

And slide 38.

There was no response to this discussion at paragraph 155 about using Matsui. That is different in kind from what everyone else did. I heard no argument from Defendant on that point.

Thank you, Your Honor.

MR. COLVIN: Your Honor, may I address the Matsui point?

THE COURT: I've heard all the argument I need to hear, counsel.

With regard to Document 202, Samsung's motion to strike pursuant to *Daubert* the Brogioli report, that motion is denied.

With regard to Document 210, Netlist's motion to strike the Robins rebuttal report, that motion is denied.

We're going to take a short recess, counsel. We still have the motions to strike on McAlexander's opening report and Robins' opening report tied to the validity issues, and I'll take those up when I come back.

With that we stand in recess.

(Brief recess.)

THE COURT: Be seated, please.

All right, counsel. Let's next take up Plaintiff's motion to strike portions of the opening report of Doctor McAlexander, Document 212.

Let me hear from Netlist on this, please.

MR. SHEASBY: Your Honor, I'll go section by section based on the table of contents, with your permission. I will not repeat arguments. I will focus on where I think is the most important elements of the motion.

THE COURT: All right. Let's go.

MR. SHEASBY: And proper claim construction opinion under the guise of written description.

If we could have slide 19 in deck 12, Mr. Huynh.

As the slide's being pulled up, this relates to Mr. McAlexander's discussion of the 'drive' term. It's, in particular, in 237 of the -- of his expert reports. And what he does is he quotes --

If you go to slide 20.

-- what he does is he sort of glosses on Judge Payne's construction of the 'drive' term, and then uses his reasoning to inform his claim interpretation. And it's -- per se, it's not that he can't have it in his expert report; he should just not be discussing any of the reasoning of Judge Payne in his testimony to the jurors.

The other major point that I'd like to discuss is the construction of 'memory module'. So in his -- he, in effect,

construes 'memory module' as the presence of large amounts of DRAM and -- the large amounts of DRAM and flash on the same module, and he does two things. He discusses -- he does two things. One is he discusses what he describes as the intent of the inventors based on other patents -- so this is paragraph 292 and 2088 are the most two dangerous examples. He talks about the clear intent of the inventors as to what their invention on -- based on other patents, and that's a violation of Federal Circuit precedent. That's *Vas-Cath versus Mahurkar*. In my mind, those are the two major issues that I'd like the Court to direct the Court's attention to in this section.

THE COURT: All right.

MR. SHEASBY: Is there -- if there's no other specifics the Court has on that section, I'll move onto the equitable defenses issue.

THE COURT: That's fine.

MR. SHEASBY: Equitable defenses have no place in this Court.

If you go to slide -- sorry.

No place in front of the jury.

If you go to slide 40.

He discusses a number of issues relating to enforceability at a number of paragraphs throughout his report. So beginning at paragraph 8, he does it again

beginning at page 130 and 142, and unclean hands at page 1118 and -- paragraph 1188, and again at paragraph 2099 he discusses standards based estoppel. None of those issues are appropriately in front of the jury, and anything having to do with estoppel or having to do with laches or disclosure to the JEDEC being in violation is solely for the Court, not for the jury.

The third issue is the relationship between JEDEC and Netlist.

Slide 46.

He engages in this practice of accusing Netlist of violating the standards of JEDEC by apparently alleging that we drafted claims that are associated with JEDEC. First off, there's a couple of issues with this. One, Mr. McAlexander has no experience in JEDEC. He's not an expert in JEDEC. There's a different expert in JEDEC. It's Mr. Halbert.

Second, his discussion of Netlist strategy has no place in his report. He doesn't know what Netlist strategy was. He's a technical expert.

And the third issue is they admitted that these facts go to prosecution laches, which is not before the Court. That we -- this sort of subtext that we did something improper or we acted inappropriately in front of JEDEC has no role in front of the jury.

If Your Honor has no further questions, I'll proceed.

THE COURT: The question I would have for you at this juncture, Mr. Sheasby, is are there issues that are appropriate for submission to the jury that the underlying facts related to these equitable defenses relate to? In other words, if there's a factual basis for an asserted equitable defense and that factual basis relates to a proper jury issue, then while arguing for the equitable relief, it would be improper, it would also be improper to strike the underlying facts that support something else that is a proper jury issue.

MR. SHEASBY: Yeah. So I think there's two issues. We can only go by what's in the report, and the report discusses it for the purpose of prosecution laches, and that's what Samsung's opposition says--it's for the purpose of laches and inequitable conduct. So they haven't identified any factual reason why it would be relevant.

The second issue is because the Court had concluded yesterday that no one's making a standard essentiality claim for these patents, any discussion of behavior in JEDEC and patenting the technology that was being developed by JEDEC is just an ad hominem because he agrees that JEDEC -- the patents are not standard essential, and our expert agrees that the patents are not standard essential as well.

So based on that, I don't see any -- listen. Is it possible to come up with some theoretical way in which the jury could care about the way JEDEC is -- operates and

Netlist's role in JEDEC? Yeah. If you stretch things far enough, they can say, Oh, it relates to damages, it relates to infringement. But that is so tenuous in comparison to the toxicity of talking to us about copying from JEDEC based on meetings we attended.

There is actually no validity allegation that we copied from JEDEC left in the case. There's no non-infringement allegation based on JEDEC in the case because these patents are standard essential. So a clever lawyer will say there's something that this relates to, but any relationship beyond laches and inequitable conduct is so attenuated that I think on the balance of -- the balancing test it should be stricken.

THE COURT: The balance of the equities with regard to the equitable defenses?

MR. SHEASBY: Yeah.

THE COURT: Okay. Let's go on to Issue 3.

MR. SHEASBY: This is the alleged printed publications issue. This issue has been resolved by the Court's motion for summary judgment motion yesterday, and I don't believe there is anything to rule on at this point because the motion for summary judgment would constrain what he's able to do in that section as well.

Other litigation. This is consistent with what the Court previously ruled on.

If you go to slide 66.

He has this extensive discussion of ITC complaints involving patents that are not the patents at issue in this case and I think, consistent with the Court's ruling, that should be stricken as well.

And then the last issue is an issue that's I think purely a legal one.

If you go to slide 73.

This is the undisclosed prior art defenses -- undisclosed invalidity defenses. The question is, in your invalidity contentions can you essentially say there's a lack of written description for every single word in the patent, which is effectively what they've done, and then at a later point in time pull out one word and use that as the basis for claiming there's no written description. And it's just -- it's going to be a policy issue for the Court to decide whether just allowing people to say everything lacks a description is a sufficient disclosure.

And then the other issue that's just a purely legal issue is that this was never disclosed. The combination of FB-DIMM with these patents in Elpida was never disclosed in the invalidity contentions. This is at paragraph 165. And it should be stricken.

THE COURT: All right. Thank you, counsel.

MR. SHEASBY: Thank you, Your Honor.

THE COURT: Let me hear responsive argument from

Samsung.

MR. LIVEDALEN: May I approach, Your Honor?

THE COURT: You may.

All right counsel. Proceed when you're ready.

MR. LIVEDALEN: Thank you, Your Honor. Brian Livedalen for Samsung.

So I want to start off with the written description issue. It appears, after the briefing, that Netlist has significantly pulled back its original request. If you recall, the original motion essentially said that Mr. McAlexander was arguing claim construction because he was adopting the claim scope that Netlist was providing and asserting its infringement case against the accused products.

As we see here on slide 4, there is absolutely nothing wrong with what Mr. McAlexander did. As we see in the *01 Communique Lab versus Citrix* case, "If the plaintiff attempts to expand the claims for infringement, the defendant is able to argue that the claim is, therefore, invalid." That's exactly what Mr. McAlexander is doing. So he's taking Mr. Doctor Mangione-Smith's assertions on infringement and assuming if that's a proper scope of the claim, then the claim is, therefore, invalid. That's proper under Federal Circuit law.

There's another case from the Federal Circuit that is the exact same thing. They ask the question--if the claim covers

a particular accused product, which here was the Solofill accused products, is written description supported; and the answer was no.

So, again, this is not an issue of rearguing claim construction; it's simply a question of whether the claims with the breadth as alleged by Netlist and its expert, are those claims supported by the written description of the patent.

THE COURT: Counsel, you said that it appeared that the Plaintiff had narrowed their complaints with regard to this section of the motion.

MR. LIVEDALEN: Yes. Only --

THE COURT: I heard counsel for Plaintiff say they were going to argue what they thought was the most pressing, but I didn't hear you say he was giving up or surrendering on any of what's in the briefing, so --

MR. LIVEDALEN: Well, let me rephrase that, then. We in our opposition we identified the *Rivera* case, identified the *01 Communique* case, and in response in their reply Netlist did not refute those cases, did not find a different case that said what Mr. McAlexander was doing. And Mr. Sheasby here discussed two terms, the 'drive' term and the 'memory module term', and for neither of those terms did he say that there was an improper claim construction because Mr. McAlexander was using the scope defined by Netlist for infringement.

Does that answer your question, Your Honor?

THE COURT: I guess it does.

I don't mean to be picky about this, but with my last name I'm very sensitive of how it's been mispronounced over the years. And the last two days I've heard Cheeseby, I've heard Sheezby. I've heard Sheasby.

You know, counsel, tell us how you want everybody to pronounce your last name, and I'll make sure they do it the right way every time.

MR. SHEASBY: Jason Sheasby is how the name is said.

THE COURT: That's what I thought.

So no more Mr. Cheeseby, no more Mr. Sheezby; call the man what he wants to be called just as you would like your name to be pronounced properly. Okay?

.

MR. LIVEDALEN: Thank you, Your Honor. Appreciate the clarification.

THE COURT: All right. Let's move on.

MR. LIVEDALEN: So with respect to the 'drive' term, we understand that neither expert should be discussing the opinions from Judge Payne, Your Honor. So to the extent that both experts are going to be discussing the rationale or analysis that got to Judge Payne's construction with respect to the 'drive' term, we're fine with that. We understand -- I believe that's subject to the Court's standing -- in one of

the standing MILs.

THE COURT: There's an order in every claim construction opinion that goes out of this district, at least this division, that says you're not to refer to anything other than the actual constructions before the jury or during a trial.

MR. LIVEDALEN: That's right, Your Honor.

THE COURT: So notwithstanding what's in anybody's reports, anybody that does that is violating a direct order of the Court and the Court will assume that they are doing it at their own peril.

Do you want to talk to me about these equitable defenses?

MR. LIVEDALEN: Sure.

With respect to the --

THE COURT: You think there's a nexus with a live and appropriate jury issue here, or --

MR. LIVEDALEN: Yes, Your Honor.

THE COURT: -- how would it be on the question I asked opposing counsel?

MR. LIVEDALEN: Let's go to slide 26, please.

So Mr. Sheasby said earlier that there is absolutely no connection between Netlist's relationship with JEDEC and the issues in this case. That is not correct.

So we see here that there were proposals in JEDEC. We see them listed, committee items 158.01, et cetera. Those

were proposals that were provided to JEDEC, proposals that the inventor Mr. Hyun Lee had access to as a member of JEDEC and was a participant in meetings where these items were discussed. And those exact items then are now found in the accused -- or found in the standard that's implemented by the accused product.

So this is relevant for two reasons. One, the inventor indicated that the item numbers and the discussions here in these committee item numbers were not his invention. He very clearly said they were not his invention. Yet those exact same features show up in the accused products.

And so we see that here. In the deposition testimony of Mr. Lee, he says -- when asked about technology that are in these disclosures, he said they are completely different.

And so not only the fact that his testimony relates to these disclosures, but that he was intimately aware of them because he actually attended the meetings and knew what they were and then indicated that they were different, that goes directly to the issue of infringement because, like I mentioned before, the technology that's discussed in these disclosures is the exact same technology that ended up in the standard in which the Samsung products complied with.

Relatedly, there is other issues that I mentioned that do exactly overlap. For example, the prosecution history of the patents, that is obviously discussed with respect to the issue

of laches. But if we see here and later in -- bear with me one second.

Later in other parts of Mr. McAlexander's report, he incorporates by reference the discussion of the prosecution history that's discussed in the laches section because, obviously, those issues overlap with other things, like written description, for example, and other issues.

So, again, to the extent there's underlying factual issues, we see no issues with those being discussed before the jury. Obviously Mr. McAlexander will be not making conclusions about equitable issues such as laches.

THE COURT: I sure hope not.

All right.

MR. LIVEDALEN: All right. We talked about -- we do agree that the printed publication issue is moot.

I believe the next one is the issues with JEDEC.

THE COURT: Yes.

MR. LIVEDALEN: So, first and foremost, the paragraphs that Netlist is seeking to strike say basic things like Netlist did not contribute to the standard. So you heard Mr. Sheasby yesterday indicating that Netlist's inventions made it into the standard and that now Samsung is infringing by practicing the standard, but the reality is that Netlist did not contribute. And so that issue is relevant. It's not only relevant to the issue of non-infringement; it's relevant

to the issue of no willfulness; it's relevant to the issue of no inducement. So the fact that Netlist did not make any contributions to the standard is directly related to JEDEC.

The other issue that it's related to is obviously damages, which we talked about yesterday, that the contributions of others in comparison to Netlist is also relevant.

It's also relevant to rebut Netlist's invention story. Again, as Mr. Sheasby said yesterday, Netlist is going to tell a story where it developed its inventions; those inventions then made it into the standard. We believe that Netlist's relationship to JEDEC, the fact that it had access to the standards as it was being developed, are critical to rebut those allegations.

And as we see here, we see -- this is an issue that we just talked about. Mr. Hyun Lee, the inventor of the '506 Patent, had access to JEDEC materials that were being presented. Those materials were before the provisional of the '506 Patent. And so there is an issue about exactly where did these inventions come from, and they did not come from Mr. Lee. So, again, that is not -- that is directly a relevant issue that needs to be presented to the jury. It's certainly not prejudicial.

And I guess that's one thing, zooming out a little bit, the entire basis for Netlist's motion is that

Mr. McAlexander's discussions of Netlist's relationship with JEDEC are prejudicial, but they've not identified any issues that are actually prejudicial with respect to Netlist's participation, access, or things it contributed or did not contribute to JEDEC. All those are just the parts of this case. There is no prejudice.

Your Honor, do you have any other questions about the JEDEC issue?

THE COURT: No. Let's move on, please.

MR. LIVEDALEN: Okay. With respect to the other litigation, I think we're in the same spot we talked about earlier today. Mr. McAlexander has no intent to talk about other litigation for purposes of the jury trial. With respect to equitable issues, that may be a different story. And, again, if Netlist opens the door on certain issues, we would at that point in time request Your Honor's guidance.

THE COURT: Well, let me call your attention to paragraphs 170 through 174. Mr. McAlexander presents two different time lines, and they reference various other litigations and details as to the outcome of those litigations. That's exactly the kind of thing that he shouldn't be putting in front of the jury. I don't know how he presents these time lines that are populated with these references to other litigation and does that consistent with the Court's directives.

MR. LIVEDALEN: Yes, Your Honor.

So I think the time lines that we can present, we -- obviously the time line includes other things than litigation, so we can certainly take out portions that reference other litigation there. But I believe, again, to the extent this issue relates to a laches issue, then those would properly be included in the timeline.

THE COURT: What about the last issue in the briefing, invalidity arguments not disclosed during discovery?

MR. LIVEDALEN: Right. So --

THE COURT: I assume you are going to tell me those are for background and the state of the art.

MR. LIVEDALEN: Your Honor, there's two issues there.

Go to slide 18, please.

So there's two issues that Netlist identifies in this section. The first, as you correctly identified, that there is certain references, whether textbooks or other types of prior art that identify the background of the invention. That's exactly what Mr. McAlexander uses it for.

And we see here in the *Allergan v. Teva* case, it says, "Courts in the Eastern district have permitted defendants to rely on references that were not included in their invalidity contentions when the references were not proposed as invalidating prior art, but were directed to other purposes,

such as showing the state of the art at the time of the invention."

So, again, Mr. McAlexander is not using these references to say the claims are obvious, the claims are invalid; he's simply using them to set the stage for what the technology was about, what the state of the art was at the time to put into proper context the invention in this case.

THE COURT: Anything else you need to present as far as argument?

MR. LIVEDALEN: Yes, Your Honor.

If we can go to the next slide.

So Mr. Sheasby argues essentially that Samsung's invalidity contentions were insufficient because, in his view, they needed to be providing more detail in terms of why Samsung alleged that there was no written description for certain claim terms.

Number one, the local rules don't require that there be any special explanation as to the basis for written description. And if you -- it sort of makes sense. Right? Because the idea is there's no written description. We identify a term; we say there is no support for this term. It's very clear what we are identifying in terms of our contention when we identify a term that lacks written description.

And as we see here in the *USAA v. PNC Bank* case, the

Defendant in that case PNC says that its identification of the relevant claim terms put USAA on notice of enablement defenses, and the Court agrees with PNC.

So in this case there's already a question about the sufficiency of disclosures under § 112, and the Court found that the identification of the term 'the lack of written description' in the invalidity contentions were sufficient.

And as we see here, the *USAA* court cited in *Ultravision Technologies v. Lamar Advertising*, and here the Court said, "Defendant argues that particular terms are indefinite and, consequently, every asserted claim" --

THE COURT: Slow down. In other words -- you don't need to read it to me. I can read it on the screen.

MR. LIVEDALEN: Okay.

THE COURT: I see it.

MR. LIVEDALEN: Okay. Thank you.

And so not only did we identify the terms that we believe lacked written description, but we also identified two different theories. If there was a theory that the term just was not disclosed, we identified those; but we also put -- as we see here there's an asterisk, and we said -- on the next slide, slide 22, we provided an asterisk: "In addition, claims reciting these terms are invalid under § 112 because the entirety of the specification does not support the breadth of the claims as written."

So we provided the terms that were not at all disclosed in the specification, and we identified the terms that we thought were too broad and not supported by the specification. So there are two particular categories in § 112 terms that we identified. So the notion that we didn't put Netlist on notice of these is just not true.

And, finally, I would say there's absolutely no prejudice. Right? There's -- the evidence that the experts need to analyze is the patent, and Netlist had the patent since it filed its complaint. We put these defenses in the expert reports, and Netlist had the opportunity to rebut them in their respective expert reports.

THE COURT: All right. Anything further?

MR. LIVEDALEN: No, Your Honor. I believe that's it.

THE COURT: Anything further from Netlist on this motion?

MR. SHEASBY: Two things.

The first issue, the overlap between the equitable defenses and the defenses at law.

Madam Courtroom Deputy, if I could have the elmo.

This is what they pointed to, which was a relationship between the JEDEC standards and non-infringement. The JEDEC standards are no longer prior art in this case. They can't be used for invalidity purposes.

And so what they're really trying to do is somehow connect references that have been excluded for validity purposes, which they're going to use for inequitable conduct, and still try to inject them into the case, and that is exactly why there's no clear, concrete connection between the equitable issues and the issues in this case.

The second issue -- and, by the way, that goes for their arguing about JEDEC as well. Any discussion of JEDEC in Mr. McAlexander's opening report which goes to validity has no role anymore because JEDEC references are not prior art in this case. They've been excluded under a *Sotera* stipulation.

Last issue.

Can we go to slide 74 of my deck?

This is not prior art as background. This is presenting a combination that was not disclosed. It has nothing to do with background. This is an undisclosed combination.

Thank you, Your Honor.

THE COURT: All right.

MR. LIVEDALEN: Your Honor, could I clarify the issue about -- in the status of prior art? I think he's saying certain issues are off the table. I think that's incorrect.

THE COURT: All right. Take a second and give me your position on it.

MR. LIVEDALEN: Thank you, Your Honor.

So the issue is not -- we're not using these limited JEDEC materials as prior art whatsoever. The issue is --

THE COURT: What are you using them for?

MR. LIVEDALEN: So, one, the first JEDEC reference, the inventor was asked about whether this JEDEC disclosure is his invention. He said no. That same disclosure is now in this published standard. So it's not an issue of prior art; it's an issue of non-infringement.

The separate issue is a § 102(f) issue, which is a derivation issue, which certainly cannot be brought by the Patent and Trademark Office. There's no question about that.

And so, again, we're not arguing that anything's prior art, that it was published, that it was accessible; we're saying that the inventor took the invention from somewhere else, from someone else, and put it in his specification. Again, that's not an issue that's covered by *Sotera*; that's not certainly anything that was covered --

THE COURT: Does that mean you're saying the inventor stole the concept from somebody else.

MR. LIVEDALEN: Derivation you can use the term that you want to use, but the idea is that derivation, § 102(f), that the invention is not the inventor's. That's the requirement. So we just -- we need to show that there was access and that he obtained the invention from somewhere else.

THE COURT: All right.

MR. LIVEDALEN: Thank you, Your Honor.

MR. SHEASBY: Your Honor, I hate to ping-pong this, but to be clear --

THE COURT: This is the last go-round.

MR. SHEASBY: This will become clear at the trial.

They have no independent § 102(f) derivation defense. The § 102(f) deviation defense just references the anticipation defense based on JEDEC materials that has been stricken. So they can't evade a *Sotera* stipulation by just using the same combination and calling it a § 102(f) to avoid § 102(a).

Thank you.

THE COURT: I have a feeling I'm going to -- no matter how I rule on this motion, I'm going to hear about this in the middle of the trial.

MR. SHEASBY: I think that's right.

THE COURT: But that may be unavoidable.

Okay. With regard to Document 212, in the first issue concerning alleged improper claim construction, I find that the following portions of Mr. McAlexander's report do cross the line into improper claim construction, and I'm going to exclude them from his report. They are paragraphs 178 through 196, 198 through 209, 210 through 220. The balance of the motion seeking to curtail or strike portions of his report based on improper claim construction are denied.

With regard to the equitable defenses, I'm going to try to announce this ruling in a manner that's consistent with the way I've approached this issue earlier in the pretrial.

As to the equitable defenses themselves, as to legal conclusions, as to legal terms and language that is inherently legal in nature, I'm going to grant the motion. As to the factual underpinnings that stop short of making a legal defensive assertion, I'm going to deny the motion. I'm not going to curtail the witness from talking about the events that factually, in the Defendant's view, may at the bench trial give rise to an equitable defense, but I am going to prevent the witness and counsel from talking about actual legal defenses.

I don't want to hear the word 'laches' in front of the jury. I don't want to hear the word 'estoppel' in front of the jury. I don't want to hear any ultimate conclusions that might be properly argued at the bench trial at a later date in front of the jury. But what gives rise factually to those assertions, whether it was delay, whether it was slumbering on someone's rights, whether it was disclaiming something, and taking a position that that no longer applied, the underlying factual allegations are not excluded. But the witness needs to be particularly careful--and I'm instructing Defense counsel to make sure the witness understands this--I don't want to hear judgments, conclusions, legal terminology, they

are estopped because of this, they have committed laches because of that. Anything like that is going to run afoul of this ruling and is going to be met with swift condemnation by the Court.

So if you have any doubts about where the line that I'm drawing is, please let me know. I've tried to draw this same line at various points earlier in this pretrial proceeding, and this ruling is intended to be consistent with my earlier rulings on this same subject matter.

With regard to the Netlist relationship with JEDEC, I'm going to strike paragraphs 170 through 174 because they are replete with references to litigation in the Central District of California and before the ITC, but I'm not otherwise going to constrain the witness.

That being said, I think everybody understands that it's been established clearly on the record through this pretrial proceeding that none of the asserted patents are standard essential, they are not encumbered with a RAND obligation, and how there's going to be relevant references to the standard-setting organization JEDEC is going to have to be very carefully presented to have any real relevance or application here.

I'm not going to go further other than I'll expect if there is an excessive use of the JEDEC standard beyond something that relates to a jury issue, I'll hear objections

about it in real-time and I'll take it up. But the JEDEC standard has a significantly reduced relevance in this case in light of the fact that there are no standard essential patents, there's no third-party beneficiary enforcement of a RAND obligation, and that largely negates the relevance of JEDEC. I'm not going to say there's no relevance, but I think it's very clear that the prominence, or lack thereof, of JEDEC has been significantly curtailed through the process of what we've gone through yesterday and today. The printed publication issue seems to be out by mutual agreement that it's moot.

The improper arguments on written description, I'm going to deny that portion of the motion.

I've already struck paragraphs 170 through 174 based on their reference to irrelevant other litigation. I would do it again if I hadn't already done it.

There's a request to strike paragraphs 222 and 242 through 244 of the McAlexander exhibits E and F. That's denied.

And with regard to the invalidity issues not disclosed during discovery, I'm not going to strike them, but I am going to limit their use, as the case law clearly delineates, as background material relevant to the technology and the state of the art, and that's the only proper use for them in this trial. And if there's an effort to use them or assert them

beyond those targeted parameters, I'll expect objections to be raised at the time. But I'm not going to expressly strike these provisions of the expert's report.

All right. That's the ruling on Document 212.

We'll go to Document 213 next. This is the Plaintiff's motion to strike the portions of Dr. Gabriel Robins' opening expert report.

Also, counsel, there is a recently filed motion for sanctions regarding Doctor Robins and his deposition and a communication with the district's hotline. I think it's appropriate to hear those concurrently. So when we are talking about Document 213, I also want to hear argument on the Plaintiff's motion for sanctions here as well, and we'll take those up concurrently.

MR. SHEASBY: Your Honor, may I grab something from the bench?

THE COURT: You may.

MR. SHEASBY: Thank you, Your Honor.

THE COURT: Although this is the bench, that's not the bench, but you can grab whatever it is.

MR. SHEASBY: It is what?

THE COURT: You said grab something from the bench. This is the bench, not that.

MR. SHEASBY: A bench.

Your Honor, may it please the Court.

I'm going to go through the Robins sections relatively quickly, but I did want to flag that the behavior with Robins is different in kind from the behavior with Mr. McAlexander.

So Mr. -- Doctor Robins doesn't just discuss a prolix amount of background art; he literally cites the art in his discussion. This is paragraph 461. This is in his invalidity discussion.

So what he's doing is intermingling together the, quote, background art and his cited art to combine it into just, you know, a hundred-reference combination. And so what -- the guidance that we need from the Court is that this is not just a situation in which we can say, Okay, don't talk about the background art; talk about your invalidity contention -- just your combination --

THE COURT: Let me ask a point of clarification. Is it Defendant's intent to present Doctor Robins live, or is he going to be presented by deposition?

MR. COLVIN: Doctor Robins will be here live, Your Honor.

THE COURT: Okay. Well, that gives the Court more flexibility than being limited to his deposition. But go ahead, Mr. Sheasby.

MR. SHEASBY: So what I don't know what to do is that when someone -- there's speeding and there's reckless driving. Mr. McAlexander sped by discussing Elpida in

something. This goes beyond speeding. This is extreme reckless driving, in which the actual invalidity analysis, the predicates for why he says the patents are invalid is larded with -- it's not just background art that was disclosed and not elected; it's undisclosed art that no one's ever seen before, that wasn't disclosed in fact discovery, that we didn't have the opportunity to engage on. And that goes beyond a level that I sort of am comfortable with just saying it sort of goes to the heart of the analysis. And it's not just once; it happens repeatedly paragraph after paragraph after paragraph--paragraph 461, paragraph 470, paragraph 486. The danger of it is --

Let me give you an example. This is paragraph 391 where he talks about a 2006 creation of stack memory by Samsung. It's not elected prior art, it's not part of his combination, and yet he's discussing it as if it has some reference to invalidity of these references. In effect, he's using an undisclosed Samsung reference to an attempt to invalidate the patent.

If they would have disclosed that they thought this was, quote, background art or would have listed it in their elected prior art references, it would be one thing for us to be able to defend it. We don't even have that option because we weren't able to depose anyone who was the author of this reference because they didn't disclose it.

So think about this. They disclosed hundreds of references in their invalidity contentions. This is not a fight about them using references in their invalidity contentions as background; this is the fact that Doctor Robins pulled off the shelf hundreds of additional undisclosed references, including Samsung's own designs, which were never disclosed before, and is using them as the predicate to his invalidity defense.

Like I said, this is beyond something that can be cured by just saying don't discuss the background art. These paragraphs need to be stricken because his invalidity analysis is based on undisclosed art, I would respectfully submit.

Your Honor, that's the first issue. If Your Honor has no questions, I'll go on.

THE COURT: Please do.

MR. SHEASBY: So the second issue is JEDEC presentations as prior art. The JEDEC presentations as to Doctor Robins have not been stricken because the PTO has not initiated yet, so the *Sotera* stipulation will not apply. And so the problem that we have --

Let's go to slide 17.

-- is there is no foundation from Mr. Gabriel Robins that -- Doctor Robins that JEDEC materials -- these JEDEC materials were publicly available. He has no role in JEDEC. He's never participated in JEDEC. Doctor -- Mr. McAlexander,

who he also relies on, has no experience in JEDEC process.

And so this is an expertise issue which is, as to these particular references, there's no one with knowledge who can lay the foundation that they were publicly available. And so this does bubble up, which is there a factual dispute as to whether they are publicly available. There may be a factual dispute that survives summary judgment, but what there is not is any competent expertise from Doctor Robins that they were publicly available. And so Doctor Robins should not be able to state that these are publicly available. He has no experience with JEDEC and he wasn't involved in it.

That's the second issue, Your Honor. I can move onto the third.

THE COURT: Please do.

MR. SHEASBY: Thank you, Your Honor.

This is another foundational issue, which is, and I'm flagging this for the Court, is this is his combination, it's 10 separate JEDEC documents that he says should be treated as one reference. He has no analysis of a motivation to combine these 10 references.

You're not going to be able to -- it's a fight we're likely going to have in front of the jury, but I want to flag this issue that when he gets up and talks about these 10 references, there will be an objection that there's no discussion of motivation to combine these 10 references, and

there isn't because it doesn't exist in his report. And so we ask that he be forbidden from claiming -- from going outside his report and discussing a motivation to combine these 10 references because it doesn't exist.

Next issue. So one of the ways they get around this idea of undisclosed prior art is they talk about -- they combine the knowledge APA, admitted prior art, with the JEDEC presentations and with knowledge of a POSITA, and he does that repeatedly where he combines it with knowledge of POSITA.

Here's the problem. This is what his knowledge of a POSITA is. It's all the references that were not disclosed. So what he's done is he's laundered all these undisclosed references into his invalidity opinions, not just his background, but because his knowledge of the art, knowledge of a POSITA is these undisclosed references.

So that's where I'm talking about. This is not just run-of-the-mill speeding; he talks about knowledge of a POSITA as the basis for his combination, and his knowledge for a POSITA is all of these undisclosed references.

And so what happened here is someone made up -- and I understand why. It's highly, highly effective. If you just say knowledge of a POSITA, they say, We can throw anything into that bucket and we'll get away with it, but his knowledge of a POSITA is based on all these undisclosed references which makes the disclosure process in the Court just irrelevant,

because if they can dump anything you want into knowledge of a POSITA, you haven't disclosed anything.

And so, like I said, this goes beyond just regular knowledge of the background information; this is a misuse of undisclosed references.

And again, I'm just flagging--it occurs again and again and again in his report; not just once.

Improperly comparing prior art references to the embodiments of the invention. In my mind, this is not appropriate at all. When he's putting up a figure from our patent and he's comparing it to a figure in a prior art reference, not as background but in a section between -- where he says the similarity of Kim and the purported invention of the patents-in-suit by showing a figure, that's completely inappropriate for a jury. The figure is not the claim. He needs to compare Kim to the claims. And doing something like this gives the jury a false impression that this is the analysis they should do, in which they match a figure from one to another.

I want to be very clear. Is it an appropriate use for figures in our patents in discussion of prior art and discussion of infringement? There certainly is. Has this gone well beyond what's appropriate? I would respectfully submit that it has.

Next issue--improperly attempting to alter invalidity

contentions in the expert report. In their invalidity contentions, they claim that an array die is any DRAM that can be stacked. That's what they claimed in their invalidity contentions. That's a problem with them because -- problem for them on infringement because on infringement they want to say a DRAM that can be stacked is not an array die because it's not a DRAM circuit.

So what they did is they did this creep in which they created a new construction in which they say our invalidity contentions are based on stacking of things other than DRAM; they are based on stacking of non-volatile memory. But the problem with that is that they never disclosed that as their invalidity contentions.

So this is another flag. When he gets up and says, I think the patent's obvious because it would be obvious to stack non-DRAM memory, that is nowhere in the invalidity contentions; nowhere. It does not exist.

Here's what their argument for why with they exist is. This is all the prior art reference they disclose, and they point out that in this massive list there are two references that in addition discussing stacking DRAM also stack non-volatile memory. They never disclose that their theory was non-volatile memory. They never disclosed any of their theory in that the patent is made obvious by references that involve DRAM that can be stacked with TSVs. They need to

change that theory to accord to their new non-infringement theory, but they should be held to their invalidity contentions.

In my mind that's the take-away from this, which is what's in the invalidity contentions has to be what Doctor Robins says. He can't say, Well, under Netlist's theory it's invalid, because that's not what the invalidity contentions say. The invalidity contentions are very explicit that they are applying it to 'DRAM circuit' -- they are applying it to DRAM with TSV in it, which is that.

So this is an example of don't let them change the position. In the invalidity contentions, if stacked DRAM with TSV satisfies the claims, they can't turn around and now say their invalidity contentions is Netlist's interpretation or switch it to non-DRAM stacking.

Those are the issues on the motion to strike. And I'll now deal with the motion for exclusion as well, with your permission, Your Honor.

THE COURT: All right.

MR. SHEASBY: Your Honor, let me begin by saying that I'm not the easiest person to sit for a deposition with, and I understand that I can be an aggressive examiner.

THE COURT: You have that reputation, Mr. Sheasby.

MR. SHEASBY: What I can say is that this has gone beyond the ebb and flow that exists in aggressive examination.

In other words --

If I could have the elmo.

This destroyed our ability to obtain effective cross examination from this witness on critical, critical issues on injecting ad hominem statements about what was known in the prior art into his deposition. He knows why he did this. Because I can't publish this to the jury, because if I publish it to the jury it just makes his improper point again and again. And so by injecting constantly into his answers not just ad hominem, not just prolix answers, but actually affirmatively injecting advocacy to destroy the ability to cross examine the witness goes far beyond this; goes far beyond what's normally allowed.

We actually sought the magistrate's intervention twice on this issue, and after we sought the magistrate's guidance, he literally did the same thing again and again and again. This is immediately after he was admonished where he started doing again; the exact same thing that he was admonished not to do. And he did it throughout the day; refusing to answer basic questions about claim terms, asking me for my definitions of terms, for terms that are -- he's saying are the plain and ordinary meaning, evading, answering completely different questions from what I asked him, spending 14 minutes looking at random JEDEC documents when I didn't ask him about JEDEC documents, staying silent for over and hour and I think 15

minutes of the 7-hour deposition he sat in stone silence refusing to answer the question. And I let him. And I let him because that's what I properly needed to have done.

But just like this expert report, the behavior in this deposition has gone far beyond, far beyond what is appropriate in this court, and it destroys my ability to cross examine him. How can I ever put this up in front of a jury on cross examination? Because he's just feeding in his improper argument.

So sometimes people are clever. They make a big deal of things later and they didn't seem that big of a deal at the time. I can tell you, I emphasized again and again that this was a big deal. I went overboard. I admonished this witness more than I should have to answer these questions directly. I called the Court twice to deal with this issue. I told counsel that I was going to seek sanctions. I don't know what I can do because this behavior really impairs my cross examination.

THE COURT: When was this deposition taken?

MR. SHEASBY: The deposition was taken in December, Your Honor. In January. In January, Your Honor.

THE COURT: Okay.

MR. SHEASBY: And I specifically talked to the magistrate, and I said to the magistrate this is not about me just getting a do-over once he's had time to sort of calm

down. I said that I was going to seek preclusive sanctions if this continued, and it continued. And you see it again and again and again of this evasion. And so I cannot effectively cross examine the witness with this behavior. It's impossible.

So at this point I do request a remedy that for the questions that he was not responsive on, he should be precluded from answering.

THE COURT: What else, Mr. Sheasby?

MR. SHEASBY: That's all I have, Your Honor.

THE COURT: Let me hear from Samsung on both the motion regarding striking portions of his opinion as well as the sanctions motion.

And I understand, Mr. Colvin, you were the opposing counsel at this deposition?

MR. COLVIN: Yes, Your Honor.

THE COURT: You're the right person to argue this.

MR. COLVIN: May I approach, Your Honor?

THE COURT: You may.

Mr. Colvin, I'm particularly interested in your response to the allegation that the witness is laundering undisclosed invalidity references as known prior art, and I'd like you to make sure you give me plenty of explanation on that.

MR. COLVIN: Will do, Your Honor. And that will come up straight away, so we'll just get right to it.

THE COURT: Okay. That's a good approach.

MR. COLVIN: So two things to note there, Your Honor. First, Doctor Robins has an extensive technology overview section in his report. I've got the table of contents of that highlighted here on the screen for you. And Netlist's motion, part of their paragraphs that they're asking to strike are in this section here.

Now, Doctor Robins, as plainly seen when you look at his report and look at this technology overview section, it's voluminous. There's a lot there. He starts with Boolean Logic and goes all the way through the evolution of modern memory systems to get to where we are today. We think that is appropriate and we think, regardless of the Court's other rulings, he should be allowed to march through that or at least --

THE COURT: We don't have time for a tutorial on Western civilization.

MR. COLVIN: Understood, Your Honor.

But we do think Doctor Robins should be able to pick and choose from this section to present a technology overview and a bit of history for the jury. That's all this section will ultimately do.

THE COURT: History's one thing. Asserting invalidity based on something the other side hasn't been given fair notice of is completely a different issue.

MR. COLVIN: Understand, Your Honor.

THE COURT: That's why I've given the speech a couple of times today about background is background and it's not a license to use something improperly by calling it background.

MR. COLVIN: Understood, Your Honor. And we hear that loud and clear, and I can short circuit a bit of this.

We will not use any of the references that they've complained about that were not disclosed in the invalidity contentions to check a box. We will not use them when discussing prior art that could check a box. So we're going to talk about the background, and then when we start talking about prior art and checking boxes for invalidity and obviousness combinations, we're not going to launder all of these undisclosed references. We're not going to do it, Your Honor.

And so if we look specifically at paragraph 461 -- now this is just a part of paragraph 461. Paragraph 461--Mr. Sheasby showed it to you--is long and it includes a long list. And I can put it on the elmo, if Your Honor would like to see the whole thing. But 461 starts by saying, "The concept of dividing stack arrays would have been obvious in view of a knowledgeable POSITA." POSITA knew about this from a few different ways. He knew about it consistent with the '060 Patent's admission that it was known. So that's the admitted

prior art.

THE COURT: Let me ask you this, Mr. Colvin.

MR. COLVIN: Yes, Your Honor.

THE COURT: Are you going to put this witness on at trial?

MR. COLVIN: I believe I am, Your Honor.

THE COURT: Well, let me change it. You are going to present this witness at trial. I don't want some other lawyer who didn't make that very clear concession on the record in front of me to put this witness on and then, Oops, I didn't know that this was as targeted and limited as Mr. Colvin said when I have to show them the transcript.

MR. COLVIN: Understood, Your Honor.

THE COURT: So he's your witness. Okay?

MR. COLVIN: Yes, Your Honor.

THE COURT: All right.

MR. COLVIN: But just to explain it, and I won't belabor this point, but that additional set of prior art reference was just, again, Doctor Robins showing his work that when he says a POSITA knows something, he's not just saying it; he put those references in there to show that. So that's with respect to checking the box.

Now, with respect to motivation to combine, the *Acoustic Tech* case says that motivation to combine prior art reference can come from knowledge of the skill in the art. Doctor

Robins does a similar thing there. Motivation for dividing, in this case we use the same example. Motivation for dividing arrays into groups consistent with Kim, JEDEC presentations, and the knowledge of the art. So he's got these several places where the motivation comes from. He said it would have been clear and well-known to a POSITA to reduce the load on drivers, and he continues. Then he says, 'for example', and he has a similar list.

Again, this is Doctor Robins showing his work, but we're only going to draw from the well of things that were disclosed in the invalidity contentions, Your Honor, for motivation.

THE COURT: All right. Continue, please.

MR. COLVIN: I'd like to talk about Kim.

So Kim was a reference that was disclosed in the invalidity contentions. Doctor Robins uses it. And there is an objection to this figure of Kim.

Now, we understand that there's a rule that you can't compare prior art to figures, but that's not really what's being done here. When you look at this, what Doctor Robins puts in the middle are essentially claim terms, and so what's being compared here is the claim terms to the prior art shown on the right and the claim terms to an embodiment of the asserted patents on the left. And so because of that, we think that it's perfectly appropriate. We think it will be helpful to the jury to be able to show where, you know, data

ports are here and where data ports are there; to show where certain die interconnects are here and where die interconnects are there. So what we think, because of what is being compared here is claim terms, we believe that that's okay, Your Honor.

Now, there is a -- opposing counsel makes an argument that this Jeddeloh and Ma reference that we use to show that you can stack die that are not DRAMs, they have an argument that that wasn't disclosed.

And if I can have the elmo for a moment.

Opposing counsel showed you this figure from slide 38 of his presentation, and this is the invalidity contentions, and he uses this to say that we only disclose that DRAM dies can be stacked and that we did not disclose that non-DRAM dies can be stacked. But when you look at this figure, I mean, this image of these invalidity contentions, what you'll see is that we're not talking about DRAM dies here; we're talking about array dies that can be stacked.

And we cite a whole bunch of evidence in our invalidity contentions where we show that array dies can be stacked, and we show it to support Netlist's original construction that DRAM dies are array dies. And we also have things like Jeddeloh and Ma if.

I can have the slides, please.

We show Jeddeloh and Ma also disclose stacking of die;

they just disclose things other than DRAM. This is where it shows up in Doctor Robins' report. These very same citations that Doctor Robins used are the exact same citations that were disclosed in the invalidity contentions. So we believe everything is above board there, Your Honor. No one's trying to hide the ball. We cited it in our invalidity contentions and Doctor Robins uses it.

With respect to 'chip select signal', Doctor Robins shows where the patents support his understanding of the plain and ordinary meaning. I think this is a very similar argument to what we had with respect to the previous Robins motion and the previous Doctor Brogioli motion. He's looking -- he is forming his opinion of what the plain and ordinary meaning is, and he's looking at the specification because that's what the case law tells him he's supposed to do.

This was a term that was argued at claim construction, but the Court construed it to be plain and ordinary meaning. And in that situation, the Plaintiff had argued for it to be one or more array dies that are selected, and Doctor Robins and the Court rejected that and found plain and ordinary meaning. And so Doctor Robins' construction of this is completely consistent with this entire picture from claim construction to the patent specification.

So that will conclude my argument on the motion to strike, Your Honor. Unless you have further questions, I'll

talk about the deposition.

THE COURT: Let me hear about the motion for sanctions.

MR. COLVIN: So first, Your Honor, we obviously have not put in a response to this yet. The motion they served only last week, and so we haven't gotten our full response into this yet.

THE COURT: Are you comfortable giving me an oral response now, or are you wanting to defer further consideration of this until you file a written response?

MR. COLVIN: I can certainly give you a preview, but to the extent we're going to dig into specific questions that are buried in the transcript, I would like to give you a full written response on that so that we can lay it out all on one page.

THE COURT: Well, we really don't have the luxury of a tremendous amount of excess time here. I don't want you to tell me or, more importantly, tell the Federal Circuit I made you argue this when you hadn't had a chance to file your response, but we're not going to do it twice. I'm happy -- you know, I'm happy to hear your argument today. I don't intend to drill down into exact quotes.

I've seen, by and large, the tone and the tenor of the deposition, and I don't intend to go any further than that, and it would certainly be helpful if we could get this issue

handled today. I just don't want to hear later about how your due process rights were trampled on by the Court because you didn't get to file your response within the time limits that the rules allow. So let's go left or let's go right, but we're not going to do both.

MR. COLVIN: Okay, Your Honor. Let me tell you what we think about the motion.

THE COURT: Okay. And I'll take it you're ready to go forward.

MR. COLVIN: So first, Your Honor, this motion is quite tardy. The deposition you heard from Mr. Sheasby was a long time ago, at least six weeks, probably even more.

THE COURT: And there wasn't a motion to compel filed with the Court here.

MR. COLVIN: That's correct, Your Honor.

THE COURT: I'm well aware of that.

MR. COLVIN: And so they filed this motion. We think they should have filed it right after the deposition if they were going to file it. They filed it here, but they haven't asked for leave, they made no attempt to show good cause for filing it when they did, and so we think as a matter of course that should toss out their motion. They just haven't complied with what they need to comply with in order to even get it properly on file in front of Your Honor.

Also they barely cite any case law. They cite I think

two cases, and I think neither one are particularly relevant. They made no attempt to show the required bad faith. And they seek to preclude what amounts to Doctor Robins' -- you know, basically all of his analysis. And I haven't gone through to pick out what would remain that he could talk about, but it is an extreme sanction that they're asking for, Your Honor.

And Doctor Brogioli did the same thing. And so we obviously don't have a paper in for Your Honor to look at all that and to look at the transcript, but he did. Doctor Brogioli's deposition was a day or two before Doctor Robins' deposition. Doctor Robins sat in on Doctor Brogioli's deposition. And Doctor Brogioli did the same thing.

Their motion starts or is sort of grounded in Rule 37, but Rule 37 doesn't apply because there was no order from Judge Love. Judge Love admonished the witness. He also had some things to say to Mr. Sheasby. But Rule 37 doesn't apply. Rule 30(d)(2), which is what we would base our response on, requires bad faith, and they've made no attempt to show bad faith.

And so those would be the contours of our response, Your Honor. I think big picture-wise, Doctor Robins is going to be in the chair right beside you, Your Honor. Certainly he's going to have to answer the questions that are put to him in front of the jury. And we certainly expect him to do that. And I know the Court expects him to do that, and that's what

he's going to do.

So certainly, Your Honor, I think that the idea that he wouldn't be allowed to testify in this court on any of these issues is just very extreme. We think there -- that is not supported by any case law, they really haven't made an attempt to support it with the case law, they haven't analyzed it under the correct rules. If they were really serious about this, Your Honor, they would have filed their motion right when it happened.

THE COURT: All right. Mr. Sheasby, do you have some brief follow-up?

MR. SHEASBY: Just brief follow-up, Your Honor, and it's relating to the motion to strike.

If we can go to slide 24.

The -- I do appreciate the representation that no undisclosed art would be used for the obviousness. Here's the problem. His combinations have something called knowledge of the POSITA. So he's going to talk about knowledge of the POSITA in his, quote, background about how it would be obvious to stack, how it would be obvious to do that. And then he leads it straight in and he says to the ladies and gentlemen of the jury, And this missing element, the knowledge of the POSITA would hear that. Jurors aren't stupid. They just heard 15 minutes of him talking about the knowledge of the POSITA in the background.

And so if he is going to -- the remedy we would seek, then, is that his knowledge of the POSITA should be removed from these limitations. In the alternative, for him to say knowledge of the POSITA, he should be citing to a disclosed prior art reference. That would be the only way to remedy this. This has gone beyond just it's going to be background. His knowledge -- what he says is knowledge of the POSITA is the undisclosed prior art.

So this idea that background is separate from the combination has no connection to reality. The combination is the improperly disclosed art because it says knowledge of the POSITA. So it's -- it does us no good to say he's not going to discuss prior art with the combination because his prior art -- undisclosed prior art -- his undisclosed prior art is the combination.

That's all I wanted to say, Your Honor.

THE COURT: All right. Thank you.

All right. With regard to Document 213 and with regard to the motion for sanctions brought by Plaintiff concerning the deposition of Doctor Robins, which I've seen because it's been attached to the motion, here's the ruling of the Court.

I'm going to hold Samsung to the concessions and the representations Mr. Colvin made on the record today. Mr. Colvin's going to put on this witness as the sole counsel interacting with Mr. Robins during the trial. Beyond that,

I'm going to deny the motion to strike.

With regard to the motion for sanctions, it appears to me that there are unclean hands on both sides. There was no motion to compel brought promptly after the deposition was taken. This motion for sanctions was not filed immediately or with leave of the Court. But by the same token, a cursory review of the transcript reveals that the witness was highly and consistently non-responsive, and I am bothered by the inability of Plaintiff to have a clean deposition to work off of and to use for potential impeachment with this witness given the way the witness comported himself.

And it's also clear to me that though there is not a written order from Magistrate Judge Love, this district has a long-standing and well-working hotline procedure. That hotline procedure was invoked. There was -- there were directives given to the witness that, by and large, were not complied with. So I don't think the Plaintiffs handled this the right way, and I don't think the witness comported himself in a proper way in the deposition.

So with regard to the motion for sanctions, here's what the Court's going to order. I'm not overtly sanctioning the witness, but I am ordering that the witness appear for a new deposition of not more than five hours in duration to be taken in this building with Mr. Colvin defending for the Defendant, and either Mr. Baxter or Ms. Truelove taking the deposition

for the Plaintiff. Mr. Sheasby is not to be present or to participate in that new deposition. And it will be taken at a time consistent and in cooperation with Magistrate Judge Payne's schedule so that he is here, and he can be an immediate on-the-scene hotline resource if any problems come up in that deposition.

And I'm instructing Plaintiff's counsel to make it abundantly clear to Doctor Robins that his responses -- well, Defendant's counsel, I should say, make it abundantly clear to Doctor Robins that he is to be responsive to the questions asked and that he is to avoid inserting material answers responses, opinions, or any other statements not called for by the question.

And we're going to get a clean deposition of this man and we're going to have a clean deposition by the time this goes to trial. And Mr. Colvin's going to put him on direct, and he's not going to launder undisclosed prior art references into some kind of exaggerated everybody in the prior art knows all these hundreds of thousands of things, as Mr. Sheasby professed fear of, that I think are fairly dealt with with the representations Mr. Colvin made.

So based on that totality, I'm going to deny any effective striking of the motion -- or striking of the report, rather, but I want the presentation of this witness handled in the way I've said. And I want a clean deposition with

responsive answers and not a verbal fistfight back and forth involved in generating that new deposition.

Again, I think there's fault on both sides, and I'm more worried about how we get a clean deposition to conduct this trial with than I am about sanctioning anybody.

That's my order. Are there any questions about it?

MR. SHEASBY: Nothing from Plaintiff, Your Honor.

MR. COLVIN: No, Your Honor; understood.

THE COURT: I think that's fair to both sides.

Both sides will bear their own costs as to this additional deposition. And Defendants will bear the cost of producing Doctor Robins to be here live in this building for that deposition.

All right. That completes Documents 212 and 213.

Let me just say this. There is a motion to bifurcate the trial as to willfulness. I'm going to deny that motion. The impact of willfulness has been significantly limited by the Court's prior rulings in the case. The Court's reviewed the motion carefully. The Court does not have the latitude to make two trials out of one, and any efficiencies or cost savings that might have previously and theoretically been achieved by a bifurcation, in my view and in light of the pretrial rulings and where this case stands, are of no real impact or value, and the cost and the expense of a second trial on a bifurcated willfulness issue is not warranted.

That motion is clearly denied.

Now, I have on my list, counsel, that there is a pending motion by the Plaintiff to supplement the reports of Kennedy and Brogioli. I assume that relates solely to the late-breaking deposition of Samsung's officer Mr. Lee? Is that the case?

MR. CORDELL: No.

THE COURT: I'm a little bit unclear as to what this motion to supplement is about.

MR. CORDELL: There are really two parts to it. I believe they do want to add some quotes from President Lee's deposition having nothing to do with why we had the deposition in the first place, but that is part of the motion. But Mr. Kennedy also wants to put in the Nanoco and KIP agreements as part of his damages analysis. He's very candid when he says that he's doing it to show that Samsung will pay a lot of money to settle one of these cases, and he's quite candid about that. It's in his reports.

THE COURT: Okay. Let me hear from the movant.

Mr. Sheasby, tell me what about those representations you disagree with and what else has not been stated is covered by your motion.

MR. SHEASBY: Sure.

So there is the late-breaking deposition of Doctor Kim that was supplementation. I didn't hear any argument from

counsel that that was improper. There also is a Nanoco and KIP agreement. I do believe that those are proper, but Mr. Kennedy is not going to rely on either Nanoco or KIP in his actual trial testimony, so that becomes moot.

THE COURT: Say that again.

MR. SHEASBY: So the President Lee issue is live, the supplementation of that deposition. As to Nanoco and KIP, those were argued yesterday and the Court denied the motion. But I can tell you just from a timing standpoint that Nanoco and KIP are not going to be presented to the jury.

THE COURT: So if they're not going to be -- if Kim and Nanoco aren't going to be presented to the jury, the scope of what this potential supplement would be?

MR. SHEASBY: President Lee's deposition.

THE COURT: No more, no less?

MR. SHEASBY: No more, no less.

THE COURT: Okay. I have some concerns about logistics, and I'll hear from both sides on this. I mean, if I grant this, even in a limited way, then there's going to be a supplemental report generated and then there's going to be a request to depose the expert on the supplemental report and then there's going to be requests for a responsive report from the competing experts, and then there's going a request for the Plaintiff to depose that expert on that supplemental responsive report.

We're going to pick a jury in two weeks, give or take. Is there a way to get this done in addition to Mr. Baxter and Ms. Truelove taking another deposition in the meantime.

MR. SHEASBY: Yeah. So it is baked. So President Lee's deposition was taken, the reports have been served by Mr. Kennedy and Doctor Brogioli, the supplemental reports, they were served before their depositions, so they have been already examined on these supplemental reports. This happened a long time ago, and so it's already been baked. There should be no back and forth.

If there was going to be a response to those reports, that would have already been done by Samsung's counsel because this happened quite a long time ago. The deposition was taken, the reports were immediately supplemented, they were deposed after the reports, so it's all done, Your Honor. This is just implementation of the fact that a deposition was ordered at the end of fact discovery.

THE COURT: All right. Let me hear from Mr. Cordell on this.

The cake is fully baked, as Mr. Sheasby says it is, or is it still in the oven.

MR. CORDELL: I think it's still in the oven. I think that they -- if you recall, we came here, President Lee's deposition was ordered on the narrowed topic of a set of meetings that happened in March of -- I forget the exact year.

THE COURT: With regard to the JDLA.

MR. CORDELL: Correct. Instead, what we had were a whole bunch of questions about other things, and the testimony that's relied here for damages purposes had nothing to do with those meetings. So if, in fact, they are permitted -- if Mr. Kennedy is now going to be permitted to rely on that testimony, then we have to have Mr. Meyer offer a rebuttal.

So Your Honor's vision of the cascade of events that will happen after this is, at least in part, accurate. It is true that we already had the supplemental reports of Kennedy when we took his deposition, but we opposed those supplements. We took the position that they were improper. But we'll suffer that, but we will ask for the right to have Mr. Meyer offer a rebuttal. It's pretty much that simple.

And I would offer another way out, which is this: In their motion, one of the things that Netlist points to is the idea that President Lee's testimony is redundant of other witnesses. So in particular for the Brogioli--is how I think you pronounce his name, with a soft G--Doctor Brogioli's report, they specifically say that another witness by the name of Jihwan Kim had offered the same testimony. So I'm not sure why we have to have President Lee's testimony on that subject at all.

THE COURT: Wasn't this deposition limited in time? I know I ordered it to be taken remotely. How long was the

deposition of Doctor Lee? I'm a little reluctant to call him President Lee, but --

MR. CORDELL: Doctor Lee is great, Your Honor.

It did run a little beyond the two hours. I was attempting to afford Mr. Sheasby some latitude so that he could ask whatever questions. Given that we all convened, I think we ended up at about two hours and 20 minutes on the tape. There was a little back and forth there as we were trying to get him to wrap it up, but, you know, I tried my best to preserve my record. And I did caution Mr. Sheasby that he was over his time several times, but he managed to wrap it up.

THE COURT: So tell me, Mr. Cordell, your suggestion other than just a 'no' to the Plaintiff here is what?

MR. CORDELL: Well, my suggestion is a 'no' to the Plaintiff that they themselves say they don't really need this testimony, and if they believe that, then they should stand pat on what they had before.

THE COURT: Well, and if the Court's not inclined to accept your suggested 'no', do you have an alternative proposal?

MR. CORDELL: I would ask for leave to file a short supplement to Mr. Meyer's report so that we can respond in particular to what Mr. Kennedy has said. With respect to Doctor Brogioli, I would have to go back and look at it. We

may need a short response. Again, I'm talking about a couple of pages from Mr. McAlexander in response. They supplemented two experts.

THE COURT: Okay.

MR. CORDELL: They supplemented Kennedy and Brogioli, so one on damages, one on infringement.

THE COURT: And the responsive experts on your side would be Meyer and McAlexander?

MR. CORDELL: I believe that's right. I'm sorry. Robins.

THE COURT: Okay. And the proposed supplemental reports from Kennedy and Brogioli have already been written and tendered.

MR. CORDELL: They have -- I think Mr. Kennedy submitted two different supplemental reports. He has a second and a third, I believe.

THE COURT: Well, I have no intention of allowing the supplemental reports of the Plaintiff's witnesses to bring in the Kim or Nanoco agreements. The only basis I would consider allowing it would be as to the Lee deposition.

Now, I haven't seen the proposed supplemental reports. Are they limited to the Lee deposition or are they proposed to go into these other areas that Mr. Sheasby's talked about with --

MR. CORDELL: So by my recollection, Your Honor, I

think Mr. Kennedy filed several supplemental reports, one of which is devoted to the Lee deposition.

THE COURT: So they are segregated into various reports?

MR. CORDELL: That's my recollection. I can double check.

THE COURT: That's correct, Mr. Sheasby?

MR. SHEASBY: Yes, Your Honor. If you just say nothing on KIP and Nanoco, it will be self-executing because we know what reports those are.

THE COURT: And if I were to grant responsive reports from the appropriate Defense witnesses, are you going to then ask me to take additional depositions on those responsive reports?

MR. SHEASBY: For Mr. Meyer, I don't think we need to. For Mr. Robins, it will be effective anyways. Doctor Robins, it will be effective anyways, because we are examining him again, so it won't require anything new. But for Mr. Meyer, definitely not.

THE COURT: In other words, you're suggesting that would be built into Ms. Truelove's deposition?

MR. SHEASBY: Yes.

MR. CORDELL: And it is the second supplemental of Mr. Kennedy served February 8, 2023, that deals with Doctor Lee's deposition.

THE COURT: Okay. All right. Well, I'll approve the supplemental report for Kennedy and Brogioli limited only to the deposition of Doctor Lee. I'll permit a responsive report from the two appropriate Defense experts of not more than five pages each; less if appropriate.

MR. CORDELL: Absolutely, Your Honor.

THE COURT: And I'll allow Plaintiff to inquire of Doctor Robins about any supplemental report or responsive supplemental report that he would author as a part of the deposition I've just ordered.

MR. CORDELL: Thank you, Your Honor.

THE COURT: Okay. With some trepidation, let me ask what else is outstanding before we get to the issue of motions in limine? Is either side aware of other motion practice that should be dealt with?

MR. SHEASBY: Your Honor, there was a trial subpoena issued to Doctor Indong Kim that we've moved to compel because they've declined to make him available. Doctor Kim resides in Northern California, he works for Samsung Semiconductor, Inc., and we believe he conducts business repeatedly in this district for, in fact, HPEE in this region and Dell, which is one district over, are both his largest clients. And so there's that motion that it has --

THE COURT: Do you think this falls within the purview of Rule 45?

MR. SHEASBY: We do, Your Honor, yes.

THE COURT: What's the response from the Defendants?

MR. McKEON: Good afternoon, Your Honor. Mike McKeon for Samsung.

We believe that this subpoena is invalid. And, first of all, you know, Mr. Kim -- or Doctor Kim, I should say, was deposed already in this case at length, so there is a healthy deposition transcript already. But Doctor Kim lives in and works in San Jose, California, and as we're his employer, Samsung Semiconductor, Inc., is also located.

And, of course, as Your Honor knows, Rule 45 requires that he reside in or be employed in or regularly transacts business in Texas. None of these things are the case here. Doctor Kim, I can tell you after talking to him about this, he came to -- last year came to Texas three times to visit two different customers. He went to Austin and Houston. They were overnight trips, and he went right back. And so this is not certainly fitting the definition of regularly transacting business in Texas; not at all, Your Honor.

He is a -- not a statutorial [sic] officer of SSI. There is -- he is one of 27 VPs, and he has a team of 40 people that report to him, but there's 1500 employees of the company. And so I think that's a factor here also.

THE COURT: And when you say SSI, Mr. McKeon, you mean Samsung Semiconductor, Inc.?

MR. McKEON: Yes, Your Honor. Thank you.

THE COURT: There are a lot of Samsung entities out there. I want to make sure I get the right one.

MR. McKEON: That's right, Your Honor.

And in terms of the undue burden, Your Honor, I mean, you know, he's a hard working gentleman. He reports to the office five days a week. Forty people report to him. He's very busy. This really would be a burden. You know, it's not just a question of coming in one day; it would be a week-long process. It is 1500 miles away, and he's very active in the business for --

THE COURT: Let me ask you this. Is he an executive officer of this Samsung entity, or is he a valued employee but doesn't hold an executive position?

MR. McKEON: Your Honor, he's a vice president. He's one of 27 vice presidents of the company.

THE COURT: Okay.

MR. McKEON: And so -- but he's not certainly, you know, in the C suite and CEO or CTO, that kind of level. So, you know, we wouldn't consider him to be at the level of Rule 45 in that regard. But also, Your Honor, you know, this would be an undue burden. And they had the deposition. There is nothing additional they need from him that they've demonstrated in the paperwork.

And I would say, too, Your Honor, as a technical matter

and technicality maybe, but it's certainly required under Fifth Circuit law, that when they tendered the subpoena, they gave him a $40 check, but under Fifth Circuit law and under the rule, they've got to give him compensation for travel, hotel, travel, and they have not done that.

So I think it's a matter of Fifth Circuit law. They haven't done -- as a technical matter, complied with the subpoena -- with the Rule 45 requirements. So you could quash on that basis alone.

But we think beyond that, Your Honor, substantively, they don't really need him, frankly, and we think it's an invalid subpoena based on the facts.

THE COURT: Now, what about the other targeted individual other than Mr. Kim, or is there someone beyond Mr. Kim? It's only directed to Mr. Kim?

MR. McKEON: That's right, Your Honor.

THE COURT: Okay.

MR. McKEON: Doctor Ji is another witness we are bringing to trial, and they've requested that Doctor Ji be called in their case, so we're going to work with them.

THE COURT: That's not a problem.

MR. McKEON: That's not an issue.

THE COURT: Okay. Okay. Do you have anything else, Mr. Sheasby?

MR. SHEASBY: Yeah.

I just wanted to say that this was not vexatious. He is an officer of SSI, he was the crucial individual who managed the relationship with Netlist, and he wrote the document saying the license was entered into because of infringement risk. So he actually is a key witness. They're trying to swap Hyun-Ki Ji out for him because they are just not wanting to have the key witness be present live in front of the jury.

And so this was not vexatious. We didn't ask about this for anyone else. He's an officer, his two biggest companies are HP and Dell, and who are in this district, and I think given those unique circumstances we do fall within the rules. I think --

THE COURT: HP and Dell are not in this district. They are in the Western District.

MR. SHEASBY: In the state. I apologize.

THE COURT: As the Federal Circuit has reminded me from time to time.

MR. SHEASBY: In the state.

I just want you to know this is not vexatious. There is a particular reason we're going this. His testimony would be incredibly impactful given live. The fact that they are flying a witness out from Korea to try to diminish Mr. Indong Kim's testimony I think tells you all you need to know about that. So we would urge for this unique situation the subpoena be granted.

Thank you.

MR. McKEON: May I respond to that?

THE COURT: Twenty seconds, Mr. McKeon.

MR. McKEON: They've requested that Doctor Ji be in their case already, so -- and he's coming to be in our case, but they've requested that he be here. He's not being swapped out for anything.

THE COURT: I understand.

All right. Given the facts in this particular situation, I am not persuaded that Mr. Kim meets the requirements of Rule 45. Three trips to Texas from California in a year is hardly a regular transacting of business in person under the rule.

The Court's given a fair amount of consideration to the Plaintiff previously by ordering President Lee to be deposed, albeit remotely.

I'm going to quash this subpoena. I'm not going to require the physical appearance and participation of Mr. Kim. I don't think it squarely fits within the purview of Rule 45. I'm not worried about the fact that Netlist hasn't put a check in his hands. That's the least of the problems. But I do have serious concerns about the regularly transacting business in person when there's three rather short trips from California to Texas over the course of a year.

And it is apparently uncontested that there's a lengthy video deposition of this witness available for presentation at

trial. Vexatious news is in the eye of the beholder, and I don't think Plaintiff crosses the threshold here with this application of Rule 45, so I'm going to reject Plaintiff's request to enforce the subpoena.

Said another way, I'm going to grant Defendant's request to quash the subpoena but, either way this gentleman's not being forced to appear live at trial. Okay?

All right. Anything else before we get to motions in limine?

MR. SHEASBY: Nothing for Plaintiff, Your Honor.

MR. CORDELL: Nothing from Defendant, Your Honor. Thank you.

THE COURT: All right. We're going to take a short recess. I want counsel, including lead counsel and local counsel, to discuss where we are on limines. If they haven't been substantially narrowed given all these pretrial rulings on dispositive motions, they ought to be.

As you're both aware, both sides are aware, the Court has implemented a relatively recent standing order on limine practice in patent infringement cases. I've already put forward more than 20 adopted limine guidelines and rulings. I would hope we can get what's left in dispute between both sides down to a very small number when I come back from recess. But I want you to meet and confer over it during the recess.

I'll try to make this about 10 or 12 minutes, and then we'll be back and see where we are.

Court stands in recess.

(Brief recess.)

THE COURT: Be seated, please.

All right. I understand we're down to nine disputed MILs. The one regarding RAND is off the table pursuant to the Court's prior rulings.

Let's turn to the disputed MILs. Obviously the standard MILs the Court's already imposed *sua sponte* pursuant to the standing order are in place.

Counsel, I don't see any reason to have extended argument on these; I just need you to get up to the podium and give me the real bottom line on what these are about.

Plaintiff's MIL No. 1 regarding RAND obligations, that obviously is denied because RAND is not in the case.

Let's go to Plaintiff's MIL No. 2.

Tell me what the Defendant's objection to it is.

MR. CORDELL: The Defendant's objections are, Your Honor, that the SK hynix agreement actually provides for the degree to which the supply commitment is relevant here. It's in the agreement itself. I can show Your Honor that at paragraph 5.3 of the supply agreement there is a black and white statement that says that nothing in this agreement shall be deemed to require either party to purchase or sell any

products from or to the other party unless and until a valid purchase order is issued and accepted.

Earlier in the agreement -- there are two agreements here. In the product purchase and supply agreement there is a commitment that the parties shall agree, but it's nothing more than an agreement to agree. This is in Section 2.1.

And then, finally, there is a most-favored-nations clause in the agreement at 7.12 that tells us, again in black and white, that the valuation of this agreement for most-favored-nations purposes, which is obviously a cue between the parties when they decide whether this agreement was worth more or less to some other, that when valuing this agreement they will be -- they will exclude any and all non-patent licensing terms such as any product supply, product development, joint venture, or lending agreements.

So in the text of the agreement in 7.12, it tells us that for the -- for valuation purposes we do not look at the supply commitment and, therefore, it is -- the MIL would violate what the agreement says in black and white.

THE COURT: Go ahead, Mr. Sheasby.

MR. SHEASBY: So my point is I think a little more prosaic, which is that Mr. Meyer is going to get on the stand and says that this product supply agreement is unenforceable, and that is a question of law for the sole determination of this Court. It is not appropriate for an expert witness to

say the supply clause is unenforceable, separate about whether it's valued for the purposes of a most-favored-nations clause is irrelevant. He's saying this is not obligatory; that there was no obligation to enter into product agreements even though the agreement on its face is obligatory.

THE COURT: Okay. Well, I think to be consistent with my prior ruling on the motion to strike regarding Mr. Meyer, I need to grant this, and I intend to grant Plaintiff's MIL No. 2. That doesn't mean I can't be persuaded to grant leave in real-time if there is a reason to present it during the course of the trial in the presence of the jury. That will be the Court's ruling.

Let's go to Plaintiff's MIL No. 3. This is to preclude Samsung presenting any argument or evidence that contradicts the findings of the Central District of California summary judgment and final judgment.

What's your objection to this, Mr. Cordell?

MR. CORDELL: So, Your Honor, we obviously hashed through this yesterday, and we appreciate and understand the Court's rulings. There is a problem in that there are elements of what happened in California that are directly relevant here.

For willfulness, for -- you know, the Plaintiff wants to use the existence of the JDLA; we want to be able to show the existence of the lawsuit and the judgment from the lawsuit to

show good faith. We both have nuggets that we want to be able to use.

What Your Honor did in the Acorn trial was instructive here. We took it up to a high level where it was an abstraction analysis. We took it to where the judgment was. And we would suggest the following: That the parties be allowed to say that there was a prior license, the JDLA--that's what I think the Plaintiff will want; that a dispute broke out, and that Netlist convinced a court to terminate that agreement as of a particular date. We would like to be able to say that the jury awarded no damages, or we could say $2, which is what the judgment was ultimately; and that the parties have appealed. And that's it. And that we don't get back -- we don't into the who shot John, we don't get into the back and forth, we don't have discussion about whether the parts were there on a Tuesday when they should have been there on a Sunday. That's all subsumed into the judgment. So just those basic facts--there was a license, there was a dispute, and Netlist successfully sued to terminate the license as of a particular date, they were awarded $2 or $0, and the parties have appealed.

THE COURT: So tell me how this is relevant to willfulness, Mr. Cordell.

MR. CORDELL: Well, Your Honor, because of the gravamen of at least one of our willfulness rejoinders is that

we believe the termination was improper, and that, in fact, we should still be licensed, and we had a good faith belief that we should still be licensed in several different time periods.

So obviously before the termination of the license, the Court's already dealt with that. We should be able to say that we felt like we had a good faith belief before the termination of the license. Remember that that termination, while made effective as of a particular date, wasn't made known to the parties until quite a bit later. Judge Scarsi entered judgment finding that the agreement was terminated at a much later date, during which time Samsung would have had a good faith belief, a justifiable good faith belief that we were still licensed.

And then after that period we feel like we have a meritorious appeal. They can quibble with this. They can say that it's not part of the overall landscape, but -- or that our defenses aren't meritorious, but that's a response. That doesn't mean that we shouldn't be able to talk about it.

It's also relevant to damages, Your Honor. We need to be able to express those periods in very precise forms so that the jury will understand what it is they're awarding damages for and how to calculate them.

THE COURT: What's your response, Mr. Sheasby?

MR. SHEASBY: So a couple of issues.

One, obviously elements of the JDLA and the dispute need

to come in, but they need to be carefully constrained. Of particular toxicity is a suggestion of the Court awarding zero damages or $1 damages. That was cover damages.

Let's go to deck 21 at slide 18, Mr. Huynh.

Judge Scarsi made clear that the only damages that were being dealt with by the jury were cover damages as opposed to consequential.

Deck 21, slide 18.

And so any discussion of the damages that were awarded in the previous case --

Slide 18.

-- would be improper because the judgment expressly said it did not include the consequential damages. We are not allowed to seek consequential damages.

The next issue --

If we could have the next slide.

Under Fifth Circuit law, a pending appeal is not a defense to willfulness or infringement. A case is res judicata in the Fifth Circuit pending appeal, and so they have no basis for using the appeal as subject to defensive willfulness.

The third issue, and I think this is dispositive --

If you go to slide 20.

-- is they actually served an interrogatory response on non-willfulness, Your Honor, and that interrogatory response

makes no reference to an appeal, makes no reference to a good faith basis as for the reason why they're shocked that they are being sued for infringement. All these things that are being talked about were not disclosed in fact discovery.

Go to the next slide.

We actually took the corporate representative on the topic of why Samsung is not willful, and they said that the only thing -- they refused to say that they had any involvement whatsoever. And so their 30(b)(6) witness on willfulness said nothing about these issues and the interrogatory response didn't either.

THE COURT: Just a minute.

Ms. Degnan, I can hear what you're saying up here. You need to whisper more quietly or you-all need to go some place where it doesn't get picked up by the microphone.

Go ahead, Mr. Sheasby.

MR. SHEASBY: So when I come back to it is obviously the JDLA and elements of the parties' behavior needs to come in, but this story about zero damages, it's inaccurate and highly prejudicial. The reference to the appeal is inaccurate and highly prejudicial, and any suggestion that they had a good faith belief in the JDLA being in place and them in compliance with it, that is not in any fact disclosure that they made.

This is not a situation where we sat on our hands and

didn't ask them. This is not a situation in which they refused to answer. They disclosed in an interrogatory and in their 30(b)(6) testimony the basis for willfulness -- non-willfulness, and it wasn't anything you're hearing right now.

What we would request is the following: There was a material breach of the license agreement for failure to supply product; the license agreement has been terminated; the termination is effective as of the date that Judge Scarsi ruled on. That's what we would request on this, Your Honor.

MR. CORDELL: Your Honor, a couple of rejoinders.

First of all, the JDLA itself bars consequential damages, so we don't have to go very far to get to the bottom of that, because the agreement tells us that consequential damages will not be awarded between these parties. Again, this Court is expert at contract applications. I won't belabor this.

THE COURT: Let me ask this. I assume that neither side contests the fact that the JDLA is going to be an exhibit in this trial.

MR. CORDELL: Correct.

MR. SHEASBY: Correct.

THE COURT: Okay.

MR. CORDELL: But here's the problem, Your Honor. Now what we're getting into now is essentially a motion for summary judgment. So this is supposed to be a MIL process,

not summary judgment.

With respect to the interrogatories --

THE COURT: Let me short circuit the process, gentlemen. Plaintiff's MIL 3 is subsumed by the Court's standing MIL No. 13 that precludes the parties from introducing evidence regarding other litigations. This is clearly an action in the Central District of California and it falls within MIL 13, so there's a grant here already in place.

Now, given that the JDLA is going to be an exhibit put before the jury during the trial, I think there is a modicum of context that needs to be supplied. I also think it would be better for both sides and it would be much fairer if that came from the Court and not from either side characterizing it.

I am happy to include in my preliminary instructions to the jury that in this case you are going to see a document called the JDLA, and it is no longer in force, it has been terminated, it was terminated by an order from another court in another district, that termination was effective for purposes of this trial. I think we've already established it was July the 15th, 2020. So I think we give them that date. I don't think there's anything more I need to say about it.

Whether it's a dollar in damages or a billion dollars in damages or whether it's on appeal or whether it's pending at the U.S. Supreme Court, I don't think I need to go further

than that.

MR. SHEASBY: We agree, Your Honor.

THE COURT: And I would welcome a proposed inclusion in those preliminary instructions from both sides that I'll consider, and then from those I'll determine what I think is appropriate or not. And then if anybody wants to go into it beyond what I'm willing to instruct the jury on, then you're going to have to approach and get leave because this is already covered by MIL 13.

MR. SHEASBY: Thank you, Your Honor.

MR. CORDELL: It is, Your Honor. But if I could just belabor it for one more moment, and let me just make clear that it is not correct that we didn't set forth this -- these defenses in our interrogatory responses. They are -- they span roughly 20 pages. And in it we specifically lay out the timeline that you may recall Ms. Degnan showing us earlier where there are different periods for purposes of willful infringement, for example.

So the -- it's not simply the termination as of July 15, 2020, that will be relevant. The jury will need to hear that Judge Scarsi, for example, didn't enter judgment until January 10, 2022. So there are --

THE COURT: Put that in your proposed submission.

MR. CORDELL: Okay. We will. Thank you.

THE COURT: Let me have that within the next 24

hours by email to my staff.

MR. CORDELL: Just for clarity, Your Honor, we're not getting into the who shot John.

THE COURT: I don't see any probative value in doing that.

MR. CORDELL: Excellent. Thank you.

THE COURT: This is, in effect, a limited exception to MIL 13 simply because I think it's pretty clear that this document is going to be, if not a central part, a material part of the evidence in the case.

MR. SHEASBY: I'll just flag, this is not their non-willfulness interrogatory response.

THE COURT: Let me say this again, counsel. You submit to me what you want. I'll look at it. I'll instruct the jury to a level that I think is appropriate. If you want to go into this topic beyond what I am going to instruct the jury on, you're going to have to approach leave and convince me to grant leave.

MR. SHEASBY: I understand, Your Honor.

MR. CORDELL: Understood, Your Honor.

THE COURT: That is Plaintiff's MIL No. 3, which technically I'm going to deny because it's already subsumed by standard MIL No. 13.

Let's take up Plaintiff's MIL No. 4. This seeks to preclude Samsung from presenting any evidence that Netlist

has failed to comply with JEDEC obligations.

All right. Mr. Cordell, what's the Samsung objection here?

MR. CORDELL: With your permission, Your Honor, Mr. Mosteller will present our arguments.

THE COURT: That's fine.

MR. MOSTELLER: Thank you, Your Honor. Mr. Mosteller for Samsung.

And I think that Your Honor has already decided Plaintiff's MIL No. 4 several times over the past couple of days with reference to other JEDEC discussions and Netlist's relations with JEDEC and the scope of what Samsung will be able to present at trial with respect to that relationship.

Here, Your Honor, we're dealing with Netlist's obligations to disclose things to JEDEC. The fact that JEDEC was sitting in on the meetings--I'm sorry--that Netlist was sitting in on the meetings at JEDEC, was silent despite its obligations to disclose patents is highly relevant. It's relevant to whether Samsung knew about the patents, it's relevant to whether Samsung or Netlist thought that products that practiced these standards were implicated by these patents, and it's relevant to Netlist's perception of these patents to products that practice those standards, Your Honor.

THE COURT: Mr. Sheasby?

MR. SHEASBY: Sure.

These patents are not standard essential. There's no obligation to disclose a patent that's not standard essential. There's no allegation from a legal -- in front of the jury that goes to whether we complied with our obligations to JEDEC, because we have no obligations to JEDEC because they're not essential. The MIL should be granted.

If there's some narrow exception to talking about the relationship with JEDEC, they should seek leave to approach. But on its face that we failed to comply with JEDEC obligations, that's a reasonable MIL. There is no allegation in front of the jury that we haven't failed to comply with our JEDEC obligations. It's not pled in the complaint. It's not pled in the affirmative defenses.

THE COURT: All right. In regard to Plaintiff's MIL No. 4, I'm going to grant this MIL, but I want the parties to understand why I'm doing it. I'm granting it so that the Court can be an active gatekeeper on what does and doesn't come in before this jury about JEDEC. I have a fear that without any constraints there's going to be confusion created, especially in light of the fact that these are not standard essential patents, and there are not obligations under JEDEC because they're not standard essential patents submitted to this standard setting body. But there may be relevant use of JEDEC and there may be implications that relate to other live issues in the case.

And, quite honestly, I don't trust either side to hew the proper line without me being a gatekeeper here. So I'm going to grant it. That doesn't mean there's not going to be any mention of it during trial; it's going to be that I'm going to get to supervise and provide oversight for what's mentioned how it's mentioned in real-time. Okay?

MR. MOSTELLER: Thank you, Your Honor.

THE COURT: All right. Plaintiff's MIL 4 is granted with that explanation.

Plaintiff's MIL 5 seeks to preclude Samsung from introducing evidence that practicing the standard is a defense to infringement or willfulness.

What's the dispute here, Mr. Cordell? These are not standard essential patents.

Well, whoever's speaking for Samsung. I have to look up if I'm going to notice who's at the podium.

MR. LIVEDALEN: Thank you, Your Honor. Brian Livedalen for Samsung.

So this issue's slightly different. I think it comes in the most clarity here with respect to willfulness. So Samsung, being a member of JEDEC and being a company that practices a standard, relies on the disclosures we talked about. Netlist has an obligation also as a member to disclose patents that it believes may be standard essential. And so Netlist for its willfulness case is going to rely on the

disclosure of related patents as evidence that Samsung knew of those patents and knew that it was infringing.

In rebuttal --

THE COURT: Wait a minute. Infringing related patents and not the patents-in-suit?

MR. LIVEDALEN: I should clarify. Infringing the patents-in-suit simply by disclosure of related patents.

THE COURT: Say that again.

MR. LIVEDALEN: Sure.

THE COURT: You believe the Plaintiff's going to show infringement of the patents-in-suit by Samsung because Samsung may have disclosed other patents to JEDEC but not these patents?

MR. LIVEDALEN: So Netlist --

THE COURT: I'm trying to follow you.

MR. LIVEDALEN: Yeah, sure.

Netlist is going to argue that the disclosure of parent patents, for example, to the '918 and '054 Patents put Samsung -- gave Samsung the requisite knowledge of the '918 and '054 Patents. They may also argue that the disclosure of a parent patent, for example, and the disclosure to that patent was standard essential; also provided Samsung with knowledge or an intent that it is somehow infringing the '918 and '054 Patents.

And so what Samsung would like to do in response is to

say that JEDEC members rely on what is disclosed, but they also rely on what is not disclosed. So the fact that there was never a disclosure of the '918 and '054 Patents as standard essential is something that they can rely on for no willfulness.

THE COURT: Okay. Maybe that's where the confusion lies. I can see the nexus with the willfulness issue. I don't see the nexus with the underlying infringement issue. The jury's going to determine infringement by comparing the claims -- the elements of the asserted claims to the accused products; no more, no less.

But beyond the determination of infringement, I do see what you're arguing with regard to willfulness. I thought you were arguing with regard to infringement. Maybe you are. I don't see the connection on the infringement --

MR. LIVEDALEN: The issue I just argued, Your Honor, is with respect to the willfulness issue.

THE COURT: Okay.

MR. LIVEDALEN: With respect to infringement, we're not going to argue that practicing the standard alone immunizes Samsung from infringement. I think that's what the gist of MIL No. 5 is asking. But we do want to present evidence that Netlist seeks to argue that its patents are simultaneously infringed and not essential, and that argument that Netlist is going to present at trial is inconsistent with

statements that Netlist made both to the Court in its original complaint and positions that it made in its interrogatory responses.

THE COURT: Well, if you want to attempt to impeach a witness for the other side, you can attempt to impeach them, but that doesn't really move the needle on whether I ought to grant or deny this motion in limine.

MR. LIVEDALEN: Yes, Your Honor. That's the one point with respect to infringement; again, but we're not going to argue that simply practicing a standard immunizes us from infringement.

THE COURT: Okay. Mr. Sheasby, the allegation is that you're going to try to show willfulness here because either Samsung knew these were essential and refused to submit them to the standard, or somehow as being involved with the standard Samsung knew about -- I mean, knew about this technology and hid it from the standard body. Respond to that somehow for me.

MR. SHEASBY: Sure.

So two issues. One, willfulness is a state of mind. What you just heard is a creation of the very able lawyers who are representing Samsung. We took a 30(b)(6) witness on willfulness. We have an interrogatory on willfulness. None of what you just heard has been disclosed anywhere in fact discovery. And so they're going to be constrained by what

they say in fact discovery.

So if their corporate representative gets up and says, Oh, we're shocked, we're shocked because it wasn't in JEDEC, I'm going to approach the bench show, I'm going to show that it's not in any discovery response that they gave us, it contradicts what they said, and then I'm going to ask for it to be stricken. That's the first issue.

The second issue is that we don't intend to say that these patents are essential and, therefore, the issue that -- this convoluted theory that they're spinning -- that we're spinning to is not one that I think is going to be -- even open the door.

But I think the thing to do is not to debate it. The MIL should be granted. If they have some convoluted theory that they want to try to present because they think the door's been opened, they should approach and they can discuss that theory with the Court.

That's my response.

THE COURT: What else from Samsung?

MR. LIVEDALEN: So one thing, Your Honor. Mr. Sheasby faults Samsung for not providing an adequate interrogatory response, but all that Netlist put in its --

THE COURT: You know, counsel. We are awful late in the day for me to be hearing about discovery disputes. I mean, we cannot spend the entire trial with each of you

pointing fingers and casting accusations at the other for failing to do their job a year and a half ago as this case was being developed and discovery was begun and pursued. That's disconcerting for both of you to be focusing on old discovery failures. The time for addressing those is long gone. We did have a raft of motions to compel in this case. I've dealt with them as best I could, but I'm growing weary of hearing about discovery disputes.

MR. SHEASBY: Your Honor, I think it's reasonable. This is not a discovery dispute. They did answer the interrogatory. They should be constrained by what the answer was.

THE COURT: Well, let me say this. To the extent Samsung might be tempted to communicate to the jury that by practicing the standard they're somehow being immunized from an infringement allegation, I think that's improper. And even though Samsung says they're not going to do this, to that extent I'm going to grant this MIL.

Willfulness is a totality of the circumstances, as the Supreme Court's made clear. I don't think there's a basis upon which I can keep the Plaintiff from addressing the standard setting body and Samsung's participation in it with regard to the willfulness issue.

I'm going to grant it as to precluding the Defendant from trying to defend against infringement on the basis of

practicing the standard, but in other regards -- any other regard I'm going to deny the motion.

MR. LIVEDALEN: Thank you, Your Honor.

THE COURT: Okay. Let's go to Defendant's disputed motions in limine. The first one seeks to prevent improper argument or evidence regarding Samsung's alleged failure to supply memory to Netlist.

What's Plaintiff's objection with Defendant's MIL No. 1?

MR. SHEASBY: There will be no discussion of the damages associated with the breach, or lack thereof, but the story of the behavior between Samsung and Netlist goes directly to the willfulness allegations in the case.

If we can have the PX Indong Kim exhibit, Mr. Huynh.

THE COURT: And the breach you're talking about is what gave rise to the termination?

MR. SHEASBY: That's correct; the cutting off of our supply.

So what the record will show is that in June of 2016, as we were getting ramped up and asking for them to increase our supply --

Scroll down. Scroll down some more.

-- the Samsung entities had what is described as a right of first refusal to our patents, which means that if we went bankrupt we would have to give them our patents.

Scroll down more.

They knew at that time --

Scroll up a little more.

-- we were in financial straits.

Scroll up. Stop right there.

We knew at the time that we were in financial straits. They were aware of the importance of our patents to RDIMM, HBM, and SSD. Then what happened --

And if I could have the elmo.

-- is Samsung asked us to agree to an additional technology transfer. That's this email right here. This is in December of 2017, in which they asked us to give them additional technology for just -- without agreeing to full production or joint production with us. We declined that option.

After we declined that option, what happened was that Samsung --

And let's go to PX 1673.

After we declined that option, internally the most senior individuals at Samsung made the decision to effectively cut off our supply. And there's actually internal emails where people are writing LOL, laughing out loud, about our cutoff in supply.

Put that down, Mr. Huynh. If we go to PX 1786, please.

Sorry. We're having a technical difficulty here, Your Honor. Here it is.

So after we did that, we decline to participate in their technology transfer, this is an internal Samsung email where they said they get zero allocation, and they said --

Scroll down.

They told us they get zero allocation. And the internal email between Samsung and Netlist says that they recognize that Samsung was a hundred percent of Netlist's business.

And if you give me one moment, I'll find that slide, Your Honor. If you -- I'll get it after the argument. But there was a slide that I showed yesterday in which Samsung was apprised that if they cut off supply, it was a hundred percent of our business. They knew that if they cut off our supply, there was a real risk that we went bankrupt; and if we went bankrupt, they got our patents.

And so we don't intend to argue at all that there was damages from this behavior, but we do absolutely intend to tell the willfulness story in which Samsung knew that they had rights to our patents if we went bankrupt, they evaluated the value of our patents, as you saw in that document, they asked us to enter into another technology collaboration, we denied -- we refused to do so, and they then terminated or stopped our product supply, recognizing that it's a hundred percent of our business.

I believe this issue has been settled by *Monsanto versus Ralph*, this is 382 F.3d 1374, in which this exact issue came

up. The jury was -- affirmed for considering in the case of damages that Ralph misappropriated a license to use the seed without restrictions, inconsistent with their obligations under the agreement.

So I have no doubt that this is going to be need to be constrained by the Court. My expectation is that we will need to make a proffer for Your Honor to see exactly what we're going to say. But this is a live issue. There is facts that are prejudicial, but undue prejudice is different from prejudice.

If you turn -- let's go back to PD Indong Kim.

THE COURT: Let me ask you this, Mr. Sheasby. The point in time that you're talking about where there was a, as you put it, refusal to supply memory from Samsung to Netlist, which subsequently resulted in the termination of the JDLA, obviously those events took place. The refusal or the failure to supply memory, as you put it, obviously that took place at a time when the JDLA was in force, both sides were licensed. So as egregious as the conduct may be, in your view, the willfulness issue is not willfulness in general; it's willfulness as to infringement.

How can that relate to infringement at a time when both sides were licensed to the totality of each side's patents?

MR. SHEASBY: Sure.

THE COURT: And I understand the foundry products

may be an exception to this. But outside of that one exception, how does anything that would be covered by the license in the JDLA, how is that a willfulness of infringement?

MR. SHEASBY: So *Monsanto versus Ralph* is the answer, Your Honor, which is that the willfulness is that they tried to end our business to get our patents. That's the willfulness, because they knew they infringed our patents. They had a first line of collateral. This is their internal document analyzing it.

And scroll down. Scroll up a little. Sorry.

Then right after they say they have a first line of collateral to our assets, they then talk about the fact that we have patents covering their technology. We then decline to participate in a joint development with them, and then they cut off our supply, recognizing that was a hundred percent of our business, showing that we could go bankrupt, and at that point they would own our patents.

And so I understand it's prejudicial, Your Honor, but it's not undue prejudice, and the reason why it's not undue prejudice is because it shows behavior, it shows scienter. They knew they needed our patents, they knew they had the right to our patents if they went bankrupt, and so they could have it both ways. They could both cut off our supply, bankrupt us, get our patents, and at the same time avoid the

fact that there was infringement on RDIMM HBM, and Enterprise SSD. And, oh, by the way, there is a live, meaningful issue as to whether HBM was even licensed. So at least as to HBM this is highly relevant.

So I understand Your Honor's trepidation of trying a case within a case, but these are real facts that go to recklessness and go to scienter. Not generic scienter. This document refers specifically to our patents and to the patents that we -- and the products we are accusing of infringement. And they say it's a first line of collateral assets for their IP. So they're saying the jury can find that they knew they were using our HBM patents, they knew that they would get them if we went bankrupt, and they attempted to bankrupt us.

So I'm very sensitive that this should not be a sideshow, but this is a real defense. This is a real allegation that's not just a recapitulation of the Central District case. And I'm happy to make a proffer to Your Honor so it can be constrained, but I cannot agree to the entire MIL as granted.

THE COURT: All right. Well, I think I've heard enough on this. I am persuaded that the Court is going to have to oversee any use of this and there's going to have to be justification given in advance and leave granted before this path is gone down.

I'm going to grant the MIL. That doesn't mean I can't be persuaded that there are some circumstances and there are some

ways and means by which this argument can come in. I don't intend to keep the Plaintiff from telling its story. I do intend to keep the Plaintiff from making up a story that doesn't have any support and is entirely intended to inflame the jury. And as I sit here right now, I can think of numerous ways where this could be improperly used. And there may be some ways in which it can be properly used, but I cannot sit here in advance and give you-all clear guidance. I'm going to have to be an active gatekeeper here, and on that basis I'm going to grant the MIL.

MR. SHEASBY: And Your Honor, we will make a written proffer in advance so that it's not chaotic.

THE COURT: I granted the MIL. You better.

MR. SHEASBY: My point is I'm not going to come up to the stand -- to the bench in the middle of a direct and say now I want to do this. I'm going to in advance --

THE COURT: I need some notice, Mr. Sheasby. I need some notice on this.

Okay. Let's go to Defendant's MIL No. 2, which seeks to prevent improper argument or evidence regarding the alleged value of the supply agreement to Netlist under the JDLA.

MR. SHEASBY: So this should be denied, Your Honor. Mr. Kennedy has already done his calculation and was argued yesterday about what the value of the supply agreement in the JDLA is. It's an essential part of his damages analysis, and

he should be entitled to make it. What we should not be able to say is we should not be able to say the value of the supply agreement in the JDLA somehow equates to damages.

So there's two separate issues. The valuation of the JDLA that Mr. Kennedy did involved the value of the supply agreement. That is properly before the Court. The Court granted it -- heard extensive argument on it yesterday and it was granted. The separate issue is should we be allowed to say that represents our damages, we should not be able to say that represents our damages.

THE COURT: Mr. Cordell.

MR. CORDELL: There's no proper purpose for this, Your Honor. We had a summary judgment argument yesterday and the Court entered summary judgment that we weren't going to get into the who shot John on the JDLA. That's what this is. Mr. Kennedy actually offers a $2 billion opinion based on what he says are the damages flowing from the breach of the JDLA. That should have been done in California. It wasn't. That's not for us to revisit.

The -- if we're going to go down this road, if we're going to get into the damages issues surrounding the JDLA, then we have to be able to tell this jury that the other jury also heard this evidence and came up with zero dollars. So those two are inextricably linked. If we're going to visit this in the first instance, it's going to be a problem.

I think what Mr. Sheasby was suggesting -- well, there might be a middle ground where Mr. Kennedy could discuss the overall value of the JDLA and whether or not the $8 million represented the true valuation of that agreement. That might be fair game on both sides. But that's a different matter altogether, and it doesn't get into these consequential damages issues, it doesn't get into any of the many tails of this dog, but would be a fairly confined analysis. And if it were that, then perhaps we can see a way forward.

But even under those circumstances, when we leave the four corners of the JDLA and these experts are then allowed to wander around and talk about the valuation of the supply agreement and whether it was fulfilled or not, then we end up back with the jury finding that it was worth zero dollars.

MR. SHEASBY: And to be clear, a lot of what Mr. Cordell said is absolutely right. The Court has already constrained it based on the MTSs that were heard yesterday. What Mr. Kennedy is going to say and the only thing he's going to say is the value of the supply clause over the term of the agreement was X amount. He's not going to say that's our damages from breach; he's not going to say that it had anything to do with the breach.

So a MIL that we can't talk about what our damages were from the breach is perfectly fine. A MIL that is a surreptitious way of precluding Mr. Kennedy from being able to

see what the value of the JDLA was, which was in his report and was approved by the Court, is we don't think proper.

MR. CORDELL: This is in no way surreptitious, Your Honor. This is an extension of the summary judgment ruling you entered yesterday. That's what it is. And if we're going to get into this, if we're going to -- if Mr. Kennedy is going to take the stand and say that JDLA was worth $2 billion, then we're going to have to get into all of the facts that surrounded that dispute in California culminating in a jury verdict of zero, because that's -- the body -- I mean, the jury that heard that evidence made a decision, and if we're going to get into the underlying facts, then they ought to at least have the punchline.

MR. SHEASBY: So, Your Honor, those are two completely different issues. Mr. Kennedy is not going to talk about the damages from breach at all. He's going to talk about what the value of the -- the enterprise value of the agreement was to Netlist at the time of the hypothetical negotiation. Mr. Meyer has a contradictory opinion. It has nothing to do with damages, which were just covered damages in the Central District of California. This is an attempt to just reargue what was lost yesterday in front of Mr. Kennedy.

THE COURT: All right. I've heard enough.

On Defendant's MIL No. 2, as to Mr. Kennedy I'm going to deny this motion in limine. I don't intend to try to

constrain his expert report and his expert testimony, which has already been through the refining fire of *Daubert* practice. What has survived in his report he's entitled to testify to without this limine interfering with it.

As to anyone besides Mr. Kennedy, I'm going to grant this motion in limine and cause is going to have to be shown and leave is going to have to be granted before these topics are gone into. This is not going to be a second *Daubert* motion on Mr. Kennedy to try to limit him beyond what's already been limited, but I do appreciate that others who have not been through the *Daubert* process might take advantage of an open door here, and so I'm going to be a gatekeeper as to anyone else.

MR. CORDELL: May I ask for clarification, Your Honor? Does that involve -- does that also apply to my witness Mr. Meyer who is responding to Mr. Kennedy's opinions about the JDLA?

THE COURT: I don't intend to hamstring the Plaintiff's expert on damages or the Defendant's expert on damages beyond what I've already done through the *Daubert* practice.

MR. CORDELL: Thank you.

THE COURT: I don't intend the MIL practice to be another layer of *Daubert* exercising -- or excising. So the experts are going to be permitted to testify what's in their

reports without the application of this MIL being an impediment. Anyone else besides Kennedy and Meyer on this topic is going to have to get leave before they go into it.

MR. CORDELL: Thank you.

THE COURT: Okay?

Defendant's MIL No. 3 has to do with the Korean tax tribunal opinion.

What's the Plaintiff's objection here?

MR. SHEASBY: So the Korean tax tribunal decision was a decision --

THE COURT: This is about whether the $8 million in the JDLA was an engineering fee or --

MR. SHEASBY: Correct, Your Honor.

THE COURT: -- compensation for licenses.

MR. SHEASBY: Correct.

And so the issue is, is that Judge Scarsi found -- Judge Scarsi adopted the Korean tax tribunal's ruling. And so if they were going to contend and argue that the $8 million was a royalty and not an NRE, that would open the door to the KTT coming in. And so that's the issue we face.

There's a separate issue, and I don't think it's relevant, but there were a number of documents exchanged between the parties in the past that don't -- that were not the KTT opinion but ultimately filtered up to the KTT, which I don't think are the subject of the MIL, and I want to make

sure did not get inadvertently swept in the MIL. And I can show you an example of that. But to shortcut --

If you go to 21-30 -- -28.

THE COURT: Let me -- while you're doing that, let me observe that this MIL is a good exemplar of something that really doesn't help the Court at all. The MIL as stated is "improper argument or evidence regarding the Korean tax tribunal opinion." Well, 'improper' is in the eye of the beholder. It doesn't tell me anything. I need to know substantively what you're trying to keep out and what you're trying to allow in. And this is really counterproductive to have this kind of vague designation on limine practice.

And both sides -- I'm not singling out Samsung; I'm just taking this opportunity to make the observation. I think both of you are equally guilty.

But we need to be specific here, because I can guarantee you when these issues come up with a jury in the box in a room full of people, you're going to need specific guidance and I'm going to need to know specifically what you're asking to do. And flashing an order that you granted that improper argument couldn't come in doesn't help anybody; it just preserves more consternation and dispute for a later time, and we need to be as precise and substantive here as we can be. That's where these limine orders really serve a valid purpose.

So enough preaching on that issue but bear it in mind.

MR. SHEASBY: So there are two issues that we're going to use the KTT for. One, is that it shows that the NRE fee was not a royalty, and --

THE COURT: Let me ask you this, Mr. Sheasby. Why does the KTT need to come in at all? I mean, if you're going to show it's not, they're going to try to show that it is, and if they are going to try to show that it is, you're going to try to show that it is not. What is the relevance of it without somebody taking the initiative to throw it in the middle of the room and make an issue out of it?

MR. SHEASBY: Yeah. So the answer is the following, which is in the Central District of California there was a ruling that the $8 million was an NRE and not a royalty and it was an improper withholding. And so for them to now contest that, which is what they plan to do--their expert says it's a royalty, it's not an NRE--that's a violation of both issue preclusion by the Central District of California, and it contradicts this KTT ruling that was adopted by Judge Scarsi.

So I would prefer it not to be in at all. The problem is I can't refer to Judge Scarsi's decision adopting it, so the best -- the next best solution I have is to -- this is not the agreement. This is some random correspondence between the parties that's not subject I don't think to this MIL at all, unless I hear otherwise. But I don't know what else to do to defend against an argument that the $8 million was a royalty

other than the KTT if they're going to just contradict Judge Scarsi.

THE COURT: Well, this is not your limine motion; this is Samsung's limine motion trying to keep reference to the KTT out.

MR. SHEASBY: I understand.

THE COURT: And you're standing here telling me that you want it in; you want me to deny their MIL so you can put it in --

MR. SHEASBY: Yes.

THE COURT: -- counter to what you're just telling me.

MR. SHEASBY: No, I want to deny the MIL so that I can put it in, because Mr. Meyer says that the $8 million fee was a royalty, not an NRE. That's why it's relevant and the MIL should be denied. I'll keep it that short.

THE COURT: Well, if it gets in, both sides are going to get to characterize it in the way they think is proper, but it seems to me that -- it seems to me both sides are better off if it doesn't come in.

Now, maybe Samsung wants to wave it around as a royalty to try and show the jury $8 million is a drop in the proverbial bucket compared to the ask of the Plaintiff in this case. I've seen that happen before. And then you're going to want to show the jury that, no, it's not a royalty, it's an

engineering fee, and it shouldn't be considered as a downward drag on your damages model.

I mean, if we're going to have that fight, I guess we're going to have that fight, but I'm not going to tie one side's hands and leave the other side's hands loose and open.

MR. SHEASBY: So, Your Honor, the fight is joined. Mr. Meyer says it's a royalty, not a non-recurring engineering fee. It's in his report. So as soon as he says it, it's going to open the door to the KTT decision. Frankly, it's a violation of issue preclusion by Judge Scarsi, but putting that to one side it's going to open the door --

THE COURT: Now, I have a lot of respect for Judge Scarsi, but this is my case --

MR. SHEASBY: Yeah, I know.

THE COURT: -- it's not his case, and we're not going to try every issue that he's dealt with over again, and I'm not necessarily going to take everything that he's decided in California as gospel here. And you-all are sadly mistaken if you think that that's the end-all be-all as far as the same issues that are live before this jury.

MR. SHEASBY: I understand, Your Honor. And I --

THE COURT: Are you intending to raise this issue, Mr. Cordell or Ms. Degnan, with Mr. Meyer?

MS. DEGNAN: If I could respond.

So Mr. Meyer is talking about this $8 million payment.

He's talking about how Samsung viewed it as it entered into the agreement as part of his overall damages analysis, and Samsung did view it as a royalty and not something else.

THE COURT: So are you by this very vague and, you know, unhelpful language in the MIL "improper argument regarding the KTT opinion," are you telling me you want Mr. Meyer to be able to tell the jury that it's an engineering fee but keep Netlist from saying it's a royalty, or do you want to say it's a royalty and keep them from saying it's an engineering fee?

MS. DEGNAN: We are happy for them to say it's an engineering fee. That is what their view of it is. And they can use evidence besides this opinion from a Korean tax agency. And we're saying they want to kind of buttress their own argument about how it was not a royalty by referring to a decision by the Korean government no different than if we were to say we think our invalidity arguments are really good because the PTAB had an institution decision.

So they want to use this Korean decision by an agency to kind of give an -- you know, strength, buttress up their own argument that it wasn't a royalty.

And I must -- I want to just say two things, Your Honor. I know you can read the California decision, but the California decision did not, in our opinion, adopt the Korean tax authority's reasoning and decision. The California

decision only said there was a breach because Samsung withheld taxes and the Korean authority decided they weren't supposed to withhold taxes. There hasn't been, in my opinion, in our opinion, an adoption -- wholesale adoption by the California court.

And issue preclusion, Your Honor, has four prongs, and the fourth prong is that there are no special circumstances that would render preclusion inappropriate or unfair. And we think having this Korean decision come in as giving a sort of an attaboy to their own argument would be incredibly unfair. And so we think the Korean decision should come out.

If they want to argue it was not a royalty and that Samsung shouldn't have thought it was a royalty, they actually can use the JDLA itself--the JDLA talks about the language--but they don't need to kind of resort to a Korean decision or the California decision to say, Hey, we're right, we were right and Samsung is wrong when they believed and treated it as a royalty. So this decision shouldn't come in under any circumstances.

But again, Mr. Meyer is focused on how Samsung understood it and treated it as a royalty; not saying that we understand that ultimately a Korean tax authority did not think there should be withholding on it because it did not consider it a royalty.

THE COURT: Okay. Okay.

Well, the consideration of this disputed issue and the issuance of an opinion by the Korean tax tribunal is prior litigation and it's covered by MIL 13, and so this is basically subsumed by standard MIL 13 already. But for the purpose of redundancy, clarity, and lack of misunderstanding, I'm going to grant Defendant's MIL 3.

MS. DEGNAN: I do just want to say one more thing, if you don't mind, Your Honor.

THE COURT: It depends on what it is.

MS. DEGNAN: Mr. Sheasby had mentioned there are two other documents he showed you on the screen. They'll be dealt with as exhibits later on today. They are very similar to this issue, because every time they talk about these other two documents they want to say they are especially important because they were used, they think, in connection with this Korean tax proceeding. We think that also falls under MIL 13 as well.

THE COURT: Well, we're not arguing exhibits now, and --

MS. DEGNAN: He did mention it just now, Your Honor.

THE COURT: And if he wants to say they are particularly important, that's no different than you arguing that we strenuously believe this motion for summary judgment be granted. I mean, that doesn't help me at all.

MR. SHEASBY: And, Your Honor, I understand the

grant. I just want to flag, we are going to put a written proffer in advance of Mr. Meyer's testimony so the Court can understand how the door will be opened.

THE COURT: I'm sure that that's probably going to happen. And again, if I'm going to have to determine in advance whether this comes in and how it comes in or not at all, I need as much notice as I can.

I will say this, counsel. I don't intend my time on the bench during this jury trial to be an ongoing interpretation by me of Judge Scarsi's opinions. What he did adopt, what he didn't adopt, what he meant by this language, what he didn't mean by that language, I'm not the Court of Appeals for the Ninth Circuit, and don't want to be, so I am concerned that I'm going to get trips to the bench throughout this trial with one or the other of you telling me, Judge Scarsi said this, you have to do that, or Judge Scarsi adopted this, you have to let us tell the jury he adopted this, or something else in that vein, and I'm just not intending to be an *ad hoc* appellate review of Judge Scarsi's prior actions as a part of this trial.

MR. SHEASBY: I understand the instruction, Your Honor.

THE COURT: Well, I wish I believed that that was not going to come up, but I am convinced it probably is.

Let's go to Defendant's MIL No. 4. This seeks to

preclude, again, improper allegations--so that's very helpful--improper allegations of price fixing or other illegal activities.

You know, tell me what in the heck you're talking about here. This is not helpful at all.

MR. McKEON: Sorry. Apologize. Mike McKeon for Samsung, Your Honor.

THE COURT: If it keeps up, I'll just deny them all.

MR. McKEON: We need to improve that language and it won't happen again, I assure you of that.

You know, during one of the depositions of the 30(b)(6) witnesses at Netlist, he started calling Samsung a part of a cartel and, you know, using that kind of language. So this is really kind of the impetus to this motion.

And as I understand from the brief that was filed by Netlist, they generally agree they're not going to use, you know, language of 'corrupt' and 'evil' and 'dishonest', but they didn't specifically mention 'cartel'. So I want to make sure that's part of their agreement.

They also, Your Honor, say they're not going to call Samsung an oligopolies, but that's exactly what the excerpt from Mr. Kennedy, in his report that's exactly what he does. And he talks about, you know, price control, significant control over pricing --

THE COURT: Let me stop you, Mr. McKeon. The

Court's standard MIL No. 9 says, "The parties shall be precluded from introducing evidence, testimony, or argument referring to any person or entity as greedy, corrupt, evil, dishonest, or using any other pejorative term. The parties shall also be precluded from introducing evidence, testimony, or argument that characterizes any other person or entity and their actions as stealing, copying, misappropriating, pirating, trespassing, or similar terms."

Now, what relief do you need that's not already covered there?

MR. McKEON: Your Honor, particularly the 'oligopolies' term.

THE COURT: I mean, I can't put Webster's Dictionary in here.

MR. McKEON: Well, there's a notion that Samsung and two other entities controlled the pricing in the memory market, and as Your Honor can take judicial notice, memory is like wheat and oil; it just fluctuates due to the market conditions of supply and demand. So he makes this allegation that we're fixing prices and that's --

THE COURT: All right. This MIL is denied. Standard MIL 9 is in place. If you think there's a violation of it, you can raise it with me at the time it happens.

MR. McKEON: Thank you, Your Honor.

THE COURT: These standard MILs, ladies and

gentlemen, are not just a waste of time. I put a lot of thought into these. I don't need new MILs urged that tread the same ground or the same territory. It's just an insult to the Court and a waste of everybody's time.

Defendant's MIL No. 4 is denied.

All right. Defendant's MIL No. 5, "No argument, evidence, or question suggesting a party's corporate representative at trial is obligated to prepare on any particular topic or is charged with knowledge of others within the company."

I need some explanation here. What's the concern that Samsung has by urging this? And then I'll hear what Netlist's objection is.

MS. REARDON: Sure. Kate Reardon for Defendants, Your Honor.

THE COURT: Please go ahead.

MS. REARDON: Now, Your Honor, we filed from making suggestive comments to the jury that Samsung's corporate representatives at trial somehow have an obligation to testify on a specific topic or that they're at fault because they should be charged with knowledge of the entire company and every topic that might be relevant to the case.

Now, as Your Honor knows by listening to the arguments over the last two days, this has been a very complicated case involving two Samsung corporate entities, three different

product lines, we have 14 Samsung witnesses testify, I believe 12 of those were corporate representatives in response to 94 different topics in Netlist's corporate notice to Samsung. So it's impossible now that a Samsung witness that testifies at trial should be deemed knowledgeable of all of those topics in Netlist's corporate notice. And that's what Netlist says it expects. Netlist's opposition is very clear that they intend to cross examine Samsung's corporate representatives and exploit any corporate deposition testimony that was given.

You know, corporate representatives at trial don't get a notice of all the topics that they're intended to testify about. And Netlist contends that corporate testimony can be used for any purpose, and our position is that the testimony played at trial, the testimony can be--I'm sorry--played at trial or can be used to impeach a witness with their own prior inconsistent statement, not statements of 14 different Samsung witnesses that testified.

So, in our view, allowing this practice would be prejudicial because you can't sit and cast aspersions that Samsung as a corporation doesn't know something because the corporate witness can't testify to all the topics that Netlist asks Samsung's witnesses about.

THE COURT: You're not asking me to constrain an examination of a 30(b)(6) witness on a matter or topic they were noticed on or deposed on, are you?

MS. REARDON: Certainly not, Your Honor. They are constrained by their own testimony. But to constrain them to the testimony of all the other Samsung witnesses would not be -- would be prejudicial and unfair.

THE COURT: You're asking me to prohibit inquiry into areas that the witness has no knowledge of and wasn't noticed on simply for the purpose of making the witness look like they're inept or unprepared or incompetent. Is that what you're doing?

MS. REARDON: That's exactly correct, Your Honor.

THE COURT: Let me hear a response from Netlist, please.

MS. ZHAO: Thank you, Your Honor. Yanan Zhao for Plaintiff Netlist.

We want to be clear that Netlist does not intend to tell the jury any of the witnesses from Samsung has an legal obligation to be prepared to testify to specific topics or any other's knowledge. That's not our intention at all. But we do want to clarify that Federal Rule of Evidence 32(a)(3) is very clear that Netlist can use Samsung's corporate witness' testimony made during depositions at trial for any purpose--to examine the witness, to impeach the witness, and to (unintelligible) decision to put 12 corporate representatives during the deposition, and it's Samsung's own choice to put whatever numbers of witnesses they are going to put at trial.

So those 12 corporate representatives testified during deposition on behalf of Samsung, so it's not their personal knowledge but the collective knowledge of the company. And every single testimony they made is just all in the transcript. It's crystal clear. It's nothing surprising to opposing side; it's a fair game to allow Netlist to cross those trial witnesses based on their corporate representatives' previous testimony.

THE COURT: Okay. Anything further?

MS. ZHAO: That's it, Your Honor.

THE COURT: Anything further from Samsung?

MS. REARDON: Your Honor, all I would say is that they are certainly available to make depo designations and play testimony of corporate representatives, but to have a witness on the stand that's responsible for being cross examined with the testimony of 14 different witnesses is completely prejudicial to Samsung and their ability to prepare a witness for trial.

THE COURT: All right. Thank you, ladies.

With regard to Defendant's MIL No. 5, with regard to any topic upon which a particular 30(b)(6) witness has previously been noticed, I'm going to deny the MIL.

With regard to any topic that the particular Samsung 30(b)(6) witness has not been previously noticed, I'm going to grant the MIL, and there's going to have to be a request for

leave to go into something that this witness was not noticed on and was not deposed on. Okay?

I mean, with an international corporation the size of Samsung, it is not unreasonable that there be multiple representatives for different facets of the company's operation, and it's unfair to try and charge every 30(b)(6) witness with notice of every aspect of the operations of the company.

MR. SHEASBY: I just want to clarify something, Your Honor. The person they're planning on calling was not a witness on any topic. So they're just going to randomly call someone who was not ever designated as their corporate representative for the express purpose of being avoided from being cross examined on the bad testimony.

THE COURT: You want to address that, Mr. Cordell?

MR. CORDELL: Only, Your Honor, that they knew who our witnesses were and they chose to depose whom they chose. I don't really understand the objection. If he says something inconsistent, whether he could be impeached with someone else's testimony doesn't end the inquiry. Mr. Sheasby's very capable of arguing in closing that he contradicted what the corporate rep said.

THE COURT: Okay. Let me revisit my ruling because it was limited by its terms to persons designated as 30(b)(6) witnesses for Samsung. With regard to persons who have been

previously designated on any topic or topics as 30(b)(6) witnesses on behalf of Samsung, they're not going to be cross examined on matters beyond what they were noticed for, without prior leave of the Court.

For witnesses who would testify at trial on behalf of Samsung who were never noticed or deposed as a 30(b)(6) witness, if the Plaintiff seeks to cross examine them on a lack of knowledge regarding Samsung or any aspect of its operations, you're going to have to approach the bench and tell me why and why that's fair and get leave before you do it. And if there's a reasonable basis for it, I'll grant leave, and if there's not, I won't grant leave.

MR. CORDELL: Understood, Your Honor. Thank you.

THE COURT: All right. That appears to be the last limine matter.

Now, it's 25 minutes after 5:00. There are quite a few disputes regarding exhibits to be pre-admitted. Here's what I'm going to do, counsel. I told you two days ago now that I had yesterday and today, and that's what I have. I'm going to refer you to Magistrate Judge Payne on any disputes regarding the pre-admission of exhibits. You will schedule a subsequent time to appear before him and take up the exhibit issues we have not reached today. All other pretrial rulings on the dispositive motions and the limine issues that I've given you will stand, and I intend to memorialize those in writing as

soon as practical. That should complete the pretrial process.

We do have a late-breaking deposition to be taken. I've already given clear instructions on how that's to be done. The jury will be assembled and ready to go on the 14th of April.

Are there other things you're aware of that I need to take up with you before that time?

MS. TRUELOVE: We just had one question, Your Honor, in regards to how Your Honor intends to conduct the bench trial; if we need to be prepared to go into that directly upon the jury retiring to deliberate on the infringement.

THE COURT: Quite honestly, the jury's verdict is going to, in large part, inform whether there's going to be a need for a bench trial, and I can't tell you what that verdict's going to be at this point.

It's unlikely that after the kind of trial that I'm anticipating I will have the energy or the stamina or the desire to immediately go into the bench trial. It's unlikely.

MS. TRUELOVE: Understood.

THE COURT: It may be moot by whatever the effect of the jury's verdict is, so I'm going to reserve answering that question until we have the benefit of knowing what the verdict is.

I will say this. Unless I tell you otherwise, you do not need to be prepared to immediately go into the bench trial.

If that changes, I'll let you know. If there is going to be a bench trial necessary based on the verdict the jury renders, I don't intend to push that off for an inordinate period of time, but I don't at this point think it's likely that I will want to immediately accept the verdict, dismiss the jury, and begin the bench trial. Okay?

MS. TRUELOVE: Thank you, Your Honor.

THE COURT: Anything from Samsung, Your Honor.

MR. CORDELL: No, Your Honor. Thank you. Appreciate the time and energy this week.

THE COURT: Is LSU still in March Madness?

MR. CORDELL: They are, if you follow the lady's team.

THE COURT: I do because you stole our coach from Baylor.

MR. CORDELL: I was hoping you wouldn't point that out, but she's done great things, so we appreciate it.

THE COURT: Well, it just takes money.

All right. That completes the pretrial proceedings, other than the exhibits which will be taken up by the magistrate judge, as I've indicated. Thank you for your time and attention, counsel. You're excused.

And the Court stands in recess.

(The proceedings were concluded at 5:30 p.m.)

I HEREBY CERTIFY THAT THE FOREGOING IS A CORRECT TRANSCRIPT FROM THE RECORD OF PROCEEDINGS IN THE ABOVE-ENTITLED MATTER. I FURTHER CERTIFY THAT THE TRANSCRIPT FEES FORMAT COMPLY WITH THOSE PRESCRIBED BY THE COURT AND THE JUDICIAL CONFERENCE OF THE UNITED STATES.

S/Shawn McRoberts 03/31/2023

_____DATE_____
SHAWN McROBERTS, RMR, CRR
FEDERAL OFFICIAL COURT REPORTER

US010949339B2

(12) **United States Patent**
Lee et al.

(10) **Patent No.: US 10,949,339 B2**
(45) **Date of Patent: *Mar. 16, 2021**

(54) **MEMORY MODULE WITH CONTROLLED BYTE-WISE BUFFERS**

(71) Applicant: **Netlist, Inc.**, Irvine, CA (US)

(72) Inventors: **Hyun Lee**, Ladera Ranch, CA (US); **Jayesh R. Bhakta**, Cerritos, CA (US)

(73) Assignee: **NETLIST, INC.**, Irvine, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **15/470,856**

(22) Filed: **Mar. 27, 2017**

(65) **Prior Publication Data**

US 2017/0337125 A1 Nov. 23, 2017

Related U.S. Application Data

(63) Continuation of application No. 13/970,606, filed on Aug. 20, 2013, now Pat. No. 9,606,907, which is a (Continued)

(51) **Int. Cl.**
G06F 12/00 (2006.01)
G11C 8/12 (2006.01)
(Continued)

(52) **U.S. Cl.**
CPC ***G06F 12/00*** (2013.01); ***G11C 5/025*** (2013.01); ***G11C 5/04*** (2013.01); ***G11C 5/066*** (2013.01); ***G11C 8/12*** (2013.01)

(58) **Field of Classification Search**
None
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,249,253 A 2/1981 Gentili et al.
4,571,676 A 2/1986 Mantellina et al.
(Continued)

FOREIGN PATENT DOCUMENTS

CN 102576565 B 9/2015
JP 09237492 9/1997
(Continued)

OTHER PUBLICATIONS

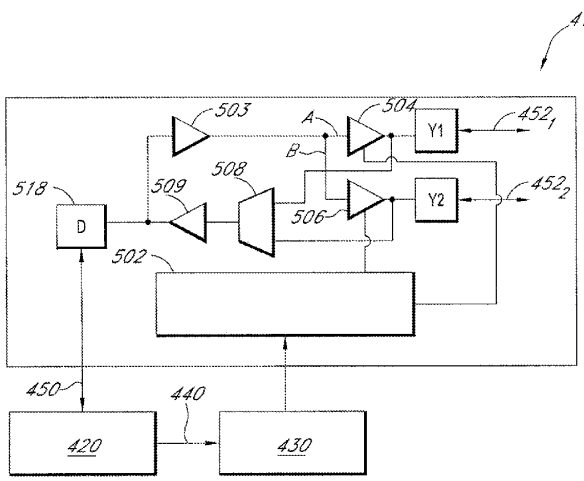
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Patent Owner's Preliminary Response, filed Apr. 9, 2017.
(Continued)

Primary Examiner — Edward J Dudek, Jr.
Assistant Examiner — Ralph A Verderamo, III
(74) *Attorney, Agent, or Firm* — Morgan, Lewis & Bockius LLP

(57) **ABSTRACT**

A memory module is configured to communicate with a memory controller. The memory module comprises DDR DRAM devices arranged in multiple ranks each of the same width as the memory module, and a module controller configured to receive and register input control signals for a read or write operation from the memory controller and to output registered address and control signals. The registered address and control signals selects one of the multiple ranks to perform the read or write operation. The module controller further outputs a set of module control signals in response to the input address and control signals. The memory module further comprises a plurality of byte-wise buffers controlled by the set of module control signals to actively drive respective byte-wise sections of each data signal associated with the read or write operation between the memory controller and the selected rank.

35 Claims, 13 Drawing Sheets



Related U.S. Application Data

continuation of application No. 12/761,179, filed on Apr. 15, 2010, now Pat. No. 8,516,185, which is a continuation-in-part of application No. 12/504,131, filed on Jul. 16, 2009, now Pat. No. 8,417,870.

(51) **Int. Cl.**

G11C 5/02	(2006.01)
G11C 5/04	(2006.01)
G11C 5/06	(2006.01)

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,592,011	A	5/1986	Mantellina et al.	
4,739,473	A	4/1988	Ng	
5,060,188	A	10/1991	Zulian et al.	
5,313,624	A	5/1994	Harriman	
5,345,412	A	9/1994	Shiratsuchi	
5,463,590	A	10/1995	Watanabe	
5,513,135	A	4/1996	Dell et al.	
5,532,954	A	7/1996	Bechtolsheim et al.	
5,537,584	A	7/1996	Miyai et al.	
5,541,448	A	7/1996	Carpenter	
5,572,691	A	11/1996	Koudmani	
5,617,559	A	4/1997	Le et al.	
5,630,096	A	5/1997	Zuravleff	
5,649,159	A	7/1997	Le et al.	
5,655,113	A	8/1997	Leung	
5,717,851	A	2/1998	Yishay et al.	
5,724,604	A	3/1998	Moyer	
5,729,716	A	5/1998	Lee et al.	
5,784,705	A	7/1998	Leung	
5,802,541	A	9/1998	Reed	
5,905,401	A	5/1999	Sher	
RE36,229	E	6/1999	Cady	
5,953,280	A	9/1999	Watsui	
5,973,392	A	10/1999	Senba et al.	
6,011,710	A	1/2000	Wieggers	
6,061,754	A	5/2000	Cepulis	
6,070,217	A	5/2000	Connolly et al.	
6,141,245	A	10/2000	Bertin	
6,173,357	B1	1/2001	Ju	
6,188,641	B1	2/2001	Uchida	
6,205,516	B1	3/2001	Usami	
6,223,650	B1	5/2001	Stuck	
6,260,127	B1	7/2001	Olarig et al.	
6,262,938	B1	7/2001	Lee	
6,438,062	B1	8/2002	Curtis	
6,446,158	B1	9/2002	Karabatsos	
6,480,439	B2	11/2002	Tokutome et al.	
6,546,476	B1	4/2003	Gillingham	
6,553,449	B1	4/2003	Dodd et al.	
6,618,791	B1	9/2003	Dodd et al.	
6,636,446	B2	10/2003	Lee	
6,683,372	B1	1/2004	Wong et al.	
6,704,910	B2	3/2004	Hong	
6,717,855	B2	4/2004	Underwood	
6,721,860	B2	4/2004	Klein	
6,747,887	B2	6/2004	Halbert et al.	
6,788,592	B2	9/2004	Nakata et al.	
6,832,303	B2	12/2004	Tanaka	
6,948,084	B1	9/2005	Manapat et al.	
7,024,518	B2 *	4/2006	Halbert	G06F 13/1689 711/101
7,093,066	B2	8/2006	Klein	
7,130,308	B2	10/2006	Haddock et al.	
7,133,960	B1	11/2006	Thompson et al.	
7,191,302	B2	3/2007	Usami	
7,225,303	B2	5/2007	Choi	
7,289,386	B2	10/2007	Bhakta et al.	
7,334,150	B2	2/2008	Ruckerbauer et al.	
7,379,361	B2	5/2008	Co et al.	
7,464,225	B2	12/2008	Tsern	
7,532,537	B2	5/2009	Solomon et al.	
7,730,254	B2	6/2010	Risse	
7,865,674	B2 *	1/2011	Gower	G06F 13/1684 365/189.02
7,881,150	B2	2/2011	Solomon et al.	
7,916,574	B1	3/2011	Solomon et al.	
7,990,746	B2	8/2011	Rajan	
8,001,434	B1	8/2011	Lee et al.	
8,081,536	B1	12/2011	Solomon et al.	
8,089,795	B2	1/2012	Rajan	
8,130,560	B1	3/2012	Rajan et al.	
8,189,328	B2	5/2012	Kanapathippillai	
8,233,303	B2	7/2012	Best	
8,244,971	B2	8/2012	Rajan	
8,250,295	B2	8/2012	Amidi et al.	
8,335,894	B1	12/2012	Rajan	
8,417,870	B2	4/2013	Lee et al.	
8,516,188	B1	8/2013	Solomon et al.	
8,689,064	B1	4/2014	Lee et al.	
8,756,364	B1	6/2014	Bhakta et al.	
8,782,350	B2	7/2014	Lee et al.	
8,856,464	B2	10/2014	Karamcheti	
9,606,907	B2 *	3/2017	Lee	G11C 5/025
2001/0008006	A1	7/2001	Klein	
2002/0048195	A1	4/2002	Klein	
2002/0112119	A1	8/2002	Halbert et al.	
2003/0070052	A1	4/2003	Lai	
2004/0098528	A1	5/2004	Janzen	
2005/0010737	A1	1/2005	Ware et al.	
2005/0257109	A1	11/2005	Averbui	
2005/0281096	A1	12/2005	Bhakta et al.	
2006/0117152	A1	6/2006	Amidi et al.	
2006/0262586	A1	11/2006	Solomon et al.	
2006/0277355	A1	12/2006	Ellsberry et al.	
2007/0070669	A1	3/2007	Tsern	
2007/0293094	A1	12/2007	Aekins	
2008/0025137	A1	1/2008	Rajan et al.	
2008/0046631	A1	2/2008	Takaku et al.	
2008/0104352	A1	5/2008	Talbot	
2008/0162790	A1	7/2008	Im	
2009/0103387	A1	4/2009	Shau	
2009/0228631	A1 *	9/2009	Marulkar	G06F 13/1605 711/100
2009/0248969	A1	10/2009	Wu et al.	
2010/0070690	A1	3/2010	Amer et al.	
2010/0091540	A1	4/2010	Bhakta et al.	
2010/0125681	A1	5/2010	Patel	
2011/0090749	A1	4/2011	Bhakta et al.	
2011/0125966	A1	5/2011	Amidi et al.	

FOREIGN PATENT DOCUMENTS

JP	H10092169	A	9/1997
JP	2010320270		12/1998
JP	2000285674		10/2000
JP	2000311485	A	10/2000
JP	2002184176		6/2002
JP	2003007963		1/2003
JP	2008046989		2/2008

OTHER PUBLICATIONS

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1003, 'Declaration of Dr. Harold Stone,' filed Dec. 22, 2017.

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1009, 'JEDEC Standard Double Data Rate DDR SDRAM Specification, JESD79 Jun. 2000,' filed Dec. 22, 2017.

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1010, 'JEDEC Standard 21-C, DDR SDRAM Registered DIMM Design Specification Jan. 2002,' filed Dec. 22, 2017.

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1011, 'JEDEC Standard DDR2 SDRAM Specification, JESD79-2B Jan. 2005,' filed Dec. 22, 2017.

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1012, 'Declaration of John J. Kelly Regarding

(56) **References Cited**

OTHER PUBLICATIONS

Records of Joint Electron Device Engineering Council JEDEC,' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1022, 'Decision Denying Institution of Inter Partes Review of U.S. Pat. No. 8,516,185, *SanDisk Corp.* v. *Netlist, Inc.*, IPR2014-01029, Paper No. 11 Dec. 16, 2014,' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1023, 'Decision Denying Institution of Inter Partes Review of U.S. Pat. No. 8,516,185, *Smart Modular Techs. Inc.* v. *Netlist, Inc.*, IPR2014-01369, Paper No. 12 Mar. 9, 2015,' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1024, 'Excerpts from the Hearing in Certain Memory Modules and Components Thereof, and Products Containing Same, Inv. No. 337-TA-1023 May 8, 2017,' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1025, 'Complainant Netlist, Inc.'s Initial Post-Hearing Brief, Certain Memory Modules and Components Thereof, and Products Containing Same, Inv. No. 337-TA-1023 May 30, 2017 excerpts relevant to '185 patent,' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1026, 'Respondents' Post-Hearing Brief, Certain Memory Modules and Components Thereof, and Products Containing Same, Inv. No. 337-TA-1023 May 30, 2017 excerpts relevant to '185 patent,' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1027, 'High-quality versions of demonstrative graphics included in Respondents' Post-Hearing Brief (Ex.1026),' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1028, 'Complainant Netlist Inc.'s Reply Post-Hearing Brief, Certain Memory Modules and Components Thereof, and Products Containing Same, Inv. No. 337-TA-1023 (Jun. 9, 2017) (excerpts relevant to '185 patent),' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1029, 'Respondents' Reply Post-Hearing Brief, Certain Memory Modules and Components Thereof, and Products Containing Same, Inv. No. 337-TA-1023 (Jun. 9, 2017) (excerpts relevant to '185 patent),' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1030, 'High-quality versions of demonstrative graphics included in Respondents' Reply Post-Hearing Brief (Ex. 1029),' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1031, 'Institution of Inter Partes Review of U.S. Pat. No. 8,516,185, *SK hynix Inc. et al.* v. *Netlist, Inc.*, IPR2017-00577, Paper No. 8 (Jul. 7, 2017),' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1032, 'Final Written Decision, *Diablo Techs., Inc.* v. *Netlist, Inc.*, IPR2014-00882, Paper No. 33 (Dec. 14, 2015),' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1033, '*Netlist, Inc.* v. *Diablo Techs., Inc.*, No. 2016-1742 (Fed. Cir. Jul. 25, 2017),' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1034, 'Netlist's Infringement Claim Chart for U.S. Pat. No. 9,606,907 Jun. 14, 2017,' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1035, 'Stone, H.S. Microcomputer Interfacing, Reading, MA: Addison Wesley, 1982,' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1039, 'Intel E7525 Memory Controller Hub (MCH) Chipset Datasheet (Jun. 2004),' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1040, 'Initial Determination, Certain Memory Modules and Components Thereof, and Products Containing Same, Inv. No. 337-TA-1023 (Nov. 14, 2017) (redacted excerpts),' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1041, 'Bruce Jacob et al., Memory System: Cache, DRAM, Disk (2008) (excerpts),' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1042, 'Direct RDRAM datasheet (2000),' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00363, Petition for Review of U.S. Pat. No. 9,606,907, filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00363, Preliminary Response, filed Apr. 10, 2018.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00364, Petition for Inter Partes Review of U.S. Pat. No. 9,606,907, filed Dec. 27, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00365, Petition for Inter Partes Review of U.S. Pat. No. 9,606,907, filed Dec. 27, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Petition for Inter Partes Review, filed Dec. 30, 2016.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Patent Owner's Preliminary Response, filed Apr. 9, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Patent Owner's Response, filed Sep. 15, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Petitioners' Objections to Evidence, filed Sep. 22, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Petitioners Motion to Exclude Certain Inadmissible Testimony of Patent Owners Expert Carl Sechen, filed Jan. 9, 2018.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Petitioners Request for Oral Hearing, filed Jan. 9, 2018.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, PO Request for Oral Argument, filed Jan. 9, 2018.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, PO Motion for Observations, filed Jan. 9, 2018.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Petitioners' Response to Patent Owner's Motion for Observations, filed Jan. 23, 2018.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, PO's Opposition to Petitioner's Motion to Exclude, filed Jan. 23, 2018.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Petitioners' Reply in Support of Motion to Exclude Certain Inadmissible Testimony of Patent Owner's Expert Carl Sechen, filed Jan. 30, 2018.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Petitioners' Updated Exhibit List, filed Feb. 9, 2018.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Patent Owner's Updated Exhibit List, filed Feb. 9, 2018.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Trial Instituted Document, filed May 15, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Order Trial Hearing, filed Jan. 25, 2018.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Hearing Transcript, filed Mar. 22, 2018.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Final Written Decision, filed May 3, 2018.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1002, 'Excerpts of FH for U.S. Pat. No. 8,756,364,' filed Dec. 30, 2016.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1003, 'Declaration of Harold Stone,' filed Dec. 15, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1007, 'JEDEC JESD79 publication (Jun. 2000),' filed Dec. 30, 2016.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1008, 'JEDEC Declaration for DDR Specification,' filed Dec. 30, 2016.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1015, 'Texas Instruments 74LS245 datasheet (2002),' filed Dec. 30, 2016.

(56) **References Cited**

OTHER PUBLICATIONS

Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1017, '2000-03 Samsung CMOS SDRM data sheet,' filed Sep. 15, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1019, 'Final Decision [IPR2014-01011],' filed Sep. 15, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1020, 'Final Decision [IPR2014-00883],' filed Dec. 30, 2016.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1023, 'Gordon_Moore_1965_article,' filed Dec. 30, 2016.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1024, 'Deposition Transcript of Carl Sechen (Dec. 8, 2017),' filed Dec. 15, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1025, 'Supplemental Declaration of Harold S. Stone (Dec. 15, 2017),' filed Dec. 15, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1026, 'Xilinx CoolRunner XPLA3 CPLD product specification (2000),' filed Dec. 15, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1027, 'Xilinx Programmable Logic Design Quick Start Handbook (2004),' filed Jan. 9, 2018.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1028, 'Xilinx TQFP package datasheet (2000),' filed Jan. 9, 2018.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1029, 'Xilinx VQFP package datasheet (2002),' filed Feb. 9, 2018.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1030, 'Xilinx CS280 package datasheet (1999),' filed Dec. 15, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1031, 'Petitioners Demonstratives,' filed Apr. 9, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 2001, 'Exhibit 2001,' filed Dec. 15, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 2002, 'Exhibit 2002,' filed Dec. 15, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 2003, 'Exhibit 2003,' filed Feb. 9, 2018.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 2004, 'Exhibit 2004,' filed Dec. 30, 2016.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 2005, 'Exhibit 2005,' filed Sep. 15, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Patent Owner's Preliminary Response, filed Apr. 10, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Petitioners' Updated Exhibit List, filed Jun. 21, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Patent Owner's Response, filed Oct. 25, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Petitioners' Objections to Evidence, filed Oct. 31, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Petitioners' Reply to Patent Owner's Response, filed Feb. 6, 2018.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Patent Owner's Objections to Evidence, filed Feb. 13, 2018.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Petitioners' Request for Oral Hearing, filed Feb. 28, 2018.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Patent Owner's Request for Oral Argument, filed Feb. 28, 2018.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Patent Owner's Motion to Exclude, filed Feb. 28, 2018.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Petitioners' Opposition to Patent Owner's Motion to Exclude, filed Mar. 14, 2018.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Patent Owner's Reply to Petitioner's Opposition to Motion to Exclude, filed Mar. 21, 2018.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Notice of Accord Filing Date, filed Jan. 10, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Trial Instituted Document, filed Jul. 7, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Order—Request for Oral Argument—37 CFR 42.70, filed Apr. 2, 2018.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 1003, 'Declaration of Harold S. Stone,' filed Jan. 5, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 1007, 'JEDEC JESD79 publication (Jun. 2000),' filed Jan. 5, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 1009, 'TI 74LS245 datasheet (2002),' filed Apr. 10, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 1010, 'JEDEC Declaration,' filed Jan. 5, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 1011, 'IPR2014-01369 Aug. 23, 2014 Petition for Inter Partes Review,' filed Jan. 5, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 1012, 'IPR2014-01029 Jun. 24, 2014 Petitioner for Inter Partes Review,' filed Apr. 10, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 1025, 'US Patent Application Publication 2011/0016250,' filed Jun. 21, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 1026, 'Netlist's post-hearing brief,' filed Jun. 21, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 1027, 'Hynix's post-hearing brief,' filed Oct. 25, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 1028, 'Demonstrative Graphics from hynix's post-hearing brief,' filed Oct. 25, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 1029, 'Netlist's reply post-hearing brief,' filed Oct. 25, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 1030, 'hynix's reply post-hearing brief,' filed Jun. 21, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 1031, 'Demonstrative Graphics from hynix's reply post-hearing brief,' filed Jun. 21, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 2001, 'Exhibit 2001,' filed Jun. 21, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 2002, 'Exhibit 2002,' filed Jun. 21, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 2003, 'Exhibit 2003,' filed Oct. 25, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 2004, 'Exhibit 2004,' filed Oct. 25, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 2005, 'Exhibit 2005,' filed Dec. 30, 2016.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 2006, 'Exhibit 2006,' filed Dec. 30, 2016.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 2007, 'Exhibit 2007,' filed Dec. 30, 2016.
Anonymous. (Dec. 1996). "Applications Note: Understanding DRAM Operation," IBM, 10 pages.
Altera, ACEX iK, Programmable Logic Device Family, Data Sheet, May 2003, Ver 3.4.
Behrens, S. "HP Printer Memory Explained", The ZonkPage, Last Updated Jan. 21, 2004. Accessed Feb. 10, 2013, Retrieved from the Internet: URL<http://warshaft.com/hpmem.htm>. 7pp.
English Translation of the Notice of Grounds for Rejection, Korean Patent Application No. 2012-7004038, dated May 11, 2016.
Examination Report, European Patent Application No. 10730021.2, dated Apr. 14, 2014.

(56) **References Cited**

OTHER PUBLICATIONS

Response to Examination Report, European Patent Application No. 10730021.2, dated Jun. 4, 2014.
Examination Report, European Patent Application No. 10730021.2, dated Apr. 29, 2015.
Inter Partes Review Case No. IPR2014-01029, Petition for Inter Partes Review of U.S. Pat. No. 8,516,185, filed on Jun. 24, 2014.
Inter Partes Review Case No. IPR2014-01029, Exhibit 1008 to Petition for Inter Partes Review, "Declaration of Charles J. Neuhauser, Ph.D. under 37 C.F.R. § 1.68," filed on Jun. 24, 2014.
Inter Partes Review Case No. IPR2014-01029, Supplemental Petition for Inter Partes Review of U.S. Pat. No. 8,516,185, filed on Jul. 23, 2014.
Inter Partes Review Case No. IPR2014-01029, Patent Owner's Preliminary Response pursuant to 37 C.F.R. § 42.107, filed on Oct. 17, 2014.
Inter Partes Review Case No. IPR2014-01029, Decision Denying Institution of Inter Partes Review 37 C.F.R. § 42.108, issued Dec. 16, 2014.
Inter Partes Review Case No. IPR2014-01369, Corrected Petition for Inter Partes Review of Claims 1-19 of U.S. Pat. No. 8,516,185, filed on Sep. 22, 2014.
Inter Partes Review Case No. IPR2014-01369, Exhibit 1008 to Corrected Petition for Inter Partes Review, "Declaration of Dr. Nader Bagherzadeh under 37 C.F.R. § 1.68," filed on Sep. 22, 2014.
Inter Partes Review Case No. IPR2014-01369, Exhibit 1013 to Corrected Petition for Inter Partes Review, "Webster's II New College Dictionary," filed on Sep. 22, 2014.
Inter Partes Review Case No. IPR2014-01369, Exhibit 1014 to Corrected Petition for Inter Partes Review, "Standard Dictionary of Electrical and Electronics Terms," IEEE 1988, filed on Sep. 22, 2014.
Inter Partes Review Case No. IPR2014-00882, Corrected Petition for Inter Partes Review of U.S. Pat. No. 7,881,150, filed on Jul. 8, 2014.
Inter Partes Review Case No. IPR2014-00882, Exhibit 1007 to Petition for Inter Partes Review, "Declaration of Dr. Srinivasan Jagannathan," filed on Jun. 22, 2014.
Inter Partes Review Case No. IPR2014-00883, Corrected Petition for Inter Partes Review of U.S. Pat. No. 8,081,536, filed on Jul. 8, 2014.
Inter Partes Review Case No. IPR2014-00883, Exhibit 1011 to Petition for Inter Partes Review, "Declaration of Dr. Srinivasan Jagannathan," filed on Jun. 21, 2014.
Inter Partes Review Case No. IPR2014-01011, Corrected Petition for Inter Partes Review of U.S. Pat. No. 7,881,150, filed on Jul. 8, 2014.
Inter Partes Review Case No. IPR2014-01011, Exhibit 1007 to Petition for Inter Partes Review, "Declaration of Dr. Srinivasan J Jagannathan." filed on Jun. 22,2014.
Inter Partes Review of U.S. Pat. No. 7,881,150, IPR Case No. IPR2014-00882, Patent Owner's Preliminary Response Pursuant to 37 C.F.R. § 42.107, filed Oct. 7, 2014.
Inter Partes Review of U.S. Pat. No. 7,881,150, IPR Case No. IPR2014-00882, Decision—Institution of Inter Partes Review 37 C.F.R. § 42.108, issued Dec. 16, 2014.
Inter Partes Review of U.S. Pat. No. 8,081,536, IPR Case No. IPR2014-00883, Patent Owner's Preliminary Response Pursuant to 37 C.F.R. § 42.107, filed Oct. 7, 2014.
Inter Partes Review of U.S. Pat. No. 8,081,536, IPR Case No. IPR2014-00883, Decision—Institution of Inter Partes Review 37 C.F.R. § 42.108, issued Dec. 16, 2014.
Inter Partes Review of U.S. Pat. No. 7,881,150, IPR Case No. IPR2014-01011, Patent Owner's Preliminary Response Pursuant to 37 C.F.R. § 42.107, filed Oct. 7, 2014.
Inter Partes Review of U.S. Pat. No. 7,881,150, IPR Case No. IPR2014-01011, Decision—Institution of Inter Partes Review 37 C.F.R. § 42.108, issued Dec. 16, 2014.
Inter Partes Review of U.S. Pat. No. 7,881,150, IPR Case No. IPR2014-01011, Exhibit 3001 to Decision—Institution of Inter Partes Review, Excerpts from IEEE Dictionary, issued Dec. 16, 2014.
Inter Partes Review of U.S. Pat. No. 7,881,150, IPR Case No. IPR2014-01011, Exhibit 3002 to Decision—Institution of Inter Partes Review, Excerpts from IEEE Dictionary, issued Dec. 16, 2014.
Inter Partes Review of U.S. Pat. No. 7,881,150, IPR Case No. IPR2014-01011, Exhibit 3003 to Decision—Institution of Inter Partes Review, Excerpts from Oxford English Dictionary, issued Dec. 16, 2014.
Inter Partes Review of U.S. Pat. No. 7,881,150, IPR Case No. IPR2014-01011, Exhibit 3004 to Decision—Institution of Inter Partes Review, Excerpts from Oxford English Dictionary, issued Dec. 16, 2014.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Petitioner1s Reply, filed May 19, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Petitioner Request for Oral Argument, filed Jun. 2, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Patent Owner Motion for Observations, filed Jun. 2, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Patent Owner Request for Oral Argument, filed Jun. 2, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Order Trial Hearing, filed Jun. 8, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Patent Ownerls Submission on Propriety of Petitioner Reply, filed Jun. 9, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Petitioner1s Response to Motion for Observations, filed Jun. 10, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Petitioner1s Response to Netlist1s Submission, filed Jun. 13, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Patent Owner1s Objections to Petitionerls Demonstrative Exhibits, filed Jun. 23, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Patent Owner1s Demonstrative Exhibits, filed Jun. 24, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Petitioner1s Updated Exhibit List, filed Jun. 24, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Supplemental Declaration of Dr. Jagannathan, filed May 19, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, May 10, 2016 Deposition of Carl Sechen, filed May 19, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Petitioner1s Demonstrative Exhibits, filed Jun. 24, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Exhibit 2007, "Deposition of Dr. Srinivasan Jagannathan on May 25, 2016," filed Jun. 2, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Petitioner1s Reply, filed May 19, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Petitioner Request for Oral Argument, filed Jun. 2, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Patent Owner Motion for Observations, filed Jun. 2, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Patent Owner Request for Oral Argument, filed Jun. 2, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Patent Owner1s Submission on Propriety of Petitioner Reply, filed Jun. 9, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Petitioner1s Response to Motion for Observations, filed Jun. 10, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Petitioner1s Response to Patent Owner1s Submission, filed Jun. 13, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Patent Owner1s Objections to Petitioner1s Demonstrative Exhibits, filed Jun. 23, 2016.

(56) **References Cited**

OTHER PUBLICATIONS

Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Patent Owner1s Demonstrative Exhibits, filed Jun. 24, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Petitioner1s Updated Exhibit List, filed Jun. 24, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Supplemental Declaration of Dr. Jagannathan, filed May 19, 2016.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibit 1003, "Declaration of Harold S. Stone," filed Dec. 30, 2016.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Petition for Inter Partes Review, filed Jan. 5, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibit 1003, "Declaration of Harold S. Stone," filed Jan. 5, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibit 1010, Declaration of John J. Kelly Regarding Records of Joint Electron Device Engineering Council (JEDEC), filed Jan. 5, 2017.
Horowitz, "The Art of Electronics," Cambridge Univ. Press, 1989, selected pages.
Huang et al, "An Efficient Parallel Transparent BIST Method for Multiple Embedded Memory Buffers," VLSI Design 2011, p. 379.
Jacob, Bruce L.; "Synchronous DRAM Architectures, Organizations, and Alternative Technologies". University of Maryland, Dec. 10, 2002.
JEDEC Standard No. 21-C Section 4.5.7, 168 Pin Registered SDRAM DIMM Family, Release 7.
JEDEC 21-C, Section 4..6.1, 278 Pin Buffered SDRAM DIMM Family.
JEDEC Standard No. 21-C Section 4.1.2.5, Appendix E, "Specific PD's for Synchronous DRAM (SDRAM)," pp. 1-25.
JEDEC Standard, "Fully Buffered DIMM (FBDIMM): DFx Design for Validation and Test," JESD82-28, Feb. 2008.
JEDEC Standard Double Data Rate (DDR) SDRAM Specification, JESD79 (Jun. 2000).
McCluskey, Edward J., *Logic Design PrincJ,p/es with Emphasis on Testable Semicustom rCircuits*, Prentice Hall, 1986, pp. 104-107 and 119-120.
G. Moore, "Cramming more components onto integrated circuits," Electronics, vol. 38, No. 8, Apr. 19, 1965.
Non-Final Office Action, U.S. Appl. No. 13/412,243, dated Jan. 2, 2014, 20 pages.
Non-final office action, U.S. Appl. No. 13/288,850, dated Oct. 11, 2013, 24 pages.
Non-final office action, U.S. Appl. No. 13/411,344, dated Dec. 31, 2013, 28 pages.
Non-final office action, U.S. Appl. No. 13/473,413, dated Nov. 17, 2011, 46 pages.
Response to non-final office action dated Oct. 14, 2013 for U.S. Appl. No. 13/288,850, filed Jan. 13, 2014, 15 pages.
Response to non-final office action dated Dec. 31, 2013 for U.S. Appl. No. 13/411,344, filed Mar. 31, 2014, 12 pages.
Patent Owner's Response to Office Action mdated ailed Nov. 13, 2012 for Reexamination Control Nos. 95/000,578; 95/000,579, and 95/001,339, filed Jan. 14, 2013, 96 pages.
Patent Owner's Response to Office Action dated Dec. 19, 2012 for Reexamination Control No. 95/001,758, filed Mar. 19, 2013, 61 pages.
Patent Owner's Response to Office Action dated Sep. 26, 2013 for Reexamination Control No. 95/001,758, filed Nov. 26, 2013, 85 pages.
Third Party Requester's Comments after Non-Final Action dated Sep. 26, 2013 for Reexamination Control No. 95/001,758, filed Dec. 26, 2013.
Patent Owner's Appeal Brief for Reexamination Control Nos. 95/000,546 and 95/000,577, filed Oct. 2, 2013, 46 pages.
Patent Trial and Appeal Board Decision on Appeal for Reexamination Control No. 95/001/337, mailed Jan. 16, 2014, 30 pages.
Patent Trial and Appeal Board Decision on Appeal for Reexamination Control No. 95/001/381, mailed Jan. 16, 2014, 24 pages.
Action Closing Prosecution mailed Mar. 27, 2014 for Reexamination Control No. 95/001,758, filed Sep. 14, 2011, 40 pages.
Action Closing Prosecution mailed Mar. 27, 2014 for Reexamination Control No. 95/001,339, filed Jun. 8, 2010, 106 pages.
Petition for Inter Partes Review filed on Jun. 24, 2014 for U.S. Pat. No. 8,516,185, IPR Case No. IPR2014-01029, and all associated documents including cited references and expert declarations, available at https://i;itabtr1als.uspto.gov.
Petition for Inter Partes Review filed on Jun. 22, 2014 for U.S. Pat. No. 7,881,150, IPR Case No. IPR2014-00882 and IPR Case No. IPR2014-01011, and all associated documents including cited references and expert declarations, available at https://ptabtriais.usto.gov.
Petition for Inter Partes Review filed on Jun. 24, 2014 for U.S. Pat. No. 8,081,536, IPR Case No. IPR2014-00883, and all associated documents including cited references and expert declarations, available at https://_tabtrials.uspto._gov.
Reese, "Introduction to Logic Synthesis using Verilog HDL," Morgan&Claypool Publisher, 2006, pp. 1-28.
Notice of Allowance, U.S. Appl. No. 12/504,131, dated Feb. 12, 2013, 52 pgs.
Response to Non-Final Office Action dated Jan. 2, 2014, filed Apr. 2, 2004, for U.S. Appl. No. 13/287,042, filed Nov. 1, 2011, 12 pages.
Office Action dated Apr. 2, 2014, for Japanese Patent Application No. 2012-520662 and English translation thereof, 7 pages.
128Mbit Sdram, Samsung datasheet K4S281632C, Rev. 00. Mar. 2000.
Stone, H.S. Microcomputer Interfacing, Reading, MA: Addison Wesley, 1982.
Texas Instruments SN74LS245 octal bus transceivers with 3-state outputs datasheet, 2002.
U.S. District Court Northern District of California, Case No. 4:13-CV-05889-YGR, *Netlist, Inc.* v. *Smart Storage Systems, Inc., and Diablo Technologies, Inc.*, Smart Storage Systems, 1nc.'s Invalidity Contentions, dated Jun. 6, 2014.
U.S. District Court Northern District of California, Case No. 4:13-CV-05889-YGR, *Netlist, Inc.* v. *Smart Storage Systems, Inc., and Diablo Technologies, Inc.*, Exhibits E.1-E.7 to "Smart Storage Systems, 1nc.'s Invalidity Contentions," dated Jun. 6, 2014.
U.S. District Court Northern District of California, Case No. 4:13-CV-05889-YGR, *Netlist, Inc.* v. *Smart Storage Systems, Inc., and Diablo Technologies, Inc.*, Exhibits F.1-F.5 to "Smart Storage Systems, 1nc.'s Invalidity Contentions," dated Jun. 6, 2014.
U.S. District Court Northern District of California, Case No. 4:13-CV-05889-YGR, *Netlist, Inc.* v. *Smart Storage Systems, Inc., and Diablo Technologies, Inc.*, Exhibits G.1-G.6 to "Smart Storage Systems, 1nc.'s Invalidity Contentions," dated Jun. 6, 2014.
U.S. District Court Northern District of California, Case No. 4:13-CV-05889-YGR, *Netlist, Inc.* v. *Smart Storage Systems, Inc., and Diablo Technologies, Inc.*, Exhibit H to "Smart Storage Systems, 1nc.'s Invalidity Contentions," dated Jun. 6, 2014.
U.S. District Court Northern District of California, Case No. 4:13-CV-05889-YGR, *Netlist, Inc.* v. *Smart Storage Systems, Inc., and Diablo Technologies, Inc.*, Diablo Technologies, 1nc.'s Invalidity Contentions, dated Jun. 6, 2014.
U.S. District Court Northern District of California, Case No. 4:13-CV-05889-YGR, *Netlist, Inc.* v. *Smart Storage Systems, Inc., and Diablo Technologies, Inc.*, Exhibits D-1 to D6 to "Diablo Technologies, 1nc.'s Invalidity Contentions," dated Jun. 6, 2014.
U.S. District Court Northern District of California, Case No. 4:13-CV-05889-YGR, *Netlist, Inc.* v. *Smart Storage Systems, Inc., and Diablo Technologies, Inc.*, Exhibits F-1 to F-5 to "Diablo Technologies, 1nc.'s Invalidity Contentions," dated Jun. 6, 2014.
U.S. District Court Northern District of California, Case No. 4:13-CV-05889-YGR, *Netlist, Inc.* v. *Smart Storage Systems, Inc., and Diablo Technologies, Inc.*, Exhibits G-1 to G-6 to "Diablo Technologies, 1nc.'s Invalidity Contentions," dated Jun. 6, 2014.
U.S. District Court Northern District of California, Case No. 4:13-CV-05889-YGR, *Netlist, Inc.* v. *Smart Storage Systems, Inc., and Diablo Technologies, Inc.*, Exhibit H to "Diablo Technologies, 1nc.'s Invalidity Contentions," dated Jun. 6, 2014.

(56) **References Cited**

OTHER PUBLICATIONS

Anonymous. (May 2002). "Hynix HYMD512G726(L)8-K/H/L Registered DOR SDRAM DIMM, 128M×72 bits Product Description," Rev. 0.1/May, 02, pp. 1-16.

"Elipida Memory to Speak at Intel's Memory Implementers Forum Roundtable Event", Intel Developer Forum, [Online]. Retrieved from the Internet: <URSL: http://www.elpida.com/en/news/2004/02-18.html>, (Jun. 14, 2011), 1 pg.

Google, Inc. v. *Netlist, Inc.*, No. 4:08-cv-04144-SBA, Netlist 1nc.'s Answer to Complaint and Counterclaim (N.D. Ca. Filed Nov. 18, 2008).

Google, Inc. v. *Netlist, Inc.*, No. C 08-04144 SBA Google 1nc.'s Invalidity Contentions Pursuant to PAT. L.F. 3-3, dated Apr. 13, 2009.

Google, Inc. v. *Netlist, Inc.*, No. C08 04144, Complaint for Declaratory Relief, (N.D. Ca Dated Aug. 29, 2008).

1nte(6400/6402 Advanced Memory Buffer Datasheet, published Oct. 2006.

Letter from G. Hopkins Guy III, Orrick, Herrington & Sutcliffe LLP, to R. Scott Oliver, Morrison & Foerster, (Apr. 14, 2009).

Luthra et al. "Interface Synthesis Using Memory Mapping for an FPGA Platform," Proceedings of the 21st International Conference on Computer Design, 2003.

MetaRAM, Inc. v. *Netlist, Inc.* No. 3:09-cv-01309-VRW, MetaRam's Reply to Netlist's Counterclaims, (N.D. Ca. Filed Jun. 3, 2009).

MetaRam, Inc. v. *Netlist, Inc.*, No. 3:09-cv-01309-VRW, Netlist's Answer to Complaint and Counterclaims, (N.D. Ca, filed May 11, 2009).

MetaRAM, Inc. v. *Netlist, Inc.*, No. C09 01309, Complaint for Patent Infringement, (N.D. Ca. Filed Mar. 25, 2009).

Micron "DDR2 SDRAM Registered DIMM (RDIMM)," 2003 Micron Technology, Inc. 18 pages.

Micron "Synchronous DRAM Module MT18LSDT472," 1998, Micron Technology, Inc., 17 pages.

Micron Technical Note,"Decoupling Capacitor Calculations for a DOR Memory Channel," 2004, 3 pages.

Miles J. Murdocca et al., "Principles of Computer Architecture", Prentice Hall, Inc., Upper Saddle River, NJ, 2000, pp. 243-252.

Murdocca et al., "Principles of Computer Architecture," Prentice Hall, 2000, pp. 249-251.

Netlist, Inc. v. *MetaRam, Inc.*, No. 09-165-GMS, MetaRAM, 1nc.'s Answer and Affirmative Defenses to Plaintiff's Complaint, dated Apr. 20, 2009.

Netlist, Inc. v. *MetaRAM, Inc.*, No. 1:09-ccv-00165-GMS, Complaint for Patent Infringement, (D. Del. Filed Mar. 12, 2009).

Non-Final Action Closing Prosecution mailed Sep. 1, 2010, for Control No. 95/001,339, filed Apr. 10, 2010, 17 pages.

Non-Final Action Closing Prosecution mailed Jun. 21, 2011, for Control No. 95/001,381, filed Jun. 9, 2010, 34 pages.

Non-Final Action Closing Prosecution mailed Mar. 12, 2012, for Control No. 95/001,337, filed Apr. 19, 2010, 33 pages.

Non-Final Action mailed Aug. 27, 2010, for Control No. 95/000,546, filed May 11, 2010, 16 pages.

Non-Final Action mailed Sep. 8, 2010, for Control No. 95/001,381, filed Jun. 9, 2010, 17 pages.

Non-Final Action mailed Apr. 4, 2011, for Control No. 95/001,339, filed Apr. 20, 2010, 61 pages. (merged with 95/000,578 and 95/000,579).

Non-Final Action mailed Jun. 15, 2011, for Control No. 95/001,381, filed Jun. 9, 2010, 33 pages.

Non-Final Action mailed Sep. 27, 2011, for Control No. 95/001,337, filed Apr. 19, 2010, 19 pages.

Non-Final Action mailed Oct. 4, 2011, for Control No. 95/001,339, filed Apr. 20, 2010, 77 pages. (merged with 95/000,578 and 95/000,579).

Non-Final Action mailed Oct. 14, 2011, for Control No. 95/001,339, filed Apr. 30, 2010, 99 pages. (merged with 95/000,578 and 95/000,579).

Non-Final Office Action dated Nov. 16, 2011, for U.S. Appl. No. 95/001,758 filed Sep. 14, 2011, 25 pages.

Order Granting Request for Inter Partes Reexamination mailed Nov. 16, 2011, for U.S. Appl. No. 95/001,758, filed Sep. 14, 2011, 13 pages.

Order Granting Request for Inter Partes Reexamination mailed Aug. 9, 2010, for Control No. 95/000,546, filed May 11, 2010, 22 pages.

Order Granting Request for Inter Partes Reexamination mailed Aug. 27, 2010, for Control No. 95/001,337, filed Apr. 19, 2010, 21 pages.

Order Granting Request for Inter Partes Reexamination mailed Sep. 1, 2010, for Control No. 95/001,339, filed Apr. 20, 2010, 14 pages.

Order Granting Request for Inter Partes Reexamination mailed Sep. 8, 2010, for Control No. 95/000,381, filed Jun. 9, 2010, 21 pages.

Order Granting Request for Inter Partes Reexamination mailed Jan. 14, 2011, for Control No. 95/000,579, filed Oct. 21, 2010, 12 pages.

Order Granting Request for Inter Partes Reexamination mailed Jan. 18, 2011, for Control No. 95/000,577, filed Oct. 20, 2010, 17 pages.

Right of appeal Notice mailed Feb. 7, 2012, for Control No. 95/001,381, 33 pages.

Right of Appeal Notice mailed Jun. 22, 2012, for Control No. 95/001,337, filed Jun. 4, 2010, 34 pages.

PC133 SDRAM Registered DIMM Design Specification, Revision 1.1, Aug. 1999, 62 pages.

Texas Instruments, "TM2SR72EPN 2097152 by 72-Bit, TM4SRT2EPN 4194304 by 72-Bit, Synchronous Dynamic RAM Modules," 1997, 15 pages.

U.S. District Court Central District of California, Case No. CV09 06900, *Netlist, Inc.* vs. *Inphi Corporation*, Complaint for Patent Infringement, filed Sep. 22, 2009 in 10 pages.

U.S. District Court Central District of California, Case No. CV09 06900, *Netlist, Inc.* vs. *INPHI Corporation*, Defendant 1nphi Corporation's Answer to Plaintiff's Complaint for Patent Infringement, filed Nov. 12, 2009 in 6 pages.

U.S. District Court Central District of California, Case No. CV09 06900, *Netlist, Inc.* vs. *INPHI Corporation*, Defendant 1nphi Corporation's Answer to Plaintiff's First Amended Complaint for Patent Infringement, filed Feb. 11, 2010 in 9 pages.

U.S. District Court Central District of California, Case No. CV09 06900, *Netlist, Inc.* vs. *INPHI Corporation*, Defendant 1nphi Corporation's Notice of Motion and Motion for Stay Pending Reexaminations and Interference Proceeding Regarding the Patents-In-Suit; Memorandum of Points and Authorities in Support Thereof, filed Apr. 21, 2010 in 28 pages.

U.S. District Court Central District of California, Case No. CV09 06900, *Netlist, Inc.* vs. *INPHI Corporation*, Plaintiff Netlist Inc's Opposition to Defendant Inphi Corporation's Motion for Stay Pending Reexaminations and Interference Proceedings Regarding the Patents-In-Suit, filed May 3, 2010 in 23 pages.

U.S. District Court Central District of California, Case No. CV09 06900, *Netlist, Inc.* vs. *Inphi Corporation*, Plaintiff Netlist, Inc's First Amended Complaint for Patent Infringement, filed Dec. 23, 2009 in 8 pages.

U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc.* v. *Netlist, Inc.*, Defendant Netlist, 1nc.'s Claim Construction Reply Brief, filed Sep. 22, 2009 in 19 pages.

U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc.* v. *Netlist, Inc.*, Defendant Netlist, 1nc.'s Opening Claim Construction Brief, filed Jul. 28, 2009 in 21 pages.

U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc.* v. *Netlist, Inc.*, Defendant Netlist, 1nc.'s Opposition to Google Inc's Motion for Summary Judgment of Invalidity, filed Jul. 6, 2010 in 13 pages.

U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc.* v. *Netlist, Inc.*, Exhibit A to Joint Claim Construction and Prehearing Statement, filed Jun. 12, 2009 in 2 pages.

U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc.* v. *Netlist, Inc.*, Exhibit B to Joint Claim Construction and Prehearing Statement, filed Jun. 12, 2009 in 36 pages.

U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc.* v. *Netlist, Inc.*, Joint Claim Construction and Prehearing Statement, filed Jun. 12, 2009 in 5 pages.

U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc.* v. *Netlist, Inc.*, Netlist, 1nc.'s Answer to Complaint and Counterclaims, filed Nov. 18, 2008 in 9 pages.

(56) **References Cited**

OTHER PUBLICATIONS

U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc.* v. *Netlist, Inc.*, Order Re Claim Construction, filed Nov. 16, 2009 in 1 page.
U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc.* v. *Netlist, Inc.*, Plaintiff Google's Reply to Counterclaims, filed Dec. 8, 2008 in 4 pages.
U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc.* v. *Netlist, Inc.*, Stipulation Re: Additional Agreed-Upon Claim Constructions, filed Oct. 28, 2009 in 3 pages.
U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc.* vs. *Netlist, Inc.*, [Redacted] Google 1nc.'s Responsive Claim Construction Brief, filed Aug. 25, 2009 in 30 pages.
U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc.* vs. *Netlist, Inc.*, Amended Exhibit A to Joint Claim Construction and Prehearing Statement, filed Oct. 28, 2009 in 1 page.
U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc.* vs. *Netlist, Inc.*, Appendix 1 to Google's Responsive Claim Construction Brief, filed Nov. 12, 2009 in 4 pages.
U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc.* vs. *Netlist, Inc.*, Attachment 1 to Exhibit B to Joint Claim Construction and Prehearing Statement, filed Jun. 12, 2009 in 7 pages.
U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc.* vs. *Netlist, Inc.*, Attachment 2 to Exhibit B to Joint Claim Construction and Prehearing Statement, filed Jun. 12, 2009 in 12 pages.
U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc.* vs. *Netlist, Inc.*, Complaint for Declaratory Relief, filed Aug. 29, 2008 in 49 pages.
U.S. District Court Northern District of California, Case No. CV09 05718, *Netlist, Inc.* vs. *Google, Inc.*, Complaint for Patent Infringement, filed Dec. 4, 2009 in 47 pages.
U.S. District Court Northern District of California, Case No. CV09 05718, *Netlist, Inc.* vs. *Google, Inc.*, Defendant Google Inc's Responsive Claim Construction Brief, filed Aug. 4, 2010 in 108 pages.
U.S. District Court Northern District of California, Case No. CV09 05718, *Netlist, Inc.* vs. *Google, Inc.*, Exhibit A to Joint Claim Construction and Prehearing Statement under Patent L. R. 4-3, filed Jun. 25, 2010 in 2 pages.
U.S. District Court Northern District of California, Case No. CV09 05718, *Netlist, Inc.* vs. *Google, Inc.*, Exhibit B to Joint Claim Construction and Prehearing Statement under Patent L. R. 4-3, filed Jun. 25, 2010 in 23 pages.
U.S. District Court Northern District of California, Case No. CV09 05718, *Netlist, Inc.* vs. *Google, Inc.*, Google's Answer to Plaintiff's Complaint for Patent Infringement; and Assertion of Counterclaims, filed Feb. 12, 2010 in 13 pages.
U.S. District Court Northern District of California, Case No. CV09 05718, *Netlist, Inc.* vs. *Google, Inc.*, Joint Claim Construction and Prehearing Statement Under Patent Local Rule 4-3, filed Jun. 25, 2010 in 5 pages.
U.S. District Court Northern District of California, Case No. CV09 05718, *Netlist, Inc.* vs. *Google, Inc.*, Plaintiff Netlist, 1nc.'s Reply Claim Construction Brief, filed Aug. 16, 2010 in 17 pages.
U.S. District Court Northern District of California, Case No. CV09 05718, *Netlist, Inc.* vs. *Google, Inc.*, Plaintiff Netlist, 1nc.'s Reply to Defendant Google 1nc.'s Counterclaim, filed Mar. 8, 2010 in 11 pages.
U.S. District Court Northern District of California, Case No. CV09 05718, *Netlist, Inc.* vs. *Google, Inc.*, Plaintiff Netlist, Inc's Opening Claim Construction Brief, filed Jul. 16, 2010 in 162 pages.
US District Court Civil Docket; *Google Inc.* v. *Netlist Inc.*; 4:08cv04144; filed Aug. 29, 2008.
US District Court Civil Docket; *Netlist Inc.* v. *Google Inc.*; 4:09cv5718, filed Dec. 4, 2009 in 10 pages.
US District Court. Civil Docket; *Netlist Inc.* v. *Inphi Corporation*; 2:09cv6900; Date filed Sep. 22, 2009.
Vogt, Pete, "Fully Buffered DIMM (FB-DIMM) Server Memory Architecture: Capacity, Performance, Reliability, and Longevity," Intel, Feb. 18, 2004, 33 pages.
Reexam U.S. Appl. No. 95/000,546 for U.S. Pat. No. 7,289,386, filed May 11, 2010, Netlist, Inc.
Reexam U.S. Appl. No. 95/000,577 for U.S. Pat. No. 7,289,386, filed Oct. 20, 2010, Netlist, Inc.
Reexam U.S. Appl. No. 95/000,578 for U.S. Pat. No. 7,619,912, filed Oct. 20, 2010, Netlist, Inc.
Reexam U.S. Appl. No. 95/000,579 for U.S. Pat. No. 7,619,912, filed Oct. 21, 2010, Netlist, Inc.
Reexam U.S. Appl. No. 95/001,337 for U.S. Pat. No. 7,636,274, filed Jun. 4, 2010, Netlist, Inc.
Reexam U.S. Appl. No. 95/001,381 for U.S. Pat. No. 7,532,537, filed Jun. 9, 2010, Netlist, Inc.
Reexam U.S. Appl. No. 95/001,338; for U.S. Pat. No. 7,532,537, filed Apr. 19, 2010, Netlist, Inc.
Reexam U.S. Appl. No. 95/001,758 for U.S. Pat. No. 7,864,627, filed Sep. 15, 2010, Netlist, Inc.
Reexam U.S. Appl. No. 95/001,339 for U.S. Pat. No. 7,619,912, filed Jun. 8, 2010, Netlist, Inc.
Reexam U.S. Appl. No. 95/002,399 for U.S. Pat. No. 8,250,295, filed Sep. 15, 2010, Netlist, Inc.
U.S. Appl. No. 11/075,395, filed Mar. 7, 2005, Netlist, Inc.
U.S. Appl. No. 11/173,175, filed Jul. 1, 2005, Netlist, Inc.
U.S. Appl. No. 11/862,931, filed Sep. 27, 2007, Netlist, Inc.
U.S. Appl. No. 12/577,682, filed Oct. 12, 2009, Netlist, Inc.
U.S. Appl. No. 12/954,492, filed Nov. 24, 2010, Netlist, Inc.
U.S. Appl. No. 12/912,623, filed Oct. 26, 2010, Netlist, Inc.
U.S. Appl. No. 11/335,875, filed Jan. 19, 2006, Netlist, Inc.
U.S. Appl. No. 12/408,652, filed Mar. 20, 2009, Netlist, Inc.
U.S. Appl. No. 12/629,827, filed Dec. 2, 2009, Netlist, Inc.
U.S. Appl. No. 12/955,711, filed Nov. 29, 2010, Netlist, Inc.
U.S. Appl. No. 12/981,380, filed Dec. 29, 2010, Netlist, Inc.
U.S. Appl. No. 13/154,172, filed Jun. 6, 2011, Netlist, Inc.
U.S. Appl. No. 13/287,042, filed Nov. 1, 2011, Netlist, Inc.
U.S. Appl. No. 13/473,413, filed May 16, 2012, Netlist, Inc.
U.S. Appl. No. 13/032,470, filed Feb. 22, 2011, Netlist, Inc.
U.S. Appl. No. 13/287,081, filed Nov. 1, 2011, Netlist, Inc.
U.S. Appl. No. 13/971,231, filed Aug. 20, 2013, Netlist, Inc.
U.S. Appl. No. 12/422,853, filed Apr. 13, 2009, Netlist, Inc.
U.S. Appl. No. 13/412,243, filed Mar. 5, 2012, Netlist, Inc.
U.S. Appl. No. 12/422,925, filed Apr. 13, 2009, Netlist, Inc.
U.S. Appl. No. 13/183,253, filed Jul. 14, 2011, Netlist, Inc.
U.S. Appl. No. 13/745,790, filed Jan. 19, 2013, Netlist, Inc.
U.S. Appl. No. 14/229,844, filed Mar. 29, 2014, Netlist, Inc.
U.S. Appl. No. 12/504,131, filed Jul. 16, 2009, Netlist, Inc.
U.S. Appl. No. 12/761,179, filed Apr. 15, 2010, Netlist, Inc.
U.S. Appl. No. 13/288,850, filed Nov. 3, 2011, Netlist, Inc.
U.S. Appl. No. 13/411,344, filed Mar. 2, 2012, Netlist, Inc.
U.S. Appl. No. 13/952,599, filed Jul. 27, 2013, Netlist, Inc.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Paper 14, 'Patent Owner's Response,' filed Oct. 19, 2018.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1047, 'Certified translation of the Examination Decision issued by the State Intellectual Property Office of the P.R.C. (dated May 30, 2018),' filed Feb. 11, 2019.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1048, 'Certified translation of the Preliminary Opinion issued by the German Patent and Trademark Office (dated Jan. 8, 2019),' filed Feb. 11, 2019.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Paper 19, 'Petitioners' Request for Oral Argument,' filed Feb. 19, 2019.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Paper 21, 'Patent Owner Objections to Reply Evidence,' filed Feb. 19, 2019.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Paper 20, 'Patent Owner Request for Oral Argument,' filed Feb. 19, 2019.

(56) **References Cited**

OTHER PUBLICATIONS

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Paper 22, 'Patent Owner Motion to Exclude,' filed Feb. 19, 2019.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Paper 23, 'Corrected Petitioners' Reply to Patent Owner's Response,' filed Feb. 20, 2019.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Paper 28, 'Order—Trial Hearing,' filed Mar. 22, 2019.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Paper 29, 'Termination Decision Document,' filed Jun. 27, 2019.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Paper 30, 'Record of Oral Hearing,' filed Jun. 27, 2019.
Netlist, Inc., Examination Decision of Request for Invalidation, CN201080039043.0, May 30, 2018, 43 pgs.
Inter Partes Review of U.S. Pat. No. 8,516,185 B2, Case No. IPR2017-00577, Record of Oral Hearing, filed Jun. 19, 2018.
Inter Partes Review of U.S. Pat. No. 8,516,185 B2, Case No. IPR2017-00577, Final Written Decision, filed Jul. 5, 2018.
Inter Partes Review of U.S. Pat. No. 7,532,537 B2, Case No. IPR2017-00667, Final Written Decision, filed Jul. 18, 2018.
Inter Partes Review of U.S. Pat. No. 7,532,537 B2, Case No. IPR2017-00668, Final Written Decision, filed Jul. 18, 2018.
Inter Partes Review of U.S. Pat. No. 9,606,907 B2, Case No. IPR2018-00362, Order, *Conduct of the Proceedings*, filed Aug. 2, 2018.
Inter Partes Review of U.S. Pat. No. 9,606,907 B2, Case No. IPR2018-00362, and Case No. IPR2018-00363, Decision Granting Institution of Inter Partes Review, filed Aug. 6, 2018.
Inter Partes Review of U.S. Pat. No. 9,606,907 B2, Case No. IPR2018-00362, and Case No. IPR2018-00363, Exhibit 3001, 'email from Mehran Arjomand,' filed Jul. 30, 2018.
Inter Partes Review of U.S. Pat. No. 9,606,907 B2, Case No. IPR2018-00362, and Case No. IPR2018-00363, Exhibit 3001, 'email from Michael D. Hatcher,' filed Jul. 30, 2018.
Inter Partes Review of U.S. Pat. No. 9,606,907 B2, Case No. IPR2018-00364, and Case No. IPR2018-00365, Decision Granting Institution of Inter Partes Review, filed Jun. 29, 2018.
Inter Partes Review of U.S. Pat. No. 9,606,907 B2, Case No. IPR2018-00364, Patent Owner's Objections to Petition Evidence, filed Aug. 20, 2018.
Response, Article 94 3 EPC, EP 18179414.0, dated Aug. 6, 2019, 23 pgs.
Netlist, Inc., Third Party Observation, EP 3404660, Aug. 6, 2019, 34 pgs.
Office Action, EP 18179414, dated Apr. 2, 2019, 6 pgs.
Patent document filed by third party, EP 18179414, Feb. 22, 2019, 10 pgs.
Observations by third parties, EP 18179414.0, Feb. 22, 2019, 1 pg.
Patent Document filed by a third party, Feb. 20, 2019, 10 pgs.
Netlist, Inc., Observations by third parties, EP 3404660, Feb. 20, 2019, 46 pgs.
Reply to Written Opinion prepared by the EPO, EP 18179414.0, dated Jan. 17, 2019, 20 pgs.
Netlist, Inc., Information on Search Strategy, EP 18179414, Oct. 23, 2018, 1 pg.
European Search Report, EP 18179414, dated Oct. 23, 2018, 3 pgs.
European Search Opinion, EP 18179414, dated Oct. 23, 2018, 3 pgs.
Office Action, European Patent Application No. 18179414.0, dated Jul. 30, 2020, 9 pgs.
Netlist, Inc., Communication Pursuant to Article 94 (3) EPC, 18179414.0, dated Mar. 23, 2020, 5 pgs.
Netlist, Inc., United States Court of Appeals for the Federal Circuit, SK Hynix Inc, Corrected Principal Brief of Appeals, Document 26, Case: 19-2340, filed Feb. 3, 2020, 284 pgs.
Netlist, Inc., United States Court of Appeals for the Federal Circuit, SK Hynix Inc., Response and Cross-Appeal Opening Brief for Cross-Appellant Netlist, Inc., Document 28, Case: 19-2340, Mar. 23, 2020, 80 pgs.
Commission Opinion, United States International Trade Commission, Certain Memory Modules and Components Thereof, Investigation No. 337-TA-1089, Apr. 21, 2020, 30 pgs.
Netlist, Inc., Communication Pursuant to Article 94(3) EPC, EP 18179414.0, dated Mar. 23, 2020, 5 pgs.
Notice of Grant, European Patent Application No. 19159900.0, dated Sep. 17, 2020, 61 pgs.
Office Action, European Patent Application No. 18179414.0, dated Sep. 24, 2020, 5 pgs.

* cited by examiner

Figure 1A: (Prior Art)

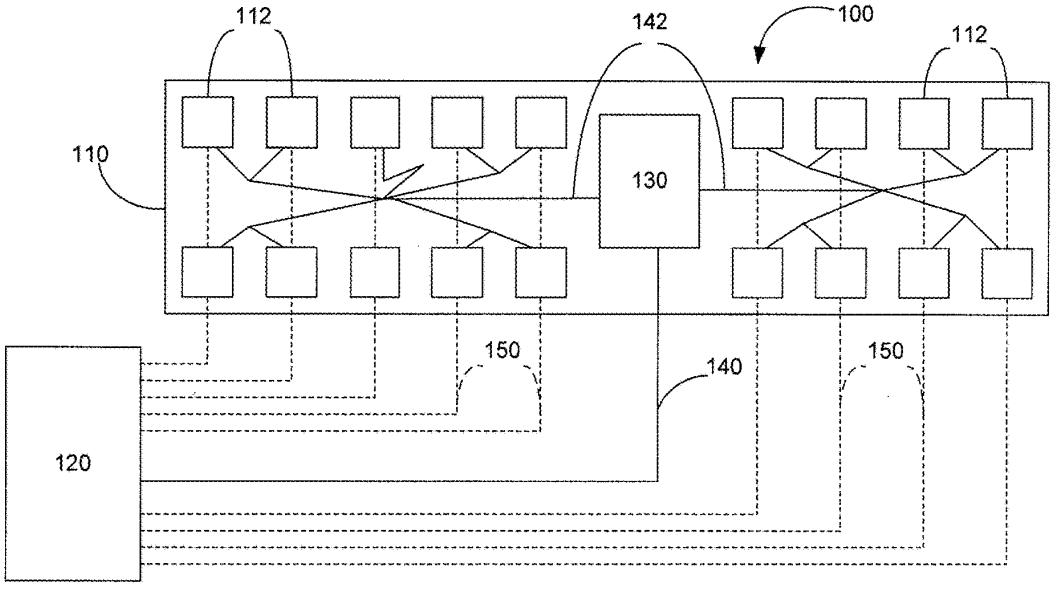


Figure 1B: (Prior Art)

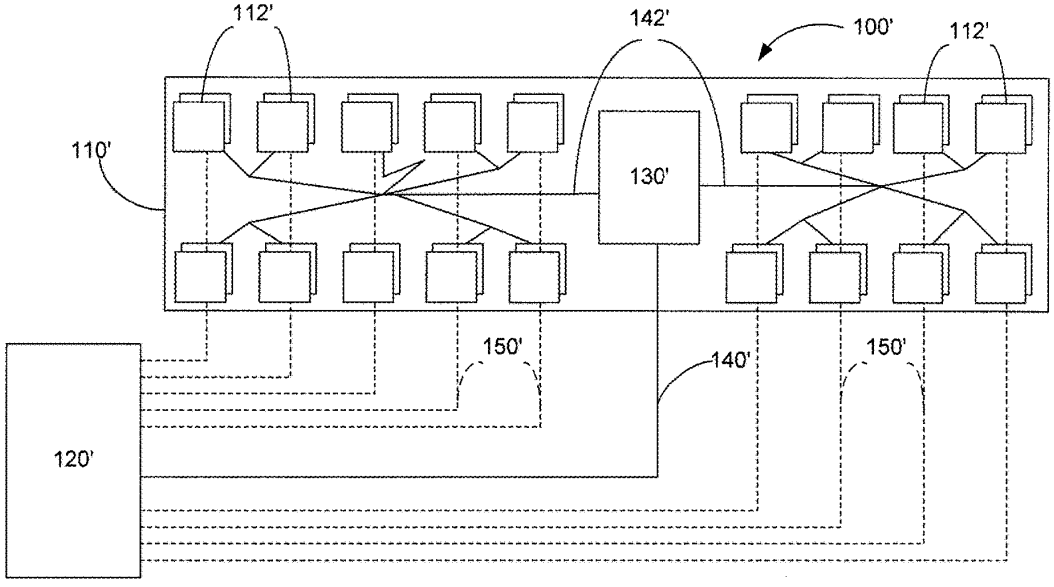


Figure 2A: (Prior Art)

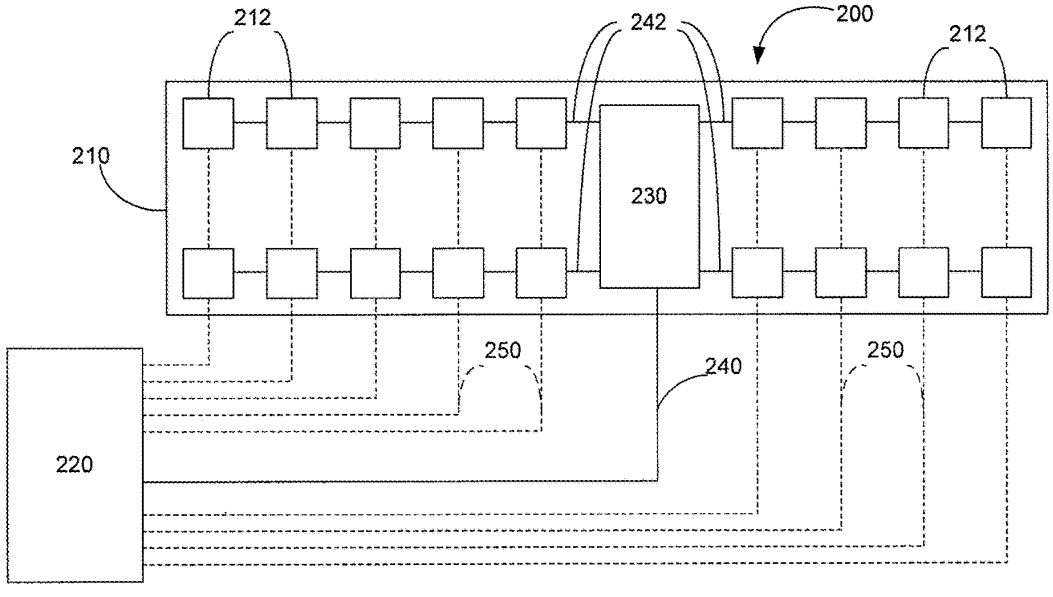


Figure 2B: (Prior Art)

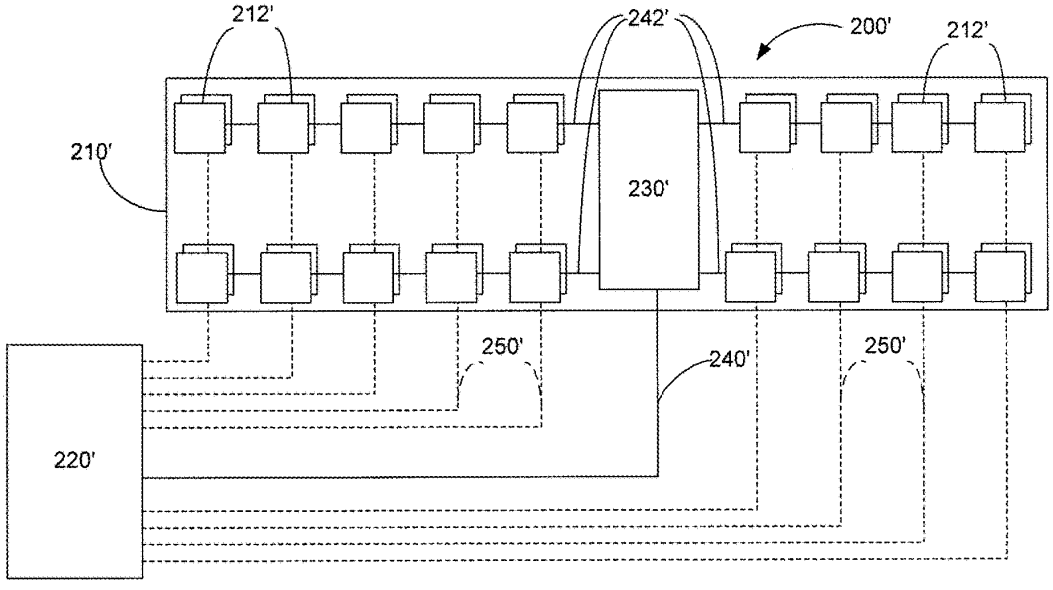


Figure 2C: (Prior Art)

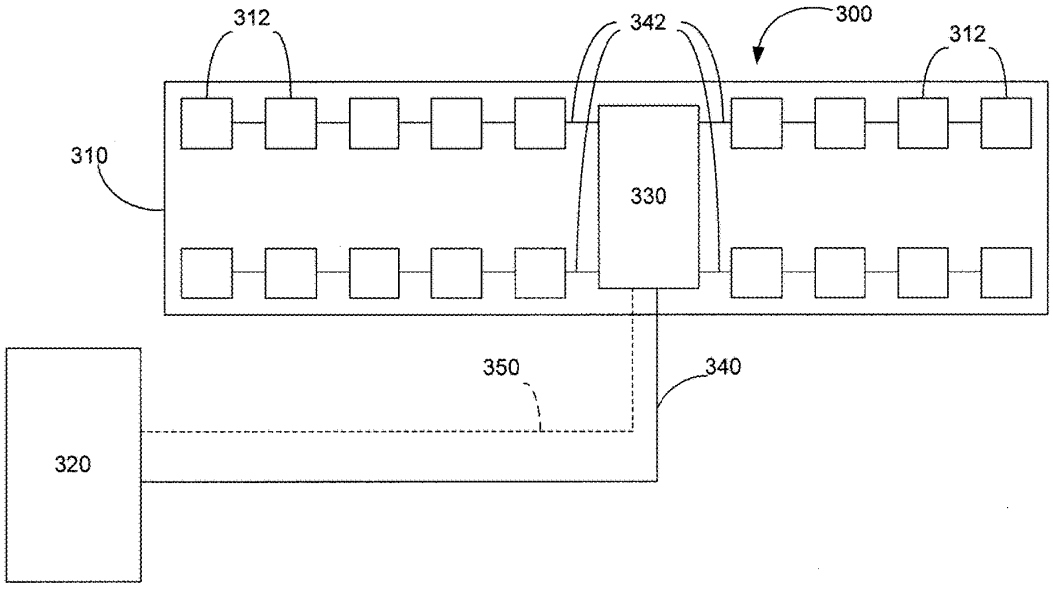
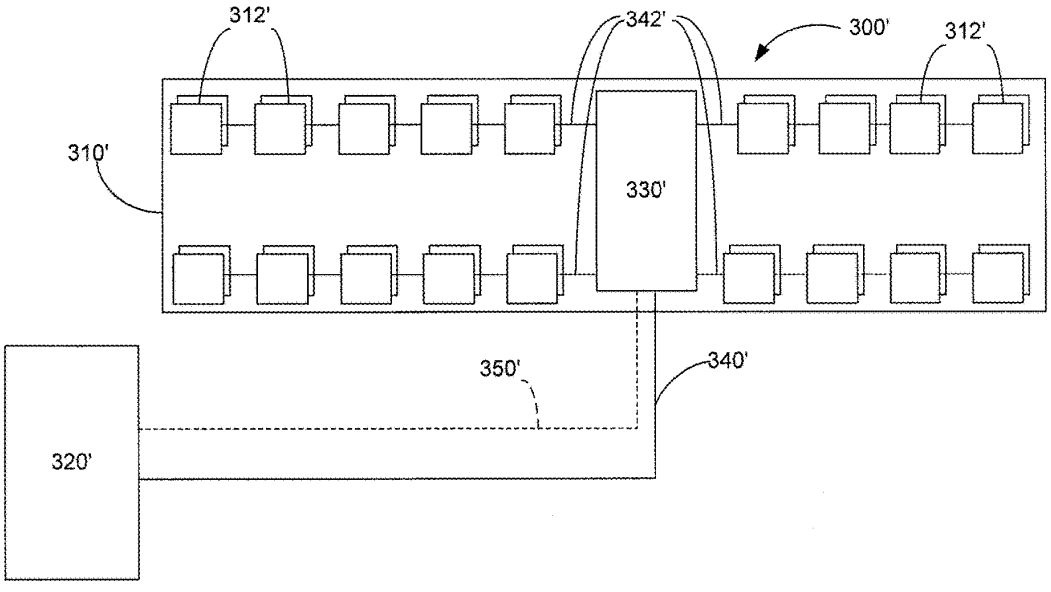


Figure 2D: (Prior Art)



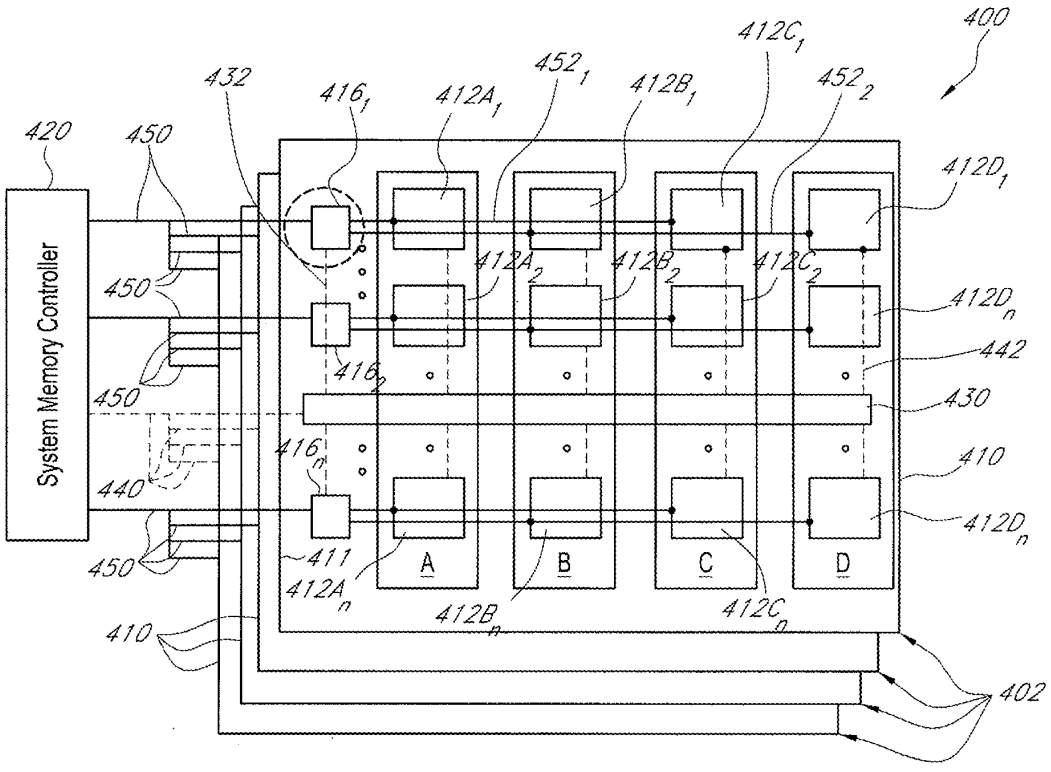

400
420
450
432
416_1
$412A_1$
452_1
$412B_1$
$412C_1$
452_2
$412D_1$
System Memory Controller
$412A_2$
$412B_2$
$412C_2$
$412D_n$
442
416_2
430
416_n
410
440
411
$412A_n$
A
B
C
D
$412B_n$
$412C_n$
402

FIG. 3A

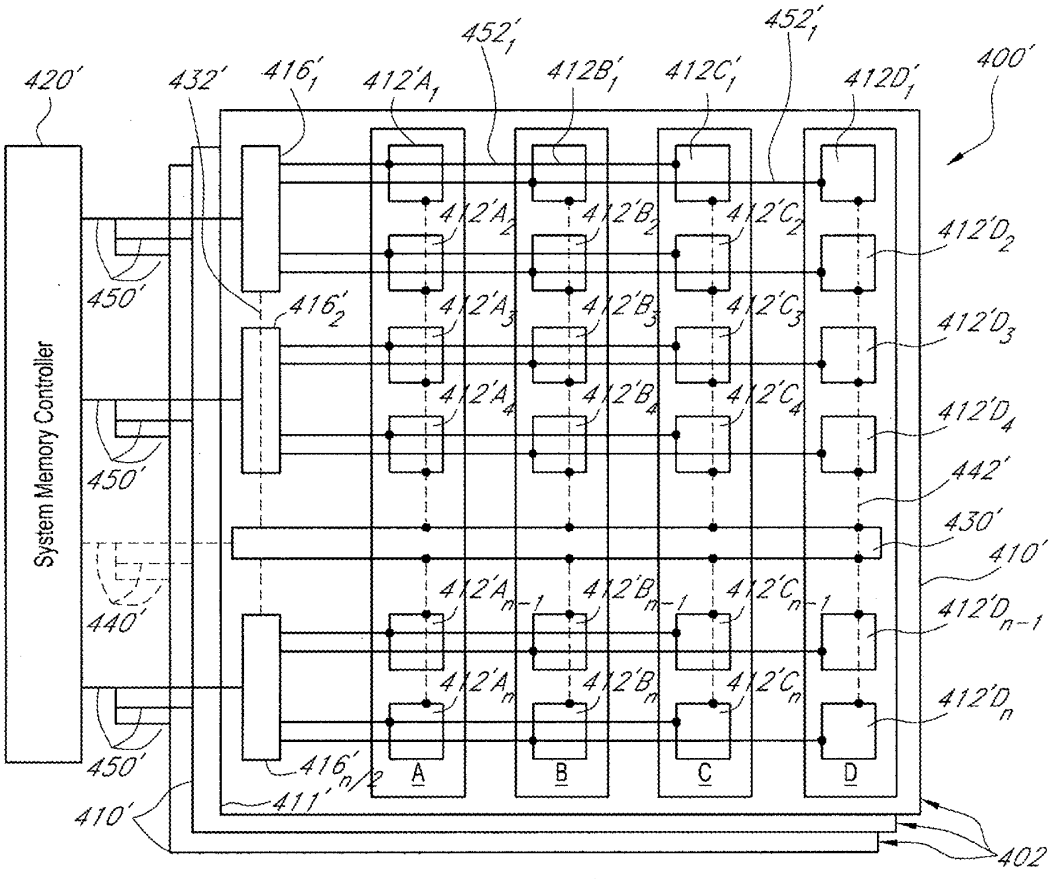

System Memory Controller
420′
432′
416′1
412′A1
452′1
412B′1
412C′1
452′1
412D′1
400′
412′A2
412′B2
412′C2
412′D2
450′
416′2
412′A3
412′B3
412′C3
412′D3
412′A4
412′B4
412′C4
412′D4
442′
430′
410′
440′
412′An−1
412′Bn−1
412′Cn−1
412′Dn−1
412′An
412′Bn
412′Cn
412′Dn
450′
416′n/2
A
B
C
D
411′
410′
402

FIG. 3B

Figure 3C:

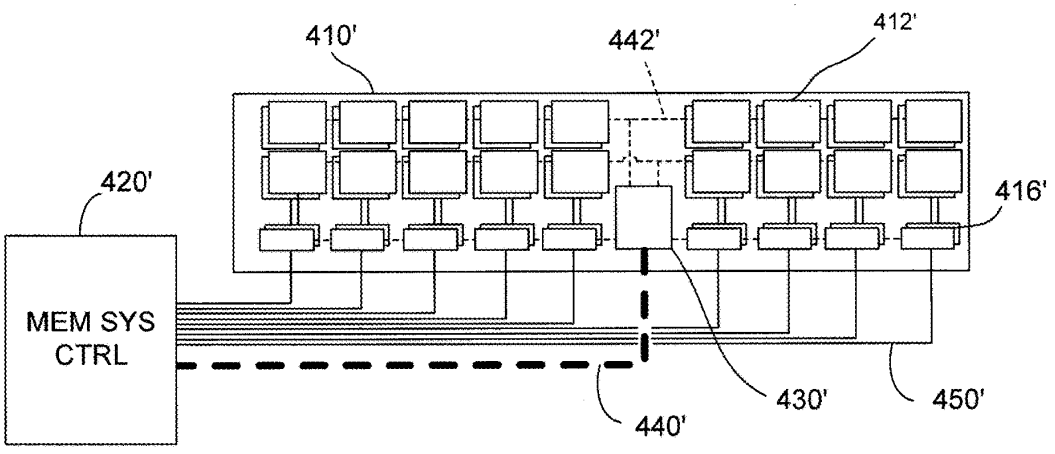
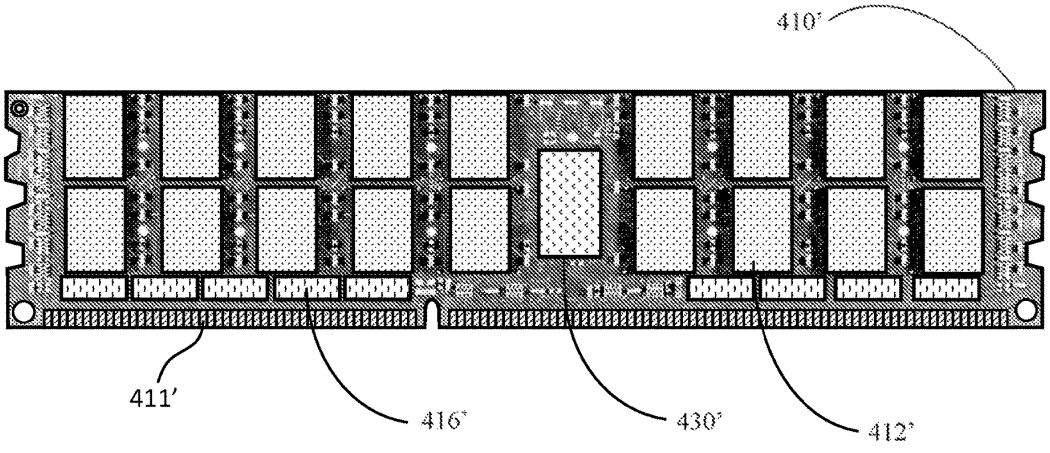


Figure 3D:



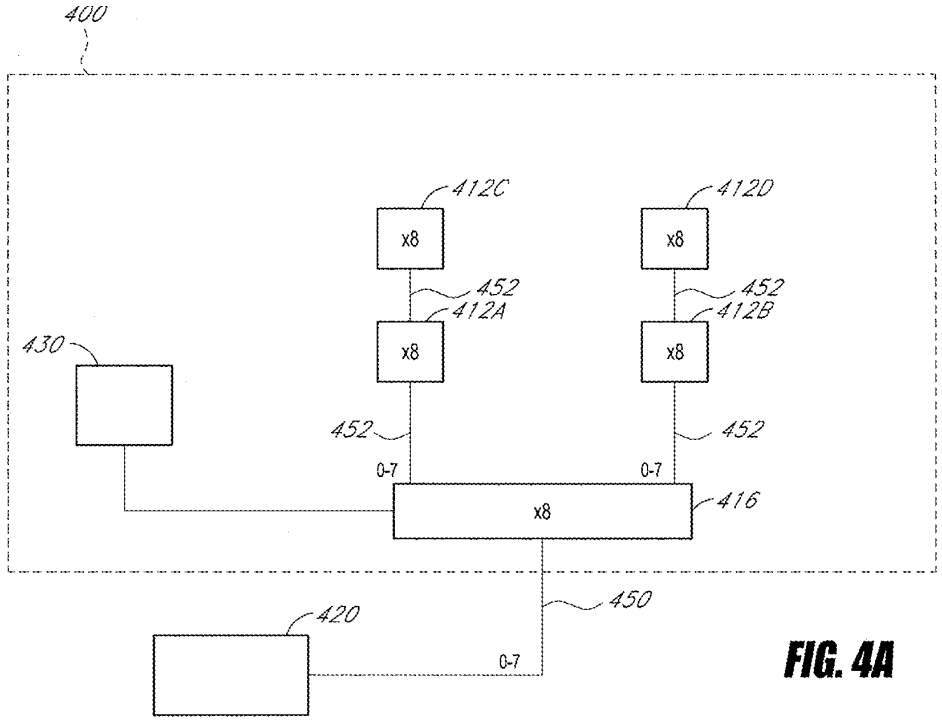

400
412C
412D
x8
x8
452
412A
452
412B
x8
x8
430
452
452
0-7
0-7
x8
416
450
420
0-7

FIG. 4A

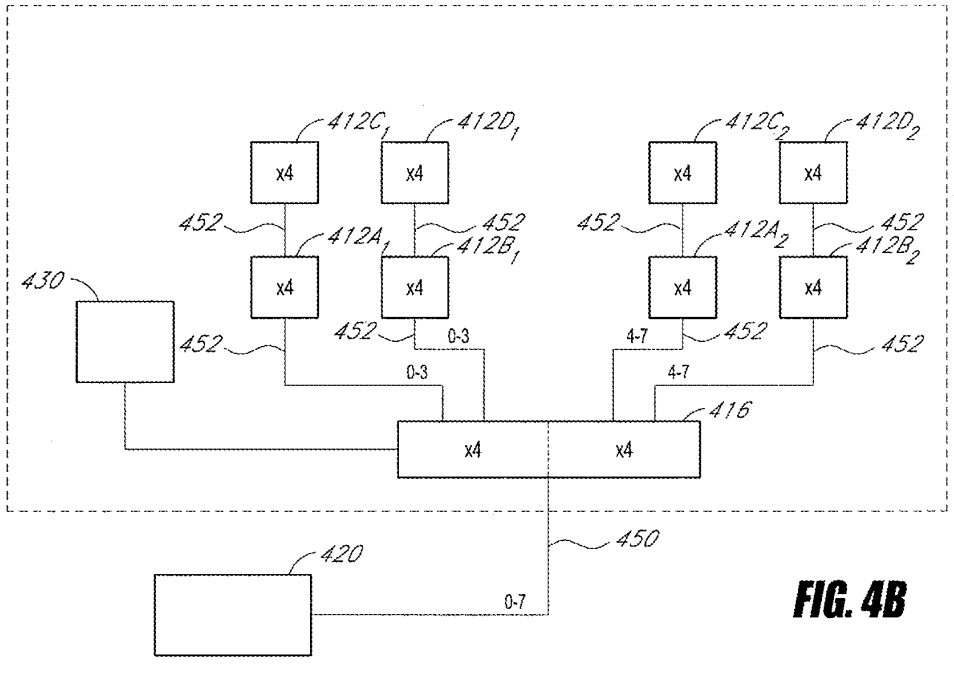

412C1
412D1
412C2
412D2
x4
x4
x4
x4
452
412A1
452
412B1
452
412A2
452
412B2
x4
x4
x4
x4
430
452
452
0-3
4-7
452
452
0-3
4-7
416
x4
x4
450
420
0-7

FIG. 4B

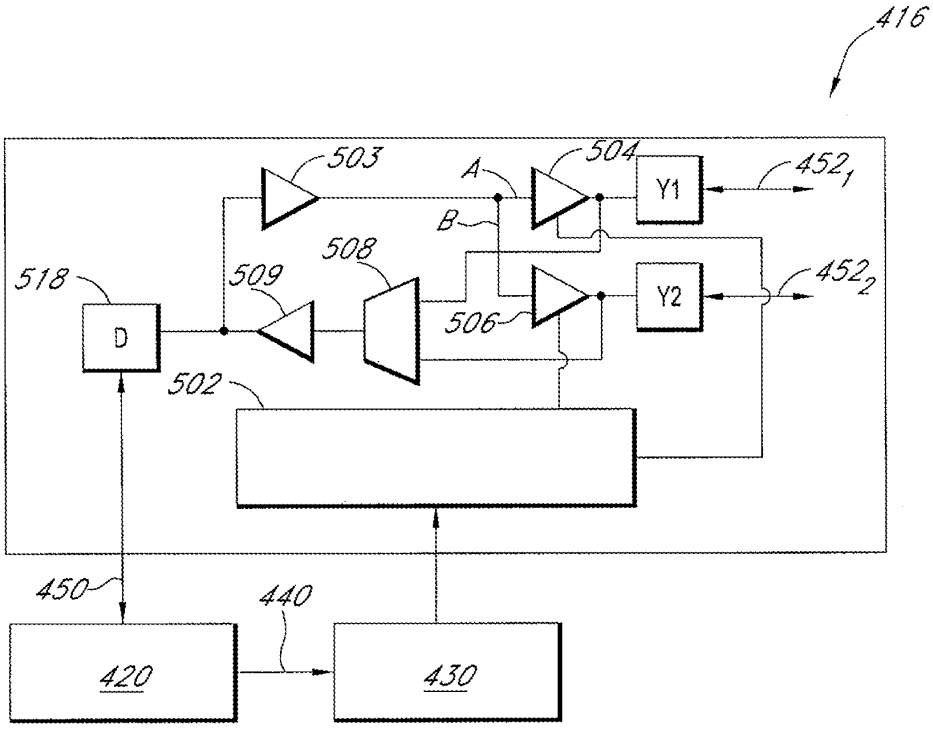

416
503
504
A
B
Y1
452_1
508
509
518
D
506
Y2
452_2
502
450
440
420
430

FIG. 5

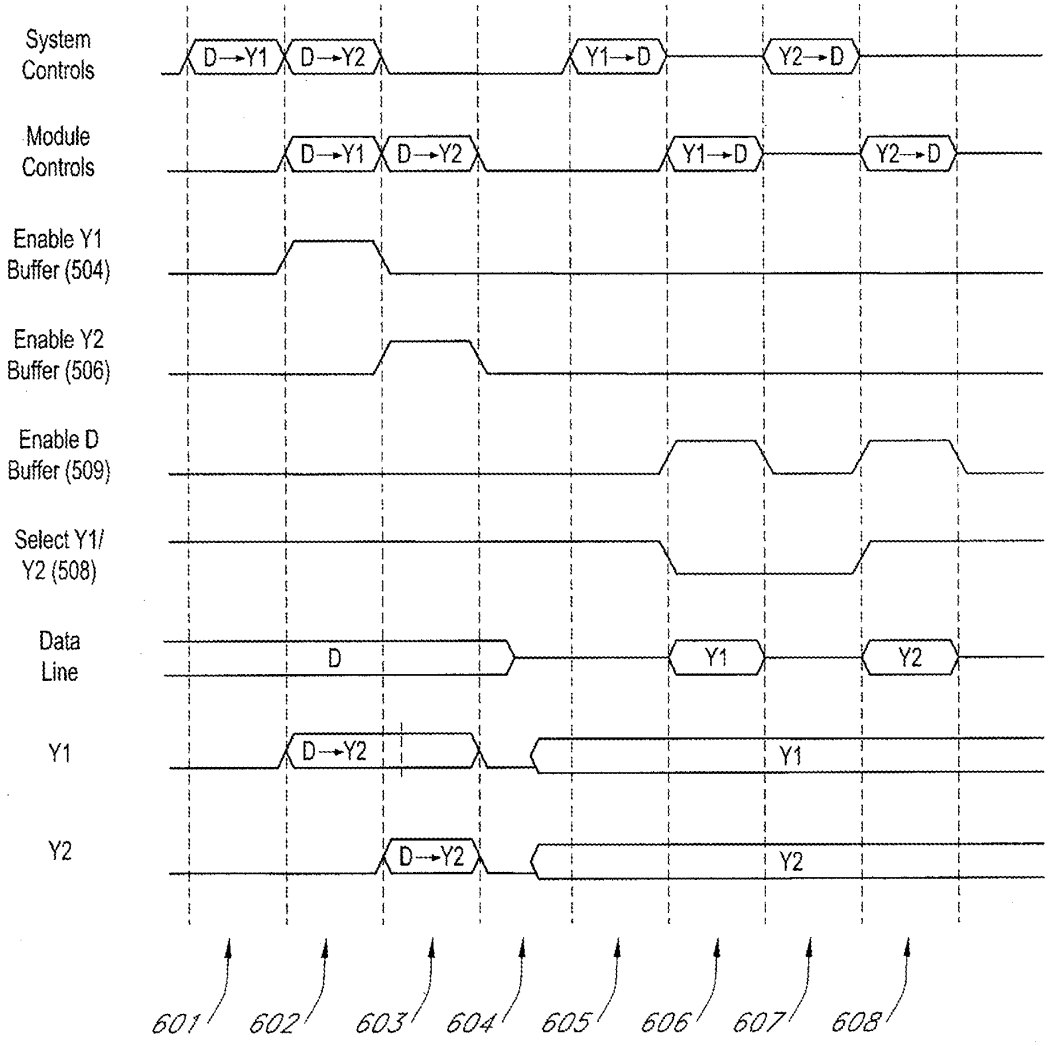

System Controls
Module Controls
Enable Y1 Buffer (504)
Enable Y2 Buffer (506)
Enable D Buffer (509)
Select Y1/ Y2 (508)
Data Line
Y1
Y2
D→Y1
D→Y2
Y1→D
Y2→D
D
Y1
Y2
601
602
603
604
605
606
607
608

FIG. 6

MEMORY MODULE WITH CONTROLLED BYTE-WISE BUFFERS

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a continuation from U.S. patent application Ser. No. 13/970,606, filed Aug. 20, 2013, to be issued as U.S. Pat. No. 9,606,907, which is a continuation from U.S. patent application Ser. No. 12/761,179, filed Apr. 15, 2010, now U.S. Pat. No. 8,516,185, which is a continuation-in-part from U.S. patent application Ser. No. 12/504,131, filed Jul. 16, 2009, now U.S. Pat. No. 8,417,870, each of which is incorporated in its entirety by reference herein.

BACKGROUND

The present disclosure relates generally to memory subsystems of computer systems, and more specifically to systems, devices, and methods for improving the performance and the memory capacity of memory subsystems or memory "boards," particularly memory boards that include dual in-line memory modules (DIMMs).

Certain types of computer memory subsystems include a plurality of dynamic random-access memory (DRAM) or synchronous dynamic random access memory (SDRAM) devices mounted on a printed circuit board (PCB). These memory subsystems or memory "boards" are typically mounted in a memory slot or socket of a computer system, such as a server system or a personal computer, and are accessed by the processor of the computer system. Memory boards typically include one or more memory modules, each with a plurality of memory devices (such as DRAMs or SDRAMs) in a unique configuration of rows, columns, and banks, which provide a total memory capacity for the memory module.

The memory devices of a memory module are generally arranged as ranks or rows of memory, each rank of memory generally having a bit width. For example, a memory module in which each rank of the memory module is 64 bits wide is described as having an "×64" or "by 64" organization. Similarly, a memory module having 72-bit-wide ranks is described as having an "×72" or "by 72" organization.

The memory capacity of a memory module increases with the number of memory devices. The number of memory devices of a memory module can be increased by increasing the number of memory devices per rank or by increasing the number of ranks. Rather than referring to the memory capacity of the memory module, in certain circumstances, the memory density of the memory module is referred to instead.

During operation, the ranks of a memory module are selected or activated by control signals that are received from the processor. Examples of such control signals include, but are not limited to, rank-select signals, also called chip-select signals. Most computer and server systems support a limited number of ranks per memory module, which limits the memory density that can be incorporated in each memory module.

The memory space in an electronic system is limited by the physical addressable space that is defined by the number of address bits, or by the number of chips selected. In general, once the memory space is defined for an electronic system, it would not be feasible to modify the memory space without an extensive design change. This is especially true for the case in which a memory space is defined by a consortium, such as the Joint Electron Device Engineering Council (JEDEC). A problem arises when a user's application requires a larger addressable memory space than the memory space that the current electronic system is designed to support.

In developing a memory subsystem, consideration is always given to memory density, power dissipation (or thermal dissipation), speed, and cost. Generally, these attributes are not orthogonal to each other, meaning that optimizing one attribute may detrimentally affect another attribute. For example, increasing memory density typically causes higher power dissipation, slower operational speed, and higher costs.

Furthermore, the specifications of the memory subsystem may be guided by physical limitations associated with these attributes. For example, high thermal dissipation may limit the speed of the operation, or the physical size of the memory module may limit the density of the module.

These attributes generally dictate the design parameters of the memory module, usually requiring that the memory system slow down operation speed if the memory subsystem is populated with more memory devices to provide higher density memory cards.

SUMMARY

In certain embodiments, a memory module is configured to communicate with a memory controller of a computer system via a set of control signal lines and a plurality of sets of data signal lines. Each set of the plurality sets of data signal lines is a byte wide. The memory module has a width of, for example, 32 bits, 64 bits, 72 bits, 128 bits, or 256 bits, etc., and comprises a printed circuit board configured to be coupled to the memory controller via the set of control signal lines and the first number of data signal lines. The module board is mountable in a memory socket of the computer system and has an edge connector comprising a plurality of electrical contacts which are positioned on an edge of the PCB and are positioned to be releasably coupled to corresponding contacts of the memory socket. The memory module further comprises memory devices such as double data rate dynamic random access memory (DDR DRAM) devices coupled to the module board and arranged in multiple ranks each of the same width (i.e., N bits) as the memory module. The memory module further comprises a module controller coupled to the module board and operatively coupled to the DDR DRAM devices via a set of registered control lines. The module controller is configured to register input address and control signals for a read or write operation received from the memory controller via the set of address and control signal lines, and to output registered address and control signals onto the set of registered control lines. The read or write operation being targeted at a specific N-bit-wide rank of the multiple N-bit-wide ranks, so that the specific N-bit-wide rank is an only N-bit-wide rank among the multiple N-bit wide ranks selected to perform the read or write operation. The module controller is further configured to output a set of module control signals in response to the input control signals, the set of module control signals including signals that are dependent on which of the multiple N-bit-wide ranks is the specific N-bit-wide rank.

The memory module further comprises a plurality of byte-wise buffers coupled to the circuit board and configured to receive the second module control signals. Each respective byte-wise buffer of the plurality of byte-wise buffers is coupled to a respective set of the plurality of sets of data signal lines and to at least one respective DDR DRAM

device in each of the multiple ranks. The plurality of byte-wise buffers are disposed on the module board at respective positions corresponding to respective sets of the plurality of sets of data signal lines. In certain embodiments, each respective byte-wise buffer includes data paths and logic controlling the data paths in response to the second module control signals so that each respective byte-wise buffer is configured to actively drive a respective byte-wise section of each N-bit wide data signal associated with the read or write operation between the respective set of the plurality of sets of data signal lines and the at least one respective DDR DRAM device in the specific N-bit-wide rank. In certain embodiments, the each respective byte-wise buffer includes configurable data paths and logic that configures the data paths in response to the second module control signals so as to enable a respective byte-wise section of each N-bit wide data signal associated with the read or write operation be communicated between the memory controller and the at least one respective DDR DRAM device in the specific N-bit-wide rank. The logic configures the data paths differently depending on which of the multiple N-bit-wide ranks is performing the read or write operation.

In certain embodiments, the byte-wise buffers are configured to present to the memory controller one DDR DRAM device load on each data line of the plurality of sets of data lines during a write operation. In certain embodiments, the control circuit controls the byte-wise buffers in accordance with a CAS latency parameter.

BRIEF DESCRIPTION OF THE DRAWINGS

A complete understanding of the present invention may be obtained by reference to the accompanying drawings, when considered in conjunction with the subsequent, detailed description, in which:

FIG. **1**A is a schematic representation of a conventional memory subsystem populated with at least one JEDEC-standard two-rank memory module;

FIG. **1**B is a schematic representation of a conventional memory subsystem populated with at least one JEDEC-standard four-rank memory module.

FIG. **2**A is a schematic representation of another conventional memory subsystem populated with at least one two-rank memory module.

FIG. **2**B is a schematic representation of another conventional memory subsystem populated with at least one four-rank memory module.

FIGS. **2**C and **2**D schematically illustrate a conventional two-rank memory module and a four-rank memory module, respectively, each comprising a memory buffer.

FIG. **3**A is a schematic representation of an example memory subsystem in accordance with an embodiment of the disclosure.

FIG. **3**B schematically illustrates another example memory subsystem in accordance with certain embodiments described herein.

FIG. **3**C schematically illustrates an example layout of the memory devices, the data transmission circuits, and the control circuit of a memory module in accordance with certain embodiments described herein.

FIG. **3**D is a photograph of an example memory subsystem in accordance with certain embodiments described herein.

FIG. **4**A schematically illustrates an example memory subsystem comprising a data transmission circuit with a bit width which is the same as that of the individual memory devices.

FIG. **4**B schematically illustrates an example memory subsystem comprising a data transmission circuit with a bit width different from that as the individual memory devices.

FIG. **5** is a schematic representation of an example embodiment of a data transmission circuit compatible with the memory subsystem of FIG. **3**A.

FIG. **6** is an example timing diagram illustrating operation of the memory system of FIGS. **3**A and **5**.

For purposes of clarity and brevity, like elements and components bear like designations and numbering throughout the figures.

DETAILED DESCRIPTION

One method for increasing memory space is based on an address decoding scheme. This method is very widely adopted in the electronics industry in designing Application-Specific Integrated Circuit (ASIC) and System-On-Chip (SOC) devices to expand system memories. Another method increases the addressable memory space without extensive alteration of the software or hardware of an existing electronics system. This method combines chip-select signals with an address signal to increase the number of physically addressable memory spaces (e.g., by a factor of 2, by a factor of 4, by a factor of 8, or by other factors as well).

These methods have several shortcomings. For example, since these methods increase the addressable memory space by directly adding memory chips, a heavier load is presented to the outputs of the system controller and the outputs of the memory devices, resulting in a slower system. Also, increasing the number of memory devices results in higher power dissipation. In addition, since an increase in the number of memory devices on each memory module alters the physical properties of the memory module while the system board remains the same, the overall signal (transmission line) wave characteristics deviate from the original design intent or specification. Furthermore, especially when registered DIMMs (RDIMMs) are used, the increase in the number of the memory devices translates to an increase in the distributed RC load on the data paths, but not on the control paths (e.g., address paths), thereby introducing uneven signal propagation delay between the data signal paths and control signal paths. As used herein, the terms "control lines" and "control paths" include address lines or paths and command lines or paths, and the term "control signals" includes address signals and command signals.

FIGS. **1**A and **1**B illustrate a prior art approach of increasing the number of memory devices. Specifically, FIG. **1**A shows a conventional memory subsystem **100** with at least one JEDEC-standard two-rank memory module **110**, such as a registered dual inline memory module (RDIMM), only one of which is shown for clarity. Each rank of the memory module **110** comprises a plurality of memory devices **112**, such as dynamic random access memory (DRAM) devices or synchronous DRAM (SDRAM) devices. A register **130** receives a plurality of control lines **140** (shown as a single solid line) from the system memory controller **120** and is connected via control lines **142** to the memory devices **112** of each rank of the memory module **110**. This memory subsystem **100** connects each data line of an array of data lines **150** (shown as dashed lines) from a system memory controller **120** to corresponding memory devices **112** in the two ranks in each memory module **110**. Therefore, during a write operation, the system memory controller **120** sees all the memory devices **112** as its load via the data lines **150**, and during a read operation, each memory

device **112** sees multiple other memory devices **112**, as well as the system memory controller **120**, as its load via the data lines **150**.

FIG. **1**B is a schematic view of another conventional memory subsystem **100'** with at least one JEDEC-standard four-rank memory module **110'** (only one of which is shown for clarity), each rank comprising a plurality of memory devices **112'**. The register **130'** receives the plurality of control lines **140'** (shown as a single solid line) from the system memory controller **120'** and is connected via control lines **142'** to the memory devices **112'** of each rank of the memory module **110'**. Each data line of the array of data lines **150'** (shown as dashed lines) from the system memory controller **120'** is connected (e.g., by four fanouts) to corresponding memory devices **112'** in the four ranks in each memory module **110'**. Therefore, as with the two-rank memory module **110** shown in FIG. **1**A, during a write operation, the system memory controller **120'** sees all the memory devices **112'** as its load via the data lines **150'**, and during a read operation, each memory device **112'** sees multiple other memory devices **112'** and the system memory controller **120'** as its load via the data lines **150'**.

For both the conventional two-rank memory module **110** and the conventional four-rank memory module **110'**, the multiple loads seen by the memory controller **120**, **120'** during write operations and the multiple loads seen by the memory devices **112**, **112'** during read operations cause significant performance issues. For example, for synchronous operation, time delays of the various signals are desired to be substantially equal to one another such that the operation of the memory module **110**, **110'** is synchronized with the system bus of the computer system. Thus, the trace lengths of the memory module **110**, **110'** are selected such that the signals are at the same clock phase. For example, the lengths of the control lines **142**, **142'** from the register **130**, **130'** to each of the memory devices **112**, **112'** are substantially equal to one another. However, for faster clock speeds, small errors in the trace lengths make such synchronous operation difficult or impossible. Therefore, these prior art techniques not only reduce the speed of the memory systems, but they also require hardware modifications to minimize any deviation of the transmission line wave characteristics from the original design specification.

FIGS. **2**A and **2**B illustrate another prior art approach of increasing the number of memory devices. Specifically, FIG. **2**A shows a conventional memory subsystem **200** with at least one two-rank memory module **210**, only one of which is shown for clarity. Each rank of the memory module **210** comprises a plurality of memory devices **212**, such as dynamic random access memory (DRAM) devices or synchronous DRAM (SDRAM) devices. A register **230** receives a plurality of control lines **240** (shown as a single solid line) from the system memory controller **220** and is connected via control lines **242** to the memory devices **212** of each rank of the memory module **210**. This memory subsystem **200** connects each data line of an array of data lines **250** (shown as dashed lines) from a system memory controller **220** to corresponding memory devices **212** in the two ranks in each memory module **210**. Therefore, during a write operation, the system memory controller **220** sees all the memory devices **212** as its load via the data lines **250**, and during a read operation, each memory device **212** sees multiple other memory devices **212**, as well as the system memory controller **220**, as its load via the data lines **250**.

FIG. **2**B is a schematic view of another conventional memory subsystem **200'** with at least one four-rank memory module **210'** (only one of which is shown for clarity), each rank comprising a plurality of memory devices **212'**. The register **230'** receives the plurality of control lines **240'** (shown as a single solid line) from the system memory controller **220'** and is connected via control lines **242'** to the memory devices **212'** of each rank of the memory module **210'**. Each data line of the array of data lines **250'** (shown as dashed lines) from the system memory controller **220'** is connected (e.g., by four fanouts) to corresponding memory devices **212'** in the four ranks in each memory module **210'**. Therefore, as with the two-rank memory module **210** shown in FIG. **2**A, during a write operation, the system memory controller **220'** sees all the memory devices **212'** as its load via the data lines **250'**, and during a read operation, each memory device **212'** sees multiple other memory devices **212'** and the system memory controller **220'** as its load via the data lines **250'**.

For the memory modules **210**, **210'**, the control lines **242**, **242'** have a "flyby" configuration. In such a configuration, control signals are sent along the control lines **242**, **242'** (e.g., in a single-path daisy-chain) from the register **230**, **230'** to the memory devices **212**, **212'** of a given rank. These control signals reach each memory device **212**, **212'** of the rank sequentially, with the control signals first reaching the memory device **212**, **212'** having the shortest control line **242**, **242'**, then reaching the memory device **212**, **212'** having the next-shortest control line **242**, **242'**, and so on. For example, a control signal may reach the memory device **212**, **212'** having the longest control line **242**, **242'** a significant period of time after the same control signal reaches the memory device **212**, **212'** having the shortest control line **242**, **242'**. For synchronous operation, the memory subsystems **200**, **200'** have the data lines **250**, **250'** configured so that the time delays of the various data signals between the memory controller **220**, **220'** and the particular memory devices **212**, **212'** are substantially tailored such that the data signals and the control signals reach the particular memory device **212**, **212'** so that operation of the memory module **210**, **210'** is synchronized with the system bus of the computer system. Such "fly-by" configurations have been described as operating in "local sync" while having "global async."

For such "fly-by" configurations, the memory controller **220**, **220'** of FIGS. **2**A and **2**B is more complicated than the memory controller **120**, **120'** of FIGS. **1**A and **1**B in that the memory controller **220**, **220'** accounts for the time delays between the various memory devices **212**, **212'** and adjusts the timing of these signals appropriately for synchronous operation. However, in some situations, the clock cycle time is approximately equal to or less than the time difference (e.g., about 900 picoseconds) between the control signals reaching the memory device **212**, **212'** having the longest control line **242**, **242'** and reaching the memory device **212**, **212'** having the shortest control line **242**, **242'**. Under such situations, synchronous operation is not achievable. Thus, the time difference between the control signals reaching the memory devices **212**, **212'** at the extremes of the control lines **242**, **242'** provide a limit to the clock speed with which the memory module **210**, **210'** may be operated. These time differences, which can be more than one clock cycle, will limit the operational speed and performance of the memory module. In addition, as with the memory subsystems **100**, **100'** of FIGS. **1**A and **1**B, the "fly-by" memory subsystems **200**, **200'** of FIGS. **2**A and **2**B suffer from large loads which result in slower clock speeds.

One recent suggestion for the "fly-by" configurations is to provide a memory buffer which handles both the control signals and the data signals. FIGS. **2**C and **2**D schematically

illustrate a conventional two-rank memory module **310** and a four-rank memory module **310'**, respectively, each comprising a memory buffer **330**, **330'**. The control lines **340**, **340'** provide conduits for control signals from the memory controller **320**, **320'** to the memory buffer **330**, **330'**, and the control lines **342**, **342'** provide conduits for control signals from the memory buffer **330**, **330'** to the memory devices **312**, **312'**. The plurality of data lines **350**, **350'** (shown as one dashed line for clarity) provide conduits for data signals from the memory controller **320**, **320'** to the memory buffer **330**, **330'**, and data lines (not shown for clarity) on the memory module **310**, **310'** provide conduits for data signals from the memory controller **320**, **320'** to the memory devices **312**, **312'**.

The configurations of FIGS. **2**C and **2**D seek to have both the data signals and the control signals going to the memory buffer **330**, **330'**. However, such configurations have significant drawbacks. To send the data signals to the various memory devices **312**, **312'**, the memory module **310**, **310'** includes an extremely large number of data lines (not shown for clarity) coupling the memory buffer **330**, **330'** to the memory devices **312**, **312'**. For example, in certain circumstances, the memory buffer **330**, **330'** for an LRDIMM is a 628-pin device, which is extremely large. In addition, the logistics of tailoring the time delays of these many data lines is complicated or difficult to provide the desired timing of data signals from the memory buffer **330**, **330'** to the memory devices **312**, **312'**. Also, the memory module **310**, **310'** utilizes significant modifications of the memory controller **320**, **320'** since the memory buffer **330**, **330'** is taking over some of the control of data signal timing that conventional memory controllers have. Even so, the memory modules **310**, **310'** of FIGS. **2**C and **2**D can only operate in asynchronous mode, not synchronous mode, due to the long fly-by times as compared to the desired clock frequencies. For example, for a fly-by delay of 1 nanosecond, if the data rate is 1 Gb/second, there is the possibility of collisions on the data lines during read/write turnaround. To combat such collisions, the data rate can be slowed down or "dead" cycles can be inserted. The memory module **310**, **310'**, as a single unit, cannot be operated in synchronous mode, but operates as locally synchronous, globally (DIMM level) asynchronous.

FIG. **3**A schematically illustrates an example memory subsystem **400** with loadreduced memory modules **402** in accordance with certain embodiments described herein. FIG. **3**B schematically illustrates another example memory subsystem **400'** with loadreduced memory modules **402'** in accordance with certain embodiments described herein. FIG. **3**C schematically illustrates an example layout of the memory devices **412'**, the data transmission circuits **416'**, and the control circuit **430'** of a memory module **402'** in accordance with certain embodiments described herein. FIG. **3**D is a photograph of an example memory subsystem in accordance with certain embodiments described herein. In FIGS. **3**A-**3**C, control lines (e.g., address and control lines **440**, **440'** coupling the system memory controller **420**, **420'** to the memory modules **410**, **410'**) are shown as dashed lines, data lines (e.g., data lines **450**, **450'** coupling the system memory controller **420**, **420'** to the memory modules **410**, **410'**) are shown as solid lines, and in FIGS. **3**A and **3**B, input/output connections are shown as black dots. In certain embodiments, as schematically illustrated by FIGS. **3**A-**3**C, the address and control lines **440**, **440'** coupling the system memory controller **420**, **420'** to the memory module **410**, **410'** (e.g., to the control circuit **430**, **430'**) are separate from the data lines **450**, **450'** coupling the system memory controller **420**, **420'** to the memory module **410**, **410'** (e.g., to the data transmission circuits **416**, **416'**). In certain embodiments, the memory subsystem **400**, **400'** is designed, for example, to deliver higher speed and higher memory density with lower thermal dissipation as compared with conventional memory subsystems. In the following discussion, aspects of the example subsystem **400** and corresponding components (e.g., memory modules **402**, memory devices **412**A, **412**B, **412**C, **412**D, data transmission circuits **416**, control circuit **430**) and of the example subsystem **400'** and corresponding components (e.g., memory modules **402'**, memory devices **412'**A_1, **412'**A_2, **412'**B_1, **412'**B_2, **412'**C_1, **412'**C_2, **412'**D_1, **412'**D_2, data transmission circuits **416'**, control circuit **430'**) should be understood to apply to certain other embodiments as well.

As schematically illustrated in FIGS. **3**A and **3**B, the example memory module **402**, **402'** comprises at least one printed circuit board **410**, **410'** and a plurality of memory devices **412**, **412'** mechanically coupled to the at least one printed circuit board **410**, **410'**. The memory module **402**, **402'** further comprises a control circuit **430**, **430'** mechanically coupled to the at least one printed circuit board **410**, **410'**. The control circuit **430**, **430'** is configurable to receive control signals from the system memory controller **420**, **420'** and to transmit module control signals to the plurality of memory devices **412**, **412'**. The memory module **402**, **402'** further comprises a plurality of data transmission circuits **416**, **416'** mechanically coupled to the at least one printed circuit board **410**, **410'** and distributed at corresponding positions relative to the at least one printed circuit board **410**, **410'**. The plurality of data transmission circuits **416**, **416'** is configurable to be operatively coupled to the system memory controller **420**, **420'** and configurable to receive module control signals from the control circuit **430**, **430'**. At least one first data transmission circuit of the plurality of data transmission circuits **416**, **416'** is operatively coupled to at least two memory devices of the plurality of memory devices **412**, **412'**. At least one second data transmission circuit of the plurality of data transmission circuits **416**, **416'** is operatively coupled to at least two memory devices of the plurality of memory devices **412**, **412'**. The at least one first data transmission circuit is configurable to respond to the module control signals by selectively allowing or inhibiting data transmission between the system memory controller **420**, **420'** and at least one selected memory device of the at least two memory devices operatively coupled to the at least one first data transmission circuit. The at least one second data transmission circuit is configurable to respond to the module control signals by selectively allowing or inhibiting data transmission between the system memory controller **420**, **420'** and at least one selected memory device of the at least two memory devices operatively coupled to the at least one second data transmission circuit.

As shown in FIGS. **3**A and **3**B, the memory subsystem **400**, **400'** is configurable to be operationally coupled to a system memory controller **420**, **420'**, of a type well-known in the art (e.g., Intel Nehalem EP, EX chipsets; AM D Opteron chipset). The memory subsystem **400**, **400'** typically comprises one or more memory modules **402**, **402'**, such as DIMMs or RDIMMs, additional details of which are shown only for one for clarity. Various types of memory modules **402**, **402'** are compatible with embodiments described herein. For example, memory modules having memory capacities of 512 MB, 1 GB, 2 GB, 4 GB, 8 GB, as well as other capacities, are compatible with embodiments described herein. In addition, memory modules having widths of 4 bytes, 8 bytes, 9 bytes, 16 bytes, 32 bytes, or 32

bits, 64 bits, 72 bits, 128 bits, 256 bits, as well as other widths (in bytes or in bits), are compatible with embodiments described herein. Furthermore, memory modules **402**, **402'** compatible with embodiments described herein include, but are not limited to, single in-line memory modules (SIMMs), dual in-line memory modules (DIMMs), small-outline DIMMs (SO-DIMMs), unbuffered DIMMs (UDIMMs), registered DIMMs (RDIMMs), fully-buffered DIMMs (FBDIMMs), mini-DIMMs, and micro-DIMMs.

The one or more memory modules **402**, **402'** comprise one or more printed circuit boards (PCBs) **410**, **410'**, which may be arranged in a vertical stack (as shown), or in a back-to-back array. Each memory module **402**, **402'** in certain embodiments comprises a single PCB **410**, **410'**, while in certain other embodiments, each of one or more of the memory modules **402** comprises multiple PCBs **410**, **410'**. In some embodiments, the PCBs **410**, **410'** are mountable in module slots (not shown) of the computer system. A PCB **410**, **410'** of certain such embodiments has at least one edge connector **411** comprising a plurality of electrical contacts which are positioned on an edge of the PCB **410**, **410'** (as shown in FIG. **3***d*) and are configured to be releasably coupled to corresponding contacts of a computer system socket to provide electrical conductivity between the system memory controller **420**, **420'** and the various components of the memory modules **402**, **401'** on the PCBs **410**, **410'**.

At least one memory module **402**, **402'** comprises a plurality of memory devices **412**, **412'** (such as DRAMs or SDRAMs). The memory devices **412**, **412'** of the memory module **402**, **402'** may advantageously be arranged in a plurality of rows or ranks. Memory devices **412**, **412'** compatible with embodiments described herein include, but are not limited to, random-access memory (RAM), dynamic random-access memory (DRAM), synchronous DRAM (SDRAM), and double-data-rate DRAM (e.g., DDR, DDR2, DDR3, etc.). In addition, memory devices **412**, **412'** having bit widths of 4, 8, 16, 32, as well as other bit widths, are compatible with embodiments described herein. Memory devices **412**, **412'** compatible with embodiments described herein have packaging which include, but are not limited to, thin small-outline package (TSOP), ball-grid-array (BGA), fine-pitch BGA (FBGA), micro-BGA (OGA), mini-BGA (mBGA), and chip-scale packaging (CSP).

In certain embodiments, the memory devices **412**, **412'** of the memory module **402**, **402'** are arranged in four ranks, although embodiments with less than four ranks (e.g., one rank, two ranks, three ranks) or more than four ranks (e.g., 6 ranks, 8 ranks) per memory module **402**, **402'** may be employed. In certain embodiments, each rank comprises eight or nine memory modules, while in certain other embodiments, other numbers of memory modules per rank may also be used. In certain embodiments, as schematically shown in FIG. **3**A, the memory devices **412** are arranged in four ranks, denoted A, B, C, and D, and each rank comprises n memory devices. For the sake of this disclosure, in the example memory subsystem **400** of FIG. **3**A, rank A comprises memory devices **412**A_1, **412**A_2, . . . , **412**A_n; rank B comprises memory devices **412**B_1, **412**B_2, . . . , **412**B_n; rank C comprises memory devices **412**C_1, **412**C_2, . . . , **412**C_n; and rank D comprises memory devices **412**D_1, **412**D_2, **412**D_n. For the sake of this disclosure, in the example memory subsystem **400'** of FIG. **3**B, rank A comprises memory devices **412'**A_1, **412'**A_2, . . . , **412'**A_n; rank B comprises memory devices **412'**B_1, **412'**B_2, . . . , **412'**B_n; rank C comprises memory devices **412'**C_1, **412'**C_2, **412'**C_n; and rank D comprises memory devices **412'**D_1, **412'**D_2, . . . , **412'**D_n.

In certain embodiments, at least one memory module **402**, **402'** comprises one or more electrical components (not shown) which may be mounted on the PCB **410**, **410'**, within the PCB **410**, **410'**, or both on and within the PCB **410**, **410'**, and are operationally coupled to one another and to the plurality of memory devices **412**, **412'**. For example, the electrical components may be surface-mounted, through-hole mounted, embedded or buried between layers of the PCB **410**, **410'**, or otherwise connected to the PCB **410**, **410'**. These electrical components may include, but are not limited to, electrical conduits, resistors, capacitors, inductors, transistors, buffers, registers, logic elements, or other circuit elements. In certain embodiments, at least some of these electrical components are discrete, while in other certain embodiments, at least some of these electrical components are constituents of one or more integrated circuits.

In certain embodiments, at least one memory module **402**, **402'** comprises a control circuit **430**, **430'** configured to be operatively coupled to the system memory controller **420**, **420'** and to the memory devices **412**, **412'** of the memory module **402**, **402'** (e.g., via lines **442**, **442'**). In certain embodiments, the control circuit **430**, **430'** may include one or more functional devices, such as a programmable-logic device (PLD), an applicationspecific integrated circuit (ASIC), a field-programmable gate array (FPGA), a customdesigned semiconductor device, or a complex programmable-logic device (CPLD). In certain embodiments, the control circuit **430**, **430'** may comprise one or more custom devices. In certain embodiments, the control circuit **430**, **430'** may comprise various discrete electrical elements; while in other embodiments, the control circuit **430**, **430'** may comprise one or more integrated circuits.

The control circuit **430**, **430'** of certain embodiments is configurable to be operatively coupled to control lines **440**, **440'** to receive control signals (e.g., bank address signals, row address signals, column address signals, address strobe signals, and rank-address or chip-select signals) from the system memory controller **420**, **420'**. The control circuit **430**, **430'** of certain embodiments registers signals from the control lines **440**, **440'** in a manner functionally comparable to the address register of a conventional RDIMM. The registered control lines **440**, **440'** are also operatively coupled to the memory devices **412**, **412'**. Additionally, the control circuit **430**, **430'** supplies control signals for the data transmission circuits **416**, **416'** (e.g., via lines **432**, **432'**), as described more fully below. The control signals indicate, for example, the direction of data flow, that is, to or from the memory devices **412**, **412'**. The control circuit **430**, **430'** may produce additional chip-select signals or output enable signals based on address decoding. Examples of circuits which can serve as the control circuit **430**, **430'** are described in more detail by U.S. Pat. Nos. 7,289,386 and 7,532,537, each of which is incorporated in its entirety by reference herein.

In certain embodiments, at least one memory module **402**, **402'** comprises a plurality of data transmission circuits **416**, **416'** mounted on the one or more PCBs **410**, **410'**, within the one or more PCBs **410**, **410'**, or both on and within the one or more PCBs **410**, **410'**. The plurality of data transmission circuits **416**, **416'** are operatively coupled to the control circuit **430**, **430'** (e.g., via lines **432**, **432'**), and configured to be operatively coupled to the system memory controller **420**, **420'** (e.g., via the data lines **450**, **450'**) upon operatively coupling the memory module **402**, **402'** to the computer system. In certain embodiments, these data transmission circuits **416**, **416'** can be referred to as "load-reducing circuits" or "load-reducing switching circuits." As used herein, the terms "load-reducing" or "load-reducing switch-

ing" refer to the use of the data transmission circuits **416**, **416'** to reduce the load seen by the system memory controller **420**, **420'** when operatively coupled to the memory module **402**, **402'**. In certain embodiments, as schematically illustrated by FIG. **3**A, the memory module **402** comprises n data transmission circuits **416**, where n is the number of memory devices per rank of the memory module **410**. For example, as schematically shown in FIG. **3**A, the memory devices **412** of the memory module **410** are arranged in four ranks of n memory devices each, and the memory module **410** comprises at least a first data transmission circuit **416**$_1$ and a second data transmission circuit **416**$_2$. The first data transmission circuit **416**$_1$ of certain such embodiments is operatively coupled to at least one memory device **412** of each rank (e.g., memory devices **412**A$_1$, **412**B$_1$, **412**C$_1$, **412**D$_1$). The second data transmission circuit **416**$_2$ of certain such embodiments is operatively coupled to at least one memory device **412** of each rank (e.g., memory devices **412**A$_2$, **412**B$_2$, **412**C$_2$, **412**D$_2$). In certain embodiments, as schematically illustrated by FIG. **3**B, the memory module **402'** comprises n/2 data transmission circuits **416'**, where n is the number of memory devices per rank of the memory module **410'**. For example, as schematically shown in FIG. **3**B, the memory devices **412'** of the memory module **410'** are arranged in four ranks of n memory devices each, and the memory module **410'** comprises at least a first data transmission circuit **416**$_1$ and a second data transmission circuit **416'2**. The first data transmission circuit **416**$_1$ of certain such embodiments is operatively coupled to at least two memory devices **412'** of each rank (e.g., memory devices **412**A$_1$, **412**A$_2$, **412**B$_1$, **412**B$_2$, **412**C$_1$, **412**C$_2$, **412**D$_1$, **412**D$_2$). The second data transmission circuit **416**$_2$ of certain such embodiments is operatively coupled to at least two memory devices **412'** of each rank (e.g., memory devices **412'**A$_3$, **412'**A$_4$, **412'**B$_3$, **412'**B$_4$, **412'**C$_3$, **412'**C$_4$, **412'**D$_3$, **412'**D$_4$). In certain embodiments, at least one data transmission circuit **416**, **416'** selectively switches between two or more memory devices **412**, **412'** so as to operatively couple at least one selected memory device **412**, **412'** to the system memory controller **420**, **420'** (e.g., the data transmission circuit **416**, **416'** is configurable to respond to module control signals by selectively allowing or inhibiting data transmission between the system memory controller **420**, **420'** and at least one selected memory device **412**, **412'**).

In certain such embodiments, the at least one data transmission circuit **416**, **416'** selectively operatively couples two selected memory devices to the system memory controller **420**, **420'**. For example, as schematically shown in FIG. **3**A, the first data transmission circuit **416**$_1$ is configurable to respond to module control signals by selectively allowing or inhibiting data transmission between the system memory controller **420** and either selected memory devices **412**A$_1$ and **412**C$_1$ or selected memory devices **412**B$_1$ and **412**D$_1$), and the second data transmission circuit **416**$_2$ is configurable to respond to module control signals by selectively allowing or inhibiting data transmission between the system memory controller **420** and either selected memory devices **412**A$_2$ and **412**C$_2$ or selected memory devices **412**B$_2$ and **412**D$_2$). Conversely, in a conventional memory module without the data transmission circuits **416**, the two or more memory devices **412** (e.g., memory devices **412**A$_1$, **412**B$_1$, **412**C$_1$, **412**D$_1$) are concurrently operatively coupled to the system memory controller **420**. A data transmission circuit **416** of certain embodiments bidirectionally buffer data signals between the memory controller **420** and the memory devices **412** corresponding to the data transmission circuit **416**. For another example, as schematically shown in FIG. **3**B, the first data transmission circuit **416'**$_1$ is configurable to respond to module control signals by selectively allowing or inhibiting data transmission between the system memory controller **420'** and either selected memory devices **412'**A$_1$ and **412'**C$_1$ or selected memory devices **412'**B$_1$ and **412'**D$_1$ and either selected memory devices **412'**A$_2$ and **412'**C$_2$ or selected memory devices **412'**B$_2$ and **412'**D$_2$), and the second data transmission circuit **416'**$_2$ is configurable to respond to module control signals by selectively allowing or inhibiting data transmission between the system memory controller **420'** and either selected memory devices **412'**A$_3$ and **412'**C$_3$ or selected memory devices **412'**B$_3$ and **412'**D$_3$ and either selected memory devices **412'**A$_4$ and **412'**C$_4$ or selected memory devices **412'**B$_4$ and **412'**D$_4$).

In certain embodiments, two or more of the data transmission circuits **416**, **416'** are mechanically coupled to the at least PCB **410**, **410'** at corresponding positions which are separate from one another. For example, as schematically illustrated by FIG. **3**A, the first data transmission circuit **416**$_1$ and the second data transmission circuit **416**$_2$ are at corresponding positions which are separate from one another (e.g., the package containing the first data transmission circuit **416**$_1$ is at a location spaced from the location of the package containing the second data transmission circuit **416**$_2$). For another example, as schematically illustrated by FIG. **3**B, the first data transmission circuit **416'**$_1$ and the second data transmission circuit **416'**$_2$ are at corresponding positions which are separate from one another (e.g., the package containing the first data transmission circuit **416'**$_1$ is at a location spaced from the location of the package containing the second data transmission circuit **416'**$_2$). In certain such embodiments, two or more of the data transmission circuits **416**, **416'** are distributed across a surface of the PCB **410**, **410'** of the memory module **402**, **402'**. In certain embodiments, the corresponding positions of two or more data transmission circuits **416**, **416'** (e.g., first data transmission circuit **416**$_1$ and second data transmission circuit **416**$_2$ of FIG. **3**A or first data transmission circuit **416'**$_1$ and second data transmission circuit **416'**$_2$ of FIG. **3**B) are along an edge **411**, **411'** of the at least one PCB **410**, **410'** such that a data transmission circuit **416**, **416'** is located substantially between the edge **411**, **411'** and at least some of the at least two memory devices **412**, **412'** to which the data transmission circuit **416**, **416'** is operatively coupled. For example, as schematically illustrated by FIG. **3**A, the first data transmission circuit **416**$_1$ is located substantially between the edge **411** and the memory devices **412**A$_1$, **412**B$_1$, **412**C$_1$, **412**D$_1$ to which the first data transmission circuit **416**$_1$ is operatively coupled, and the second data transmission circuit **416**$_2$ is located substantially between the edge **411** and the memory devices **412**A$_2$, **412**B$_2$, **412**C$_2$, **412**D$_2$ to which the second data transmission circuit **416**$_1$ is operatively coupled. For another example, as schematically illustrated by FIG. **3**B, the first data transmission circuit **416'**$_1$ is located substantially between the edge **411'** and the memory devices **412'**A$_1$, **412'**A$_2$, **412'**B$_1$, **412'**B$_2$, **412'**C$_1$, **412'**C$_2$, **412'**D$_1$, **412'**D$_2$ to which the first data transmission circuit **416'**$_1$ is operatively coupled, and the second data transmission circuit **416'**$_2$ is located substantially between the edge **411'** and the memory devices **412'**A$_3$, **412'**A$_4$ **412'**B$_3$, **412'**B$_4$, **412'**C$_3$, **412'**C$_4$, **412'**D$_3$, **412'**D$_4$ to which the second data transmission circuit **416'**$_2$ is operatively coupled.

FIGS. **3**C and **3**D illustrate the positioning of the data transmission circuits **416'** in accordance with certain embodiments described herein. In certain embodiments, the position of at least one of the data transmission circuits **416'**

is generally aligned with one or more of the memory devices **412'** to which the data transmission circuit **416'** is operatively coupled. For example, the one or more of the data transmission circuits **416'** and the memory devices **412'** to which it is operatively coupled can be positioned generally along a line that is substantially perpendicular to the edge **411'** of the PCB **410'**. In certain embodiments, the position of at least one of the data transmission circuits **416'** is generally offset from a line defined by the positions of the one or more of the memory devices **412'** to which the data transmission circuit **416'** is operatively coupled. For example, as shown in FIGS. **3**C and **3**D, the memory devices **412'** operatively coupled to a data transmission circuit **416'** can be positioned along a line that is substantially perpendicular to the edge **411'** of the PCB **410'** and the data transmission circuit **416'** can be generally offset from this line in a direction generally along the edge **411'** of the PCB **410'**. In certain such embodiments, the data transmission circuits **416'** are sufficiently small in width and breadth (e.g., 2.5 mm by 7.5 mm) to fit between the edge **411'** and the corresponding memory devices **412'** while maintaining the desired size of the memory module **400'**. Other positions and sizes of the separate data transmission circuits **416'** are also compatible with certain embodiments described herein. For example, in certain embodiments, one or more of the data transmission circuits **416**, **416'** can be positioned between two or more memory devices **412**, **412'**, or can be spaced away from an edge **411**, **411'** of the PCB **410**, **410'** with one or more memory devices **412**, **412'** between the edge **411**, **411'** and the one or more data transmission circuits **416**, **416'**.

In certain embodiments, the data transmission circuit **416** comprises or functions as a byte-wise buffer. In certain such embodiments, each of the one or more data transmission circuits **416** has the same bit width as does the associated memory devices **412** per rank to which the data transmission circuit **416** is operatively coupled. For example, as schematically illustrated by FIG. **4**A (which corresponds generally to FIG. **3**A), the data transmission circuit **416** can be operatively coupled to a single memory device **412** per rank, and both the data transmission circuit **416** and the memory device **412** per rank to which the data transmission circuit **416** is operatively coupled can each have the same bit width (e.g., 4 bits, 8 bits, or 16 bits). The data transmission circuit **416** of FIG. **4**A has a bit width of 8 bits, and receives data bits **0-7** from the system memory controller **420** and selectively transmits the data bits **0-7** to selected memory devices **412**A, **412**B, **412**C, **412**D in response to the module control signals from the control circuit **430**. Similarly, data transmission circuits **416'** of certain embodiments can function as a byte-wise buffer for associated memory devices **412'**A, **412'**B, **412'**C, **412'**D to which the data transmission circuits **416'** are operatively coupled in response to the module control signals from the control circuit **430'**.

In certain other embodiments, the bit widths of one or more of the memory devices **412** may be different from the bit widths of the one or more data transmission circuits **416** to which they are connected. For example, as schematically illustrated by FIG. **4**B (which corresponds generally to FIG. **3**B), the data transmission circuits **416** may have a first bit width (e.g., a bit width of 8 bits) and the memory devices **412** may have a second bit width which is less than the first bit width (e.g., one-half the first bit width, or a bit width of 4 bits), with each data transmission circuit **416** operatively coupled to multiple memory devices **412** per rank (e.g., two memory devices **412** in each rank). In certain such embodiments, the total bit width of the multiple memory devices **412** per rank connected to the circuit **416** equals the bit width of the circuit **416** (e.g., 4 bits, 8 bits, or 16 bits). The data transmission circuit **416** of FIG. **4**B has a total bit width of 8 bits, and receives data bits **0-7** from the system memory controller **420** and selectively transmits data bits **0-3** to a first memory device **412**A_1, **412**B_1, **412**C_1, **412**D_1 and data bits **4-7** to a second memory device **412**A_2, **412**B_2, **412**C_2, **412**D_2 in response to the module control signals from the control circuit **430**. Similarly, data transmission circuits **416'** of certain embodiments can function with different bit widths than those of the associated memory devices **412'**A_1, **412'**A_2, **412'**B_1, **412'**B_2, **412'**C_1, **412'**C_2, **412'**D_1, **412'**D_2 to which the data transmission circuits **416'** are operatively coupled in response to the module control signals from the control circuit **430'**.

In certain embodiments, by having the data transmission circuit **416** comprise or serve as a "byte-wise" buffer (e.g., as shown in the examples of FIGS. **4**A and **4**B), the data signals are synchronous with the synch clock. In addition, for certain such embodiments in which the memory module **400** experiences variations in one or more characteristics (e.g., temperature, voltage, manufacturing parameters), the memory module **400** can be designed to optimize the circuits of a smaller number of components as compared to other configurations which do not utilize byte-wide buffering (e.g., having four ranks of 8-bit memory devices and having two 4-bit buffers). In certain embodiments, the data transmission circuits **416** are used for bit slicing in which the data are defined in sections. For example, rather than defining data to be 64-bit-wide (e.g., [63:0]), the data can be defined or sliced in 16-bit-wide sections (e.g., [15:0], [31:16], [47:32], [63:48]). In certain such embodiments, not all the bits are grouped together and not all the bits produce the same behavior (e.g., logic- and/or time-wise).

One or more of the data transmission circuits **416**, in accordance with an embodiment of this disclosure, is operatively coupled to a corresponding one or more of the data lines **452** connected to one or more memory devices **412** in each of the ranks A, B, C, D. For example, in certain embodiments, each data transmission circuit **416** is connected to one or more data lines **452** connected to one corresponding memory device in each of the ranks (e.g., memory devices **204**A, **204**B, **204**C, and **204**D, as shown in FIG. **3**A). Each data line **450**, **452** thus carries data from the system memory controller **420**, through the data transmission circuits **416**, to the memory devices **204**A, **204**B, **204**C, **204**D connected to the data transmission circuits **416**. The data transmission circuits **416** of certain embodiments may be used to drive each data bit to and from the memory controller **420** and the memory devices **412**, instead of the memory controller **420** and the memory devices **412** directly driving each data bit to and from the memory controller **420** and the memory devices **412**. Specifically, as described in more detail below, one side of each data transmission circuit **416** of certain embodiments is operatively coupled to a memory device **412** in each rank (e.g., via data lines **452**), while the other side of the data transmission circuit **416** is operatively coupled to the corresponding data line **450** of the memory controller **420**.

To reduce the memory device loads seen by the system memory controller **420** (e.g., during a write operation), the data transmission circuit **416** of certain embodiments is advantageously configured to be recognized by the system memory controller **420** as a single memory load. This advantageous result is desirably achieved in certain embodiments by using the data transmission circuits **416** to electrically couple only the enabled memory devices **412** to the memory controller **420** (e.g., the one, two, or more memory

devices **412** to which data is to be written) and to electrically isolate the other memory devices **412** from the memory controller **420** (e.g., the one, two, or more memory devices **412** to which data is not to be written). Therefore, during a write operation in which data is to be written to a single memory device **412** in a rank of the memory module **400**, each data bit from the system memory controller **420** sees a single load from the memory module **400**, presented by one of the data transmission circuits **416**, instead of concurrently seeing the loads of all of the four memory devices **412**A, **412**B, **412**C, **412**D to which the data transmission circuit **416** is operatively coupled. In the example of FIG. **3**A, during a write operation in which data is to be written to two memory device **412** in two ranks (e.g., memory devices **412**A and **412**C or memory devices **412**B and **412**D), each data bit from the system memory controller **420** sees a single load from the memory module **402**, which is presented by one of the data transmission circuits **416**, instead of concurrently seeing the loads of all of the four memory devices **412**A, **412**B, **412**C, **412**D to which the data transmission circuits **416** is operatively coupled. In comparison to the standard JEDEC four-rank DEVINI configuration (see FIG. **2**A and FIG. **2**B), the memory system **402** of certain embodiments may reduce the load on the system memory controller **420** by a factor of four.

FIG. **5** schematically illustrates an example data transmission circuit **416** compatible with certain embodiments described herein. In one embodiment, the data transmission circuits **416** includes control logic circuitry **502** used to control the various components of the data transmission circuit **416**, which may include one or more buffers, one or more switches, and one or more multiplexers among other components. The illustrated embodiment of FIG. **5** is 1-bit wide and switches a single data line **518** between the memory controller **420** and the memory devices **412**. In other embodiments, the data transmission circuit **416** may be multiple bits wide, for example, 8 bits, and switch a corresponding number of data lines **518**. In a multiple bit wide embodiment, the control logic circuitry **502** may be shared over the multiple bits.

As a part of isolating the memory devices **412** from the system memory controller **420**, in one embodiment, the data transmission circuits **416** allow for "driving" write data and "merging" read data. In the operational embodiment shown in FIG. **5**, in a write operation, data entering a data transmission circuit **416** via a data line **518** is driven onto two data paths, labeled path A and path B, preferably after passing through a write buffer **503**. The ranks of memory devices **412** are likewise divided into two groups with one group associated with path A and one group associated with path B. As shown in FIG. **3**A, rank A and rank C are in the first group, and rank B and rank D are in the second group. Accordingly, the memory devices **412**A, **412**C of rank A and rank C are connected to the data transmission circuits **416** by a first one of the two data paths, and the memory devices **412**B, **412**D of rank B and rank D are connected to the data transmission circuits **416** by a second one of the two data paths. In other embodiments, the driving of write data and merging of read data may be performed over more than two data paths.

As is known, Column Address Strobe (CAS) latency is a delay time which elapses between the moment the memory controller **420** informs the memory modules **402** to access a particular column in a selected rank or row and the moment the data for or from the particular column is on the output pins of the selected rank or row. The latency may be used by the memory module to control operation of the data transmission circuits **416**. During the latency, address and control signals pass from the memory controller **420** to the control circuit **430** which produces controls sent to the control logic circuitry **502** (e.g., via lines **432**) which then controls operation of the components of the data transmission circuits **416**.

For a write operation, during the CAS latency, the control circuit **430**, in one embodiment, provides enable control signals to the control logic circuitry **502** of each data transmission circuit **416**, whereby the control logic circuitry **502** selects either path A or path B to direct the data. Accordingly, when the control logic circuitry **502** receives, for example, an "enable A" signal, a first tristate buffer **504** in path A is enabled and actively drives the data value on its output, while a second tristate buffer **506** in path B is disabled with its output in a high impedance condition. In this state, the data transmission circuit **416** allows the data to be directed along path A to a first terminal Y**1**, which is connected to and communicates only with the first group of the memory devices **412**, e.g., those in ranks A and C. Similarly, if an "enable B" signal is received, the first tristate **504** opens path A and the second tristate **506** closes path B, thus directing the data to a second terminal Y**2**, which is connected to and communicates only with the second group of the memory devices **412**, e.g., those in ranks B and D.

For a read operation, the data transmission circuit **416** operates as a multiplexing circuit. In the illustrated embodiment of FIG. **5**, for example, data signals read from the memory devices **412** of a rank are received at the first or second terminals Y**1**, Y**2** of the data transmission circuit **416**. The data signals are fed to a multiplexer **508**, which selects one to route to its output. The control logic circuitry **502** generates a select signal to select the appropriate data signal, and the selected data signal is transmitted to the system memory controller **420** along a single data line **518**, preferably after passing through a read buffer **509**. The read buffer **509** may be a tristate buffer that is enabled by the control logic circuitry **502** during read operations. In another embodiment, the multiplexer **508** and the read buffer **509** may be combined in one component. In yet another embodiment, the multiplexer **508** and the read buffer **509** operations may be split over two tristate buffers, one to enable the value from Y**1** to the data line **518** and another to enable the value from Y**2** to the data line **518**.

The data transmission circuits **416** present a load on the data lines **518** from the write buffer **503** and the read buffer **509**. The write buffer **503** is comparable to an input buffer on one of the memory devices **412**, and the read buffer **509** is comparable to an output buffer on one of the memory devices **412**. Therefore, the data transmission circuits **416** present a load to the memory controller **420** that is substantially the same as the load that one of the memory devices **412** would present. Similarly, the data transmission circuits **416** present a load on the first and second terminals Y**1**, Y**2** from the multiplexer **508** and the first tristate buffer **504** (on the first terminal Y**1**) and the second tristate buffer **506** (on the second terminal Y**2**). The multiplexer **508** is comparable in loading to an input buffer on the memory controller **420**, and the first and second tristate buffers **504**, **506** are each comparable to an output buffer on the memory controller **420**. Therefore, the data transmission circuits **416** present a load to the memory devices **412** that is substantially the same as the load that the memory controller **420** would present.

Additionally, the data transmission circuits **416** operate to ameliorate quality of the data signals passing between the memory controller **420** and the memory devices **412**. With-

out the data transmission circuits **416**, waveforms of data signals may be substantially degraded or distorted from a desired shape between source and sink. For example, signal quality may be degraded by lossy transmission line characteristics, mismatch between characteristics of transmission line segments, signal crosstalk, or electrical noise. However, in the read direction, the read buffer **509** regenerates the signals from the memory devices **412** thereby restoring the desired signal waveform shapes. Similarly, in the write direction, the first tristate buffer **504** and the second tristate buffer **506** regenerate the signals from the memory controller **420** thereby restoring the desired signal waveform shapes.

Referring again to FIG. **3**A, when the memory controller **420** executes read or write operations, each specific operation is targeted to a specific one of the ranks A, B, C, and D of a specific memory module **402**. The data transmission circuit **416** on the specifically targeted one of the memory modules **402** functions as a bidirectional repeater/multiplexor, such that it drives the data signal when connecting from the system memory controller **420** to the memory devices **412**. The other data transmission circuits **416** on the remaining memory modules **402** are disabled for the specific operation. For example, the data signal entering on data line **518** entering into data transmission circuit **416** is driven to memory devices **412**A and **412**C or **412**B and **412**C depending on which memory devices are active and enabled. The data transmission circuit **416** then multiplexes the signal from the memory devices **412**A, **412**B, **412**C, **412**D to the system memory controller **420**. The data transmission circuits **416** may each control, for example, a nibble-wide data path or a bytewide-data path. As discussed above, the data transmission circuits **416** associated with each module **402** are operable to merge data read signals and to drive data write signals, enabling the proper data paths between the system memory controller **420** and the targeted or selected memory devices **412**. Thus, the memory controller **420**, when there are four four-rank memory modules, sees four load-reducing switching circuit loads, instead of sixteen memory device loads. The reduced load on the memory controller **420** enhances the performance and reduces the power requirements of the memory system, as compared with, for example, the conventional systems described above with reference to FIGS. **1**A, **1**B and **2**A-**2**D.

Operation of a memory module using the data transmission circuit **416** may be further understood with reference to FIG. **6**, an illustrative timing diagram of signals of the memory module **402**. The timing diagram includes first through eighth time periods **601**-**608**. When the memory devices **404** are synchronous memories, each of the time periods **601**-**608** may correspond to one clock cycle of the memory devices **404**.

The first, second, and third time periods **601**-**603** illustrate write operations with data passing from the memory controller **401** to the memory module **402**. The fourth time period **604** is a transition between the write operations and subsequent read operations. The timing diagram shows a write operation to the first group of memory devices **412**A, **412**C connected to the first terminals Y**1** of the data transmission circuits **416** and a write operation to the second group of memory devices **412**B, **412**D connected to the second terminals Y**2** of the data transmission circuits **416**. Recalling the CAS latency described above, each write operation extends over two time periods in a pipelined manner.

The write to the first group of memory devices **412**A, **412**C appears in the first time period **601** when system address and control signals **440** pass from the memory controller **420** to the module controller **430**. The control circuit **430** evaluates the address and control signals **440** to determine that data is to be written to memory devices **412**A, **412**C in the first group. During the second time period **602**, the control circuit **430** supplies control signals to the control logic circuitry **502** to enable the first tristate buffer **504** and to disable the second tristate buffer **506** and the read buffer **509**. Thus, during the second time period **602**, data bits pass from the data lines **518** to the first terminal Y**1** and on to the memory devices **412**A, **412**C.

Similarly, the write to the second group of memory devices **412**A, **412**C appears in the second time period **602** when system address and control signals **440** pass from the memory controller **420** to the control circuit **430**. The control circuit **430** evaluates the address and control signals **440** to determine that data is to be written to memory devices **412**B, **412**D in the second group. During the third time period **603**, the control circuit **430** supplies control signals to the control logic circuitry **502** to enable the second tristate buffer **506** and to disable the first tristate buffer **504** and the read buffer **509**. Thus, during the third time period **603**, data bits pass from the data lines **518** to the second terminal Y**2** and on to the memory devices **412**B, **412**D.

The fifth, sixth, seventh, and eighth time periods **605**-**608** illustrate read operations with data passing to the memory controller **420** from the memory module **402**. The timing diagram shows a read operation from the first group of memory devices **412**A, **412**C connected to the first terminals Y**1** of the data transmission circuits **416** and a read operation from the second group of memory devices **412**B, **412**D connected to the second terminals Y**2** of the data transmission circuits **416**. Recalling the CAS latency described above, each read operation extends over two time periods in a pipelined manner.

The read from the first group of memory devices **412**A, **412**C appears in the fifth time period **605** when system address and control signals **440** pass from the memory controller **420** to the control circuit **430**. The control circuit **430** evaluates the address and control signals **440** to determine that data is to be read from memory devices **412**A, **412**C in the first group. During the sixth time period **606**, the control circuit **430** supplies control signals to the control logic circuitry **502** to cause the multiplexer **58** to select data from the first terminal Y**1**, to enable the read buffer **509**, and to disable the first tristate buffer **504** and the second tristate buffer **506**. Thus, during the sixth time period **606**, data bits pass from the memory devices **412**A, **412**C via the first terminal Y**1** to data lines **518** and on to the memory controller **420**.

The read from the second group of memory devices **412**B, **412**D appears in the seventh time period **607** when system address and control signals **440** pass from the memory controller **420** to the control circuit **430**. The control circuit **430** evaluates the address and control signals **440** to determine that data is to be read from memory devices **412**B, **412**D in the second group. During the eighth time period **608**, the control circuit **430** supplies control signals to the control logic circuitry **502** to cause the multiplexer **508** to select data from the second terminal Y**2**, to enable the read buffer **509**, and to disable the first tristate buffer **504** and the second tristate buffer **506**. Thus, during the eighth time period **606**, data bits pass from the memory devices **412**B, **412**D via the second terminal Y**2** to data lines **518** and on to the memory controller **420**.

Various embodiments have been described above. Although this invention has been described with reference to

these specific embodiments, the descriptions are intended to be illustrative of the invention and are not intended to be limiting. Various modifications and applications may occur to those skilled in the art without departing from the true spirit and scope of the invention as defined in the appended claims.

We claim:

1. A N-bit-wide memory module mountable in a memory socket of a computer system and configurable to communicate with a memory controller of the computer system via address and control signal lines and N-bit wide data signal lines, the N-bit wide data signal lines including a plurality of sets of data signal lines, each set of data signal lines is a byte wide, the memory module comprising:

a printed circuit board (PCB) having an edge connector comprising a plurality of electrical contacts which are positioned on an edge of the PCB and configured to be releasably coupled to corresponding contacts of the memory socket;

double data rate dynamic random access memory (DDR DRAM) devices coupled to the PCB and arranged in multiple N-bit-wide ranks;

a module controller coupled to the PCB and operatively coupled to the DDR DRAM devices, wherein the module controller is configurable to receive from the memory controller via the address and control signal lines input address and control signals for a memory write operation to write N-bit-wide write data from the memory controller into a first N-bit-wide rank of the multiple N-bit-wide ranks, and to output registered address and control signals in response to receiving the input address and control signals, wherein the registered address and control signals cause the first N-bit-wide rank to perform the memory write operation by receiving the N-bit-wide write data, wherein the module controller is further configurable to output module control signals in response to at least some of the input address and control signals; and

a plurality of byte-wise buffers coupled to the PCB and configured to receive the module control signals, wherein each respective byte-wise buffer of the plurality of byte-wise buffers has a first side configured to be operatively coupled to a respective set of data signal lines, a second side that is operatively coupled to at least one respective DDR DRAM device in each of the multiple N-bit-wide ranks via respective module data lines, and a byte-wise data path between the first side and the second side, wherein the each respective byte-wise buffer is disposed on the PCB at a respective position corresponding to the respective set of the plurality of sets of data signal lines;

wherein the each respective byte-wise buffer further includes logic configurable to control the byte-wise data path in response to the module control signals, wherein the byte-wise data path is enabled for a first time period in accordance with a latency parameter to actively drive a respective byte-wise section of the N-bit wide write data associated with the memory operation from the first side to the second side during the first time period; and

wherein the byte-wise data path includes first tristate buffers, and the logic in response to the module control signals is configured to enable the first tristate buffers to drive the respective byte-wise section of the N-bit wide write data to the respective module data lines during the first time period.

2. The memory module of claim **1**, wherein the DDR DRAM devices each has a bit width of 4 bits, wherein the at least one respective DDR DRAM device in each of the multiple N-bit-wide ranks includes a respective pair of DDR DRAM devices, and wherein a first subset of the first tristate buffers is enabled for the first time period to drive a first nibble of the respective byte-wise section of the N-bit wide write data to a first subset of the respective module data lines coupled to a first one of the respective pair of DDR DRAM devices in each of at least some of the multiple N-bit-wide ranks, while a second subset of the first tristate buffers is enabled for the first time period to drive a second nibble of the respective byte-wise section of the N-bit wide write data to a second subset of the respective module data lines coupled to a second one of the respective pair of DDR DRAM devices in each of at least some of the multiple N-bit-wide ranks.

3. The memory module of claim **2**, wherein the byte-wise data path further includes a set of write buffers configurable to receive the respective byte-wise section of the N-bit wide write data from the respective set of data signal lines before the first tristate buffers regenerate and drive the respective byte-wise section of the N-bit wide data to the second side of the each respective byte-wise buffer.

4. The memory module of claim **3**, wherein each of the write buffers is comparable to an input buffer on one of the DDR DRAM devices such that the each respective byte-wise buffer presents to the memory controller one DDR DRAM device load during the memory operation.

5. The memory module of claim **1**, wherein the DDR DRAM devices each has a bit width of 8 bits, and wherein the at least one respective DDR DRAM device in each of the multiple N-bit-wide ranks includes a single DDR DRAM device.

6. The memory module of claim **1**, wherein the logic is configurable to enable the first tristate buffers at a beginning of the first time period and to disable the first tristate buffers at an end of first time period.

7. The memory module of claim **1**, wherein the module controller is configured to use the module control signals to control timing of the first time period in accordance with the latency parameter.

8. The memory module of claim **1**, wherein the registered address and control signals include rank select signals, the rank select signals including one rank select signal for each of the multiple N-bit-wide ranks, and wherein the rank select signal received by the first N-bit-wide rank is different from the rank select signal received by any other N-bit-wide rank of the multiple N-bit-wide ranks.

9. The memory module of claim **8**, wherein each of the respective module data lines is configured to carry data from the memory controller to a corresponding memory device in each of the multiple N-bit-wide ranks during the memory write operation.

10. The memory module of claim **1**, wherein:

the module controller is further configurable to receive from the memory controller via the address and control signal lines additional input address and control signals for a memory read operation to read N-bit-wide read data from the memory controller from a second N-bit-wide rank of the multiple N-bit-wide ranks, and to output additional registered address and control signals in response to the additional input address and control signals;

the second N bit-wide rank is configurable to output the N bit-wide read data associated with the memory read

operation in response to at least some of the additional registered address and control signals;

the module controller is further configurable to output additional module control signals in response to the additional input address and control signals;

the logic in the each respective byte-wise buffer is further configurable to control the byte-wise data path in response to the additional module control signals, wherein the byte-wise data path is enabled for a second time period to actively drive a respective byte-wise section of the N bit wide read data associated with the memory read operation from the second side to the first side during the second time period in response to the additional module control signals;

the byte-wise data path further includes second tristate buffers configurable to be enabled by the logic to drive the respective byte-wise section of the N-bit wide read data to the respective set of data signal lines during the second time period; and

the second tristate buffers are disabled during the first time period and the first tristate buffers are disabled during the second time period.

11. A N-bit-wide memory module mountable in a memory socket of a computer system and configured to communicate with a memory controller of the computer system via address and control signal lines and N-bit wide data signal lines, the N-bit wide data signal lines including a plurality of sets of data signal lines, the memory module comprising:

a printed circuit board (PCB) having an edge connector comprising a plurality of electrical contacts which are positioned on an edge of the PCB and are configured to be releasably coupled to corresponding contacts of the memory socket;

double data rate dynamic random access memory (DDR DRAM) devices coupled to the PCB and arranged in multiple N-bit-wide ranks, each N-bit-wide rank includes n DDR DRAM devices;

n/2 data transmission circuits mounted on the PCB, wherein each of the n/2 data transmission circuits is disposed on the PCB at respective positions corresponding to a respective set of data signal lines among the plurality of sets of data signal lines, and wherein the each of the n/2 data transmission circuits has a first side configured to be operatively coupled to the respective set of data signal lines, a second side that is operatively coupled to a respective pair of DDR DRAM devices in each of the multiple ranks via respective module data lines, and data paths between the first side and the second side;

a module controller coupled to the PCB and operatively coupled to the DDR DRAM devices, the module controller being configurable to receive from the memory controller via the address and control signal lines input address and control signals for a memory write operation to write N-bit-wide write data from the memory controller into a first N-bit-wide rank among the multiple N-bit-wide ranks, and to output registered address and control signals in response to the input address and control signal, wherein the registered address and control signals cause the first N-bit-wide rank to perform the memory write operation by receiving the N-bit-wide write data, wherein the module controller is further configurable to transmit module control signals to the n/2 data transmission circuits in response to the input address and control signals; and

wherein, in response to the module control signals, each respective data transmission circuit is configurable to enable the data paths for a first time period in accordance with a latency parameter to actively drive a respective section of the N-bit wide write data associated with the memory operation from the first side to the second side during the first time period, wherein a first subsection of the respective section of the N-bit wide write data output from the second side of the each respective data transmission circuit is written into a first one of the respective pair of DDR DRAM devices in the first N-bit wide rank, while a second subsection of the respective section of the N-bit wide write data output from the second side of the each respective data transmission circuit is written into a second one of the respective pair of DDR DRAM devices in the first N-bit wide rank; and

wherein the data paths includes first tristate buffers, and the respective data transmission circuit in response to the module control signals is configured to enable a first subset of the first tristate buffers to drive the first subsection of the respective section of the N-bit wide write data to a first subset of the respective module data lines coupled to a first one of the respective pair of DDR DRAM devices in each of at least some of the multiple N-bit-wide ranks, and to enable a second subset of the first tristate buffers to drive a second subsection of the respective section of the N-bit wide write data to a second subset of the respective module data lines coupled to a second one of the respective pair of DDR DRAM devices in each of at least some of the multiple N-bit-wide ranks.

12. The memory module of claim **11**, wherein the registered address and control signals include a rank select signal for each of the multiple N-bit wide ranks, the rank select signal for the first N-bit wide rank is different from the rank select signal for any other N-bit wide rank of the multiple N-bit wide ranks.

13. The memory module of claim **12**, wherein each of the respective module data lines is configured to carry data from the memory controller to a corresponding memory device in each of the multiple N-bit-wide ranks during the memory operation.

14. The memory module of claim **13**, wherein the data paths further include a set of write buffers configurable to receive the respective section of the N-bit wide data from the respective set of data signal lines before the respective section of the N-bit wide write data is regenerated and driven by—the first tristate buffers to the respective module data lines.

15. The memory module of claim **14**, wherein each of the set of write buffers is comparable in loading to an input buffer on one of the DDR DRAM devices such that the each respective data transmission circuit presents to the memory controller one DDR DRAM device load during the memory write operation.

16. The memory module of claim **11**, wherein the module controller is configurable to use the module control signals to control timing of the first time period in accordance with the latency parameter.

17. The memory module of claim **11**, wherein the each respective data transmission circuit is configured to present to the memory controller one DDR DRAM device load on each of the respective set of data lines during the memory write operation.

18. The memory module of claim **11**, wherein:

the module controller is further configurable to receive from the memory controller via the address and control signal lines additional input address and control signals

for a subsequent memory read operation to read N bit-wick: read data from a second N-bit-wide rank of the multiple N-bit-wide ranks, and to output additional registered address and control signals in response to the additional input address and control signals;

the second N-bit-wide rank of the multiple N-bit-wide ranks is configurable to output the N-bit-wide read data associated with the memory read operation in response to at least some of the additional registered address and control signals;

the module controller is further configurable to output additional module control signals in response to the additional input address and control signals; and

the each respective data transmission circuit is further configurable to enable the data paths for a second time period to actively drive a respective section of the N-bit wide read data associated with the memory read operation from the second side to the first side during the second time period in response to the additional module control signals;

the data paths further includes second tristate buffers configurable to be enabled during the second time period to drive the respective section of the N-bit wide read data to the respective set of data signal lines; and

the each respective data transmission circuit is further configurable to keep the second tristate buffers disabled during the first time period and to keep the first tristate buffers disabled during the second time period.

19. A N-bit-wide memory module mountable in a memory socket of a computer system and configured to communicate with a memory controller of the computer system via address and control signal lines and N-bit wide data signal lines, the N-bit wide data signal lines including a plurality of sets of data signal lines, the memory module comprising:

a printed circuit board (PCB) having an edge connector comprising a plurality of electrical contacts which are positioned on an edge of the PCB and are configured to be releasably coupled to corresponding contacts of the memory socket;

double data rate dynamic random access memory (DDR DRAM) devices coupled to the PCB and arranged in multiple N-bit-wide ranks;

a module controller coupled to the PCB and operatively coupled to the DDR DRAM devices, the module controller being configurable to receive from the memory controller via the address and control signal lines first address and control signals for a first memory operation to read first N-bit-wide data from a first N-bit-wide rank among the multiple ranks and to subsequently receive second address and control signals for a second memory operation to read second N-bit-wide data from a second N-bit-wide rank among the multiple ranks, the module controller being further configurable to output first registered address and control signals for the first memory operation in response to receiving the first address and control signals and to output second registered address and control signals for the second memory operation in response to receiving the second address and control signals, wherein the first registered address and control signals cause the first N-bit-wide rank to output the first N-bit-wide data associated with the first memory operation, and the second registered address and control signals cause the second N-bit-wide rank to output the second N-bit-wide data associated with the second memory operation, wherein the module controller is further configurable to output first module control signals for the first memory operation in response to receiving the first address and control signals and to output second module control signals for the second memory operation in response to receiving the second address and control signals; and

a plurality of buffers coupled to the PCB and configured to receive the first module control signals and to receive the second module control signals, wherein each respective buffer of the plurality of buffers has a first side that is operatively coupled to a respective set of data signal lines, a second side that is operatively coupled to at least one respective DDR DRAM device in each of the multiple N-bit-wide ranks via respective module data lines, and data paths between the first side and the second-side, wherein the each respective buffer is disposed on the PCB at a respective position corresponding to the respective set of the plurality of sets of data signal lines;

wherein the each respective buffer further includes logic configurable to enable the data paths for a first time period to actively drive a respective section of the first N-bit wide data associated with the first memory operation from the second side to the first side during the first time period, and to enable the data paths for a second time period subsequent to the first time period to actively drive a respective section of the second N-bit wide data associated with the second memory operation from the second side to the first side during the second time period, wherein the data paths are disabled after the first time period and before the second time period.

20. The memory module of claim **19**, wherein the DDR DRAM devices each has a bit width of 4 bits, wherein the at least one respective DDR DRAM device in each of the multiple N-bit-wide ranks includes a respective pair of DDR DRAM devices, and wherein a first nibble of the respective section of the first N bit wide data is output by a first one of the respective pair of DDR DRAM devices in the first N-bit wide rank, while a second nibble of the respective section of the first N bit wide data is output by a second one of the respective pair of DDR DRAM devices in the first N-bit wide rank.

21. The memory module of claim **19**, wherein the module controller is configured to use the first module control signals to control timing of the first time period in accordance with a latency parameter and to use the second module control signals to control timing of the second time period in accordance with the latency parameter.

22. The memory module of claim **19**, wherein the each respective buffer is configured to present to the at least one respective DDR DRAM device a load that is similar to that of the memory controller during the first or the second memory operation.

23. The memory module of claim **19**, wherein the first registered address and control signals include a first set of rank select signals corresponding to respective ones of the multiple N-bit wide ranks, the first set of rank select signals including a first rank select signal received by the first N-bit wide rank that is different from each of the other rank select signals in the first set of rank select signals, and wherein the second registered address and control signals include a second set of rank select signals corresponding to respective ones of the multiple N-bit wide ranks, the second set of rank select signals including a second rank select signal received by the second N-bit wide rank that is different from each of the other rank select signals in the second set of rank select signals.

24. The memory module of claim **23**, wherein each of the respective module data lines is configured to carry data from

the memory controller to a corresponding memory device in each of the multiple N-bit-wide ranks.

25. The memory module of claim **19**, wherein the data path includes at least a first input buffer and a second buffer for each bit of the respective section of the first or second N-bit wide data, wherein the first buffer is configured to receive the each bit from the second side of the each respective buffer, and the second buffer is configured to regenerate and drive the each bit to the first side of the each respective buffer, wherein at least the second buffer is a tristate buffer configurable to be enabled by the logic in response to the first or second module control signals, and wherein the first buffer is comparable to an output buffer on one of the DDR DRAM devices.

26. The memory module of claim **19**, wherein the data paths include a set of input buffers configurable to receive the respective section of the first or second N-bit wide data from the respective set of module data lines and a set of output buffers configurable to regenerate and drive the respective section of the first or second N-bit wide data to the first side, wherein at least the set of output buffers are tristate buffers configurable to be enabled by the logic during the first time period in response to the first module control signals and during the second time period in response to the second module control signals, and wherein at least the set of output buffers are disabled by the logic after the first time period and before the second time period.

27. A memory module mountable in a memory socket of a computer system and configurable to communicate with a memory controller of the computer system via address and control signal lines and data signal lines, the data signal lines including a plurality of sets of data signal lines, each set of data signal lines is n bit wide, the memory module comprising:

a printed circuit board (PCB) having an edge connector comprising a plurality of electrical contacts which are positioned on an edge of the PCB and are configured to be releasably coupled to corresponding contacts of the memory socket;

a module controller coupled to the PCB and configurable to receive from the memory controller via the address and control signal lines first address and control signals for a memory write operation and to subsequently receive second address and control signals for a memory read operation, the module controller being further configurable to output first registered address and control signals and first module control signals for the memory write operation in response to the first address and control signals, and to output second registered address and control signals and second module control signals for the memory read operation in response to the second address and control signals; and

double data rate dynamic random access memory (DDR DRAM) devices coupled to the PCB and configurable to perform the memory write operation by receiving write data in response to the first registered address and control signals, and to perform the memory read operation by outputting read data in response to the second registered address and control signals; and

a plurality of n-bit-wide data buffers coupled to the PCB and configured to receive the first module control signals and subsequently the second module control signals, wherein each respective n-bit-wide data buffer of the plurality of n-bit-wide data buffers has a first side that is operatively coupled to a respective set of data signal lines, and a second side that is operatively coupled to a respective n-bit-wide section of the DDR DRAM devices via respective module data lines, wherein:

the each respective n-bit-wide data buffer includes a first set of input buffers configurable to receive a respective n-bit section of the write data from the respective set of data signal lines, a first set of tristate buffers configurable to drive the respective n-bit section of the write data to the respective module data lines, a second set of input buffers configurable to receive a respective n-bit section of the read data from the respective module data lines, a second set of tristate buffers configurable to drive the respective n-bit section of the read data to the respective set of data signal lines, and logic configurable to control at least the first set of tristate buffers and the second set of tristate buffers;

the logic in response to the first module control signals is configured to enable the first set of tristate buffers for a first time period corresponding to the memory write operation to drive the respective n-bit section of the write data;

the logic in response to the second module control signals is configured to enable the second set of tristate buffers for a second time period corresponding to the memory read operation to drive the respective n-bit section of the read data;

the first set of tristate buffers are disabled during the second time period; and

the second set of tristate buffers are disabled during the first time period.

28. The memory module of claim **27**, wherein the logic is configurable to keep the first set of tristate buffers and the second set of tristate buffers disabled when the memory module is not targeted by the memory controller for any memory operations.

29. The memory module of claim **27**, wherein:

the DDR DRAM devices include a plurality of ranks of DDR DRAM devices;

the each respective n-bit-wide buffer is coupled to at least one respective DDR DRAM device in each of the plurality of ranks via the respective module data lines;

the first registered address and control signals cause one rank of DDR DRAM devices to perform the memory write operation by receiving the write data; and

the second registered address and control signals cause one rank of DDR DRAM devices to perform the memory read operation by outputting the read data.

30. The memory module of claim **29**, wherein:

the at least one respective DDR DRAM device in each of the plurality of ranks includes a respective pair of DDR DRAM devices;

a first n/2-bit section of the respective n-bit section of the write_data is driven by a first subset of the first set of tristate buffers to a first one of the respective pair of DDR DRAM devices in the one rank of DDR DRAM devices performing the memory write operation;

a second n/2-bit section of the respective n-bit section of the write_data is driven by a second subset of the first set of tristate buffers to a second one of the respective pair of DDR DRAM devices in the one rank of DDR DRAM devices performing the memory write operation;

a first n/2-bit section of the respective n-bit section of the read_data is received by a first subset of the second set of input buffers from a first one of the respective pair

of DDR DRAM devices in the one rank of DDR DRAM devices performing the memory read operation; and

a second n/2-bit section of the respective n-bit section of the read_data is received by a second subset of the second set of input buffers from the second one of the respective pair of DDR DRAM devices in the one rank of DDR DRAM devices performing the memory read operation.

31. The memory module of claim **29**, wherein:

the at least one respective DDR DRAM device in each of the plurality of ranks includes one respective DDR DRAM;

the respective n-bit section of the write_data is driven by the first set of tristate buffers to the respective DDR DRAM device in the one rank of DDR DRAM devices performing the memory write operation; and

the respective n-bit section of the read_data is received by the second set of input buffers from the respective DDR DRAM device in the one rank of DDR DRAM devices performing the memory read operation.

32. The memory module of claim **29**, wherein each of the respective module data lines is configured to carry data from the memory controller to a corresponding memory device in each of the plurality of ranks.

33. The memory module of claim **29**, wherein:

the module controller is further configurable to receive from the memory controller via the address and control signal lines third address and control signals for a subsequent memory read operation, and to output third registered address and control signals and third module control signals for the subsequent memory read operation in response to receiving the third address and control signals;

the DDR DRAM devices are configurable to perform the subsequent memory read operation by outputting additional read data in response to the third registered address and control signals;

the subsequent memory read operation is performed by another rank of DDR DRAM devices that is different from the one rank of DDR DRAM devices performing the memory read operation;

the logic in response to the third module control signals enables the second set of tristate buffers for a third time period corresponding to the subsequent memory read operation to drive a respective n-bit section of the additional read data; and

the logic is further configurable to disable the second set of tristate buffers after the second time period and before the third time period.

34. The memory module of claim **27**, wherein the first set of tristate buffers are enabled for the first time period in accordance with a latency parameter.

35. The memory module of claim **27**, wherein the second set of tristate buffers are enabled for the second time period in accordance with a latency parameter.

* * * * *



(12) **United States Patent**
Chen et al.

(10) **Patent No.: US 11,016,918 B2**
(45) **Date of Patent: May 25, 2021**

(54) **FLASH-DRAM HYBRID MEMORY MODULE**

(71) Applicant: **Netlist, Inc.**, Irvine, CA (US)

(72) Inventors: **Chi-She Chen**, Walnut, CA (US); **Jeffrey C. Solomon**, Irvine, CA (US); **Scott H. Milton**, Irvine, CA (US); **Jayesh Bhakta**, Cerritos, CA (US)

(73) Assignee: **Netlist, Inc.**, Irvine, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **17/138,766**

(22) Filed: **Dec. 30, 2020**

(65) **Prior Publication Data**

US 2021/0124701 A1 Apr. 29, 2021

Related U.S. Application Data

(63) Continuation of application No. 15/934,416, filed on Mar. 23, 2018, which is a continuation of application
(Continued)

(51) **Int. Cl.**
G06F 13/36 (2006.01)
G06F 13/28 (2006.01)
(Continued)

(52) **U.S. Cl.**
CPC ***G06F 13/28*** (2013.01); ***G06F 1/185*** (2013.01); ***G06F 3/0613*** (2013.01);
(Continued)

(58) **Field of Classification Search**
CPC .. G06F 13/28; G06F 13/4027; G06F 13/1694; G06F 13/4223; G06F 12/0638;
(Continued)

(56) **References Cited**

U.S. PATENT DOCUMENTS

2,043,099 A 6/1936 Hanna
3,562,555 A 2/1971 Ahrons
(Continued)

FOREIGN PATENT DOCUMENTS

EP 2737383 A2 6/2014
KR 0130873 B1 4/1999
(Continued)

OTHER PUBLICATIONS

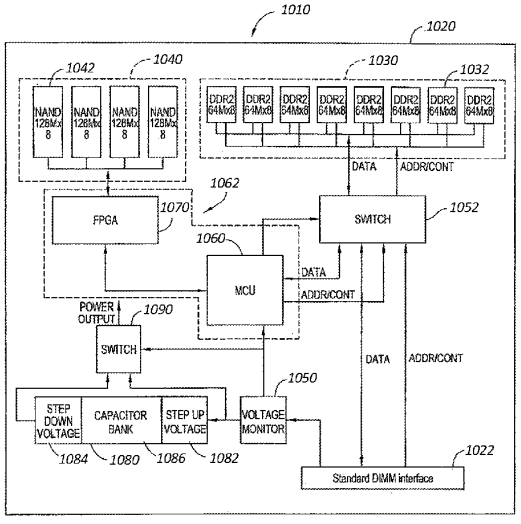
Exhibit, Decision Denying Institution of Inter Partes Review, *Smart Modular Technologies, Inc.* v. *Netlist, Inc.*, Case No. 4:13-cv-05889-YGR Document 309-3, filed Mar. 17, 2015, 23 pages.
(Continued)

Primary Examiner — Hashem Farrokh
(74) *Attorney, Agent, or Firm* — Shami Messinger PLLC

(57) **ABSTRACT**

In certain embodiments, a memory module includes a printed circuit board (PCB) having an interface that couples it to a host system for provision of power, data, address and control signals. First, second, and third buck converters receive a pre-regulated input voltage and produce first, second and third regulated voltages. A converter circuit reduces the pre-regulated input voltage to provide a fourth regulated voltage. Synchronous dynamic random access memory (SDRAM) devices are coupled to one or more regulated voltages of the first, second, third and fourth regulated voltages, and a voltage monitor circuit monitors an input voltage and produces a signal in response to the input voltage having a voltage amplitude that is greater than a threshold voltage.

30 Claims, 22 Drawing Sheets



Related U.S. Application Data

No. 14/840,865, filed on Aug. 31, 2015, now Pat. No. 9,928,186, which is a continuation of application No. 14/489,269, filed on Sep. 17, 2014, now Pat. No. 9,158,684, which is a continuation of application No. 13/559,476, filed on Jul. 26, 2012, now Pat. No. 8,874,831, which is a continuation-in-part of application No. 12/240,916, filed on Sep. 29, 2008, now Pat. No. 8,301,833, which is a continuation of application No. 12/131,873, filed on Jun. 2, 2008, now abandoned.

(60) Provisional application No. 60/941,586, filed on Jun. 1, 2007, provisional application No. 61/512,871, filed on Jul. 28, 2011.

(51) **Int. Cl.**

G06F 12/02	(2006.01)
G06F 13/16	(2006.01)
G06F 1/18	(2006.01)
G06F 12/06	(2006.01)
G06F 13/42	(2006.01)
G11C 7/10	(2006.01)
G11C 14/00	(2006.01)
G06F 3/06	(2006.01)
G06F 13/40	(2006.01)

(52) **U.S. Cl.**

CPC ***G06F 3/0659*** (2013.01); ***G06F 3/0685*** (2013.01); ***G06F 12/0246*** (2013.01); ***G06F 12/0638*** (2013.01); ***G06F 13/1694*** (2013.01); ***G06F 13/4027*** (2013.01); ***G06F 13/4243*** (2013.01); ***G11C 7/1072*** (2013.01); ***G11C 14/0018*** (2013.01); *G06F 2212/205* (2013.01); *G06F 2212/7208* (2013.01)

(58) **Field of Classification Search**

CPC G06F 12/0246; G06F 1/185; G06F 3/0613; G06F 3/0685; G06F 3/0659; G06F 2212/7208; G06F 2212/205; G11C 7/1072; G11C 14/0018

USPC 711/103, 104, 105; 710/138

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

3,916,390	A	10/1975	Chang et al.
4,234,920	A	11/1980	Ness et al.
4,420,821	A	12/1983	Hoffman
4,449,205	A	5/1984	Hoffman
4,607,332	A	8/1986	Goldberg
4,658,204	A	4/1987	Goodwin
4,882,709	A	11/1989	Wyland
4,884,242	A	11/1989	Lacy et al.
4,965,828	A	10/1990	Ergott, Jr. et al.
5,430,742	A	7/1995	Jeddeloh et al.
5,444,664	A	8/1995	Kuroda et al.
5,490,155	A	2/1996	Abdoo et al.
5,519,663	A	5/1996	Harper, Jr. et al.
5,519,831	A	5/1996	Holzhammer
5,563,839	A	10/1996	Herdt et al.
5,577,213	A	11/1996	Avery et al.
5,619,644	A	4/1997	Crockett et al.
5,630,096	A	5/1997	Zuravleff et al.
5,675,725	A	10/1997	Malcolm
5,757,712	A	5/1998	Nagel et al.
5,799,200	A	8/1998	Brant et al.
5,813,029	A	9/1998	Klein
5,870,350	A	2/1999	Bertin et al.
5,874,995	A	2/1999	Naimpally et al.
5,890,192	A	3/1999	Lee et al.
5,953,215	A	9/1999	Karabatsos
5,991,885	A	11/1999	Chang et al.
6,023,421	A	2/2000	Iadanza et al.
6,026,465	A	2/2000	Mills et al.
6,065,092	A	5/2000	Roy
6,112,310	A	8/2000	Jun et al.
6,145,068	A	11/2000	Lewis
6,158,015	A	12/2000	Klein
6,199,142	B1	3/2001	Saulsbury et al.
6,216,247	B1	4/2001	Creta et al.
6,269,382	B1	7/2001	Cabrera et al.
6,336,174	B1	1/2002	Li et al.
6,336,176	B1	1/2002	Leyda et al.
6,363,450	B1	3/2002	Lash et al.
6,421,279	B1	7/2002	Tobita et al.
6,459,647	B1	10/2002	Kengeri
6,487,102	B1	11/2002	Halbert et al.
6,487,623	B1	11/2002	Emerson et al.
6,571,244	B1	5/2003	Larson
6,614,685	B2	9/2003	Wong
6,658,507	B1	12/2003	Chan
6,691,209	B1	2/2004	O'connell
6,693,840	B2	2/2004	Shimada et al.
6,721,212	B2	4/2004	Sasaki
6,721,860	B2	4/2004	Klein
6,769,081	B1	7/2004	Parulkar
6,799,241	B2	9/2004	Kahn et al.
6,799,244	B2	9/2004	Tanaka et al.
6,810,513	B1	10/2004	Vest
6,816,982	B2	11/2004	Ravid
6,839,774	B1	1/2005	Ahn et al.
6,910,635	B1	6/2005	Miks et al.
6,944,042	B2	9/2005	Komatsuzaki
6,948,029	B2	9/2005	Yano
6,952,368	B2	10/2005	Miura et al.
7,053,470	B1	5/2006	Sellers et al.
7,062,618	B2	6/2006	Tsunoda et al.
7,089,412	B2	8/2006	Chen
7,102,391	B1	9/2006	Sun et al.
7,107,480	B1	9/2006	Moshayedi et al.
7,111,142	B2	9/2006	Spencer et al.
7,136,978	B2	11/2006	Miura et al.
7,155,627	B2	12/2006	Matsui
7,200,021	B2	4/2007	Raghuram
7,234,099	B2	6/2007	Gower et al.
7,353,325	B2	4/2008	Lofgren et al.
7,409,491	B2	8/2008	Doblar et al.
7,409,590	B2	8/2008	Moshayedi et al.
7,411,859	B2	8/2008	Sohn et al.
7,421,552	B2	9/2008	Long
7,467,251	B2	12/2008	Park et al.
7,519,754	B2	4/2009	Wang et al.
7,600,142	B2	10/2009	Ichikawa
7,613,877	B2	11/2009	Shimozono et al.
7,716,411	B2	5/2010	Panabaker et al.
7,818,488	B2	10/2010	Park et al.
7,873,750	B2	1/2011	Yabuta et al.
7,881,150	B2	2/2011	Solomon et al.
7,952,179	B2	5/2011	Chiu et al.
8,001,434	B1	8/2011	Lee et al.
8,081,536	B1	12/2011	Solomon et al.
8,086,955	B2	12/2011	Zhou et al.
8,102,614	B2	1/2012	Song et al.
8,214,616	B2	7/2012	Ware et al.
8,233,303	B2	7/2012	Best et al.
8,301,833	B1	10/2012	Chen et al.
8,407,395	B2	3/2013	Kim et al.
8,412,879	B2	4/2013	Chang et al.
8,516,187	B2	8/2013	Chen et al.
8,671,243	B2	3/2014	Chen et al.
8,677,060	B2	3/2014	Chen et al.
8,874,831	B2	10/2014	Lee et al.
8,880,791	B2	11/2014	Chen et al.
8,904,098	B2	12/2014	Amidi et al.
8,904,099	B2	12/2014	Chen et al.
9,043,677	B2	5/2015	Kong et al.
9,158,684	B2	10/2015	Lee et al.
9,361,250	B2	6/2016	Shan et al.

(56) **References Cited**

U.S. PATENT DOCUMENTS

9,436,600	B2	9/2016	Lee	
9,921,762	B2	3/2018	Amidi et al.	
9,928,186	B2	3/2018	Lee et al.	
2002/0053944	A1	5/2002	Brass et al.	
2002/0083368	A1	6/2002	Abe et al.	
2002/0199061	A1	12/2002	Friedman et al.	
2003/0028733	A1	2/2003	Tsunoda et al.	
2003/0076726	A1	4/2003	Cowles et al.	
2003/0137881	A1	7/2003	Sasaki	
2003/0147297	A1	8/2003	Shiota et al.	
2003/0158995	A1	8/2003	Lee et al.	
2003/0204776	A1	10/2003	Testin	
2003/0206478	A1	11/2003	Ayukawa et al.	
2003/0210601	A1	11/2003	Lin et al.	
2004/0088508	A1	5/2004	Ballard et al.	
2004/0163027	A1	8/2004	MacLaren et al.	
2004/0190210	A1	9/2004	Leete	
2005/0044302	A1	2/2005	Pauley et al.	
2005/0060488	A1	3/2005	Poechmueller	
2005/0132250	A1	6/2005	Hansen et al.	
2005/0141273	A1	6/2005	Park et al.	
2005/0144418	A1	6/2005	Kita	
2005/0183472	A1	8/2005	Choi	
2005/0204091	A1	9/2005	Kilbuck et al.	
2005/0249011	A1	11/2005	Maeda	
2005/0273548	A1	12/2005	Roohparvar	
2006/0039197	A1	2/2006	Khouri et al.	
2006/0069896	A1	3/2006	Sanders	
2006/0080515	A1	4/2006	Spiers et al.	
2006/0126369	A1	6/2006	Raghuram	
2006/0212651	A1	9/2006	Ashmore	
2006/0294295	A1	12/2006	Fukuzo	
2007/0070669	A1	3/2007	Tsern	
2007/0136523	A1*	6/2007	Bonella	G06F 9/4401 711/113
2007/0147115	A1	6/2007	Lin et al.	
2007/0192627	A1	8/2007	Oshikiri	
2007/0255898	A1	11/2007	Nishide et al.	
2007/0276995	A1	11/2007	Caulkins et al.	
2007/0288683	A1	12/2007	Panabaker et al.	
2008/0104344	A1	5/2008	Shimozono et al.	
2008/0126624	A1*	5/2008	Prete	G11C 7/106 710/53
2008/0126690	A1	5/2008	Rajan et al.	
2008/0147968	A1	6/2008	Lee et al.	
2008/0189479	A1	8/2008	Cope et al.	
2008/0195806	A1	8/2008	Cope	
2008/0235443	A1	9/2008	Chow et al.	
2008/0291727	A1	11/2008	Seo et al.	
2009/0031099	A1	1/2009	Sartore	
2009/0235038	A1	9/2009	Sartore	
2010/0110748	A1	5/2010	Best	
2010/0122200	A1	5/2010	Merry, Jr. et al.	
2010/0274953	A1	10/2010	Lee et al.	
2010/0322020	A1	12/2010	Kim	
2011/0078496	A1	3/2011	Jeddeloh	
2011/0161569	A1	6/2011	Shan et al.	
2011/0320804	A1	12/2011	Chan et al.	
2012/0110417	A1	5/2012	Abreu et al.	
2012/0117402	A1	5/2012	Machnicki et al.	
2012/0204079	A1	8/2012	Takefman et al.	
2012/0265952	A1	10/2012	Kurita	
2012/0271990	A1*	10/2012	Chen	G06F 12/00 711/103
2012/0317433	A1	12/2012	Ellis et al.	
2013/0019076	A1	1/2013	Amidi et al.	
2013/0086309	A1	4/2013	Lee et al.	
2013/0254456	A1	9/2013	Chen et al.	
2013/0254497	A1	9/2013	Chen et al.	
2014/0032820	A1	1/2014	Harasawa et al.	
2014/0059170	A1	2/2014	Gasparakis et al.	
2014/0156919	A1	6/2014	Chen et al.	
2014/0156920	A1	6/2014	Chen et al.	
2015/0058701	A1	2/2015	Xing et al.	

FOREIGN PATENT DOCUMENTS

KR	100606242	B1	7/2006
WO	2013016723	A2	1/2013

OTHER PUBLICATIONS

Exhibit, Decision Denying Institution of Inter Partes Review, *Smart Modular Technologies, Inc.* v. *Netlist, Inc.*, Case No. 4:13-cv-05889-YGR Document 309-4, filed Mar. 17, 2015, 20 pages.

Exhibit, Decision Denying Institution of Inter Partes Review, *Smart Modular Technologies, Inc.* v. *Netlist, Inc.*, Case No. 4:13-cv-05889-YGR Document 309-5, filed Mar. 17, 2015, 28 pages.

Exhibit, Decision Denying Institution of Inter Partes Review, *Smart Modular Technologies, Inc.* v. *Netlist, Inc.*, Case No. 4:13-cv-05889-YGR Document 309-9, filed Mar. 17, 2015,17 pages.

Exhibit, Institution of Inter Partes Review, *Sandisk Corporation* v. *Netlist, Inc.*, Case 4:13-cv-05889-YGR Document 316-9, filed Mar. 24, 2015, 29 pages.

Exhibit, Institution of Inter Partes Review, *Smart Modular Technologies, Inc.* v. *Netlist, Inc.*, Case 4:13-cv-05889-YGR Document 309-1, filed Mar. 17, 2015, 22 pages.

JEDEC Standard, Double Data Rate (DDR) SDRAM Specification, JESD79, Jun. 2000, 77 pages.

Exhibit, Letter sent via email on Dec. 6, 2013, Case 4:13-cv-03901-YGR Document 53-2, filed Jan. 6, 2014.

Exhibit, Order Denying Defendant's Motion to Stay Pending Inter Partes Review (Doc.59), *The Procter and Gamble Company* v. *Team Technologies, Inc, et al.*, Case 4:13-cv-05889-YGR Document 316-6, filed Mar. 24, 2015, 10 pages.

Exhibit, Patent Public Advisory Committee Quarterly Meeting, Appeals Statistics USPTO, Case No. 4:13-cv-05889-YGR Document 309-8, filed Mar. 17, 2015, 23 pages.

Exhibit, Reporter's Transcript of Proceedings, *Netlist, Inc* v. *Smart Modular Technologies, Inc., et al.*, Case 4:13-cv-05889-YGR Document 316-3, filed Mar. 24, 2015, 15 pages.

Exhibit, Transcript of Official Electronic Sound Recording Proceeding, *Netlist* v. *Smart Modular Technologies, Inc, et al.*, Case 4:13-cv-05889-YGR Document 305-7, filed Mar. 10, 2015, 10 pages.

File History for U.S. Appl. No. 12/240,916, filed Sep. 29, 2008, 320 pages.

File History for U.S. Appl. No. 13/905,048, filed May 29, 2013, 181 pages.

File History for U.S. Appl. No. 60/941,586, filed Jun. 1, 2007, 23 pages.

File History U.S. Pat. No. 8,671,243.

Final Office Action, dated Jun. 15, 2016, issued in U.S. Appl. No. 14/489,281, 10 pages.

Final Office Action, *SanDisk Corporation* v. *Netlist, Inc.*, U.S. Pat. No. 8,301,833, IPR2014-0099-1007 (PTAB), dated Feb. 1, 2012, 13 pages.

Final Written Decision, U.S. Pat. No. 8,671,243, IPR2017-00587-34, Paper No. 34, entered Jun. 20, 2018, 53 pages.

Final Written Decision, U.S. Pat. No. 8,874,831, IPR2017-00692, Paper No. 25, entered Jul. 5, 2018, 42 pages.

Final Written Decision, U.S. Pat. No. 7,881,150, Case No. Case IPR2014-00882, (PTAB), Paper 33, filed Dec. 14, 2015, 51 pages.

Hasan, J. et al. Efficient Use of Memory Bandwidth to Improve Network Processor Throughput, Proceedings of the 30th Annual International Symposium on Computer Architecture (ISCA'03), IEEE, 2003, 12 pages.

Inter Partes Review No. IPR2017-00692 (PTAB), U.S. Pat. No. 8,874,831, filed Jul. 26, 2012, 78 pages.

Inter Partes Review of U.S. Pat. No. 8,874,831, Case IPR2017-00692 (PTAB), filed Jul. 26, 2012, Paper No. 1, 78 pages.

International Preliminary Report on Patentabillity in PCT/US12/48750, dated Apr. 3, 2014pp. 1-8.

International Search Report and Written Opinion in PCT/US12/48750, dated Oct. 10, 2012pp. 1-10.

ISSCC 2006 / Session 7 / Non-Volatile Memory / 7.7, IEEE International Solid-State Circuits Conference, 2006, 10 pages.

(56) **References Cited**

OTHER PUBLICATIONS

Jacob, B., "Memory Systems Cache, DRAM, Disk", Morgan Kaufman Publishers, Burlington, MA, 2008, Preface and Ch. 7 pp. 315-322, 58 pages.
Jandhyala, S. et al., "Design-For-Test Analysis of a Buffered SDRAM DIMM", Semiconductor Group, Texas Instruments, Proceedings of International Workshop in Memory Technology, Design and Testing, Singapore, 13014, Aug. 1996,15 pages.
JEDEC Definition of DIMM, Exhibit 1029, IPR No. 2017-00587, Dec. 18, 2017, 2 pages.
JEDEC Global Standard for the Microelectronics Industry, Why JEDEC Standards Matter, 2014, 1 page.
JEDEC Standard No. 21-C (Release 17), Annex J: Serial Presence Detects for DDR2 SDRAM (Revision 1.3), 60 pages.
Petition for Inter Partes Review of U.S. Pat. No. 8,516,187 (on behalf of SMART Modular Technologies, Inc.), filed Aug. 22, 2014.
Petition for Inter Partes Review of U.S. Pat. No. 8,671,243, filed May 29, 2013.
Petition for Inter Partes Review of U.S. Pat. No. 8,874,831, filed Jul. 26, 2012.
Petition for Inter Partes Review of U.S. Pat. No. 8,301,833 (on behalf of SMART Modular Technologies, Inc.), filed Aug. 22, 2014. U.S. Appl. No. 60/912,321, Case No. No. IPR2017-00692-1007 (PTAB), filed Apr. 17, 2017, 42 pages.
Petition for Inter Partes Review of U.S. Pat. No. 8,516,187, IPR2014-00982 (PTAB), dated Jun. 19, 2014, 67 pages.
Petition for Inter Partes Review of U.S. Pat. No. 8,671,243, Case No. IPR2017-00587-1(PTAB), May 29, 2013, 82 pages.
Petition for Inter Partes Review, *SanDisk Corporation* v. *Netlist, Inc*., U.S. Pat. No. 8,301,833, IPR2014-00994 (PTAB), filed Jun. 20, 2014, 69 pages.
Webster's II New College Dictionary, Houghton Mifflin Company, Boston, MA, 2001, pp. 259, 1115.
Videotaped Deposition of Russel Jacob Baker, U.S. Pat. No. 8,671,243, Case Nos. IPR2017-00587-1030 (PTAB), filed Dec. 18, 2017, 268 pages.
Petitioners Demonstratives, *SK hynix Inc., et al*., v. *Netlist, Inc*., U.S. Pat. No. 8,671,243, IPR2017-00587-1037, 89 pages.
Petitioners' Reply in Support of its Motion to Exclude, U.S. Pat. No. 8,874,831, Case No. IPR2017-00692, filed Apr. 9, 2018, 8 pages.
Petitioners' Reply to Patent Owner's Response, Inter Partes Review No. IPR2017-00587-16, U.S. Pat. No. 8,671,243, filed Jan. 12, 2018, 38 pages.
Petitioners' Reply, U.S. Pat. No. 8,874,831, Case No. IPR2017-00692 (PTAB), Paper No. 15, filed Mar. 2, 2018, 35 pages.
Petitioners' Request for Hearing, U.S. Pat. No. 8,301,833, Case No. Case No. IPR2017-00649-2020 (PTAB), Paper No. 8, entered Aug. 23, 2017, 18 pages.
Prosecution History, U.S. Appl. No. 12/240,916, U.S. Pat. No. 8,301,833, *SK hynix Inc* v *Netlist Inc*, Case No. IPR2017-00649-1002 (PTAB), 320 pages.
Prosecution History, U.S. Appl. No. 13/559,476, U.S. Pat. No. 8,301,833, *SK hynix Inc* v *Netlist Inc*, Case No. IPR2017-00649-1002 (PTAB), 312 pages.
Prosecution History, U.S. Appl. No. 60/941,586, *SK hynix Inc* v *Netlist Inc*, Case No. IPR2017-00649-1005 (PTAB), 23 pages.
Provisional Application for Advance Dynamic Disk Memory Module, Specification, *SK hynix Inc*., V. *Netlist, Inc*., Case No. IPR2017-00649-1006 (PTAB), dated Dec. 8, 2005, 53 pages.
U.S. Appl. No. 60/912,321, filed Apr. 17, 2007.
U.S. Appl. No. 60/941,586, filed Jun. 1, 2007.
Requirement for Restriction Election, *SanDisk Corporation* v. *Netlist, Inc*., U.S. Pat. No. 8,301,833, IPR2014-00994-1003 (PTAB), dated Mar. 31, 2019, 7 pages.
Restriction Requirement in U.S. Appl. No. 12/240,916, dated Mar. 31, 2011.
Search Report (Updated) Prior Art Search for U.S. Pat. No. 8,301,833, Global Patent Solutions, Nov. 24, 2020, 179 pages.
Search Report Prior Art Search for U.S. Pat. No. 8,301,833, Global Patent Solutions, Nov. 2, 2020, 38 pages.
Second Amended Answer and Counterclaims to Plaintiff's Complaint for Declaratory Judgment, *Diablo Technologies, Inc*. v. *Netlist, Inc*., Case No. 4:13-CV-03901 (NDCA), YGR, filed Feb. 17, 2014, 21 pages.
Supplemental Declaration of Ronald H. Spuhler, *Netlist* v. *Smart Storage Systems, Inc. et. al*., Case 4:13-cv-05889-YGR Document 305-1, filed Mar. 10, 2015, 2 pages.
U.S. Office Action in U.S. Appl. No. 13/536,173, dated Apr. 15, 2013, pp. 1-10.
U.S. Appl. No. 12/240,916, Case No. IPR2017-00692-2018 (PTAB), dated Sep. 29, 2008, 52 pages.
Third Amended Complaint for Patent Infringement, *Netlist* v. *Smart Storage Systems, Inc. et. al*., Case 4:13-CV-05889-YGR (NDCA), filed Oct. 7, 2014, 21 pages.
Drawing by R. Jacob Bake regarding Address of DRAM and Flash, U.S. Pat. No. 8,671,243, Case IPR2017-00587-1023 (PTAB), filed Dec. 18, 2017, 1 page.
Advisory Action in U.S. Appl. No. 12/240,916, dated Mar. 13, 2012.
Amendment and Reply to Office Action, *SanDisk Corporation* v. *Netlist*, U.S. Appl. No. 13/536,173, IPR2014-00982 (PTAB), dated May 21, 2013, 24 pages.
Amendment and Response to Election Restriction, *SanDisk Corporation* v. *Netlist, Inc*., U.S. Pat. No. 8,301,833, IPR2014-00994-1004 (PTAB), dated May 20, 2011, 9 pages.
American Heritage Dictionary of the English Language, Third Ed., Houghton Mifflin Company, Boston, MA, 1996, 7 pages.
Appeals from the USPTO, PTAB in Nos. IPR2014-00882, IPR2014-00883, IPR2014-01011, US Court of Appeals for the Federal Circuit, decided Jul. 25, 2017, 8 pages.
Charles C., Multiplexers and Demultiplexers, Ex. 1024, IPR No. 2017-00587-1024(PTAB), 2003, 8 pages.
Charles C., Multiplexers and Demultiplexers, Ex. 1033, IPR No. 2017-00587-1033 (PTAB), 2003, 8 pages.
Charles C., Multiplexers and Demultiplexers, Ex. 1034, IPR No. 2017-00587-1033 (PTAB), 2003, 8 pages.
Charles C., Multiplexers and Demultiplexers, Ex. 1035, IPR No. 2017-00587-1035 (PTAB), 2003, 8 pages.
Corrected Petition for Inter Partes Review of Claims 1-30 of U.S. Pat. No. 8,301,833, Declaration of Michael F. Heafey, IPR2014-01370 (PTAB), filed Sep. 22, 2014, 4 pages.
Corrected Petition for Inter Partes Review of Claims 1-32 of U.S. Pat. No. 8,516,187, IPR2014-01371 (PTAB), Sep. 22, 2014, 67 pages.
Corrected Petition for Inter Partes Review of Claims 1-32 of U.S. Pat. No. 8,516,187, Declaration of Dr. Nader Bagherzadeh, IPR2014-01371 (PTAB), filed Sep. 22, 2014, 306 pages.
Decision Denying Institution of Inter Partes Review, U.S. Pat. No. 8,516,187, IPR2014-01371 (PTAB), Paper 12, entered Mar. 13, 2015, 22 pages.
Decision Denying Institution of Inter Partes Review, *SanDisk Corporation* v. *Netlist, Inc*., U.S. Pat. No. 8,301,833, IPR2014-00994 (PTAB), Paper 8, dated Dec. 16, 2014, 16 pages.
Decision Denying Institution of Inter Partes Review, *SanDisk Corporation* v. *Netlist*, U.S. Pat. No. 8,516,187, IPR2014-00982 (PTAB), Paper 9, dated Dec. 22, 2014, 16 pages.
Decision Denying Institution of Inter Partes Review, *Smart Modular Tech* v *Netlist Inc*., U.S. Pat. No. 8,301,833, IPR2014-01370 (PTAB), Paper 13, entered Mar. 13, 2015, 19 pages.
Decision Instituting Inter Partes Review, U.S. Pat. No. 8,671,243, Case No. IPR2017-00587-7(PTAB), Paper No. 7, entered Jun. 22, 2017, 40 pages.
Declaration of R. Jacob Baker, U.S. Pat. No. 8,874,831, Case IPR2017-00692-2016 (PTAB), filed Nov. 10, 2017, 72 pages.
Declaration of R. Jacob Baker, U.S. Pat. No. 8,671,243, Case IPR2017-00587-2016 (PTAB), filed Oct. 13, 2019, 109 pages.
Declaration of Jeff McMullen, *Netlist, Inc*. v. *Diablo Technologies, Inc*., Case No. 4: I 3-CV-05889-YGR (NDCA), Document 362-1, filed Sep. 5, 2018, 3 pages.
Declaration of Paul Min, In Inter Partes Review of U.S. Pat. No. 8,301,833, IPR2014-00994-1020 (PTAB), filed Sep. 29, 2008, 215 pages.

(56) **References Cited**

OTHER PUBLICATIONS

Declaration of Paul Min, In Inter Partes Review of U.S. Pat. No. 8,516,187, IPR2014-00982-1013 (PTAB), filed Jun. 28, 2012, pp. 240.
Declaration of Ron Maltiel, U.S. Pat. No. 8,671,243, No. IPR2017-00587-1003 (PTAB), filed May 29, 2013, 131 pages.
Declaration of Ron Maltiel, U.S. Pat. No. 8,874,831, No. IPR2017-00692-1003 (PTAB), filed Jul. 26, 2012 29, 172 pages.
Declaration of Steven J. Corr, *Netlist, Inc.* v. *SMART Storage Systems, Inc., et al*., Case 4:13-cv-05889-YGR (NDCA), Document 305-8, filed Mar. 10, 2015, 3 pages.
Deposition of Baker, Exhibit 1030, p. 78, U.S. Pat. No. 8,671,243, *SK hynix Inc., et al.* v. *Netlist, Inc*., Case No. IPR-2017-00587 (PTAB), filed Dec. 18, 2017, 268 pages.
Deposition of Ron Maltiel, U.S. Pat. No. 8,874,831, Case No. IPR2017-00692-2010 (PTAB), dated Sep. 27, 2017, 155 pages.
Joint Status Report Regarding Inter Partes Review, *Netlist* v. *Smart Storage Systems, Inc. et. al*., Case No. 4:13-CV-05889-YGR (NDCA), filed Aug. 17, 2018, 4 pages.
Material Science & Engineering, Department of Material Science and Engineering, Stanford University, Exhibit 2011, Case No. IPR2017-00692, 2 pages.
Merriam-Webster's Collegiate Dictionary, Eleventh Ed., Merriam Webster Corporation, Springfield, MA, 2003, 7 pages.
Microsoft Computer Dictionary Fifth Edition, 2002, 3 pages.
Microsoft Computer Dictionary Fifth Edition, 2002, 9 pages.
Microsoft Press, Computer Dictionary, Second Edition, 1994, 4 pages.
Microsoft Windows 2000 Professional Resource Kit, 76 pages.
Notice of Allowance in U.S. Appl. No. 12/240,916, dated Sep. 17, 2012.
Notice of Allowance in U.S. Appl. No. 13/536,173, dated Jul. 2, 2013.
Notice of Allowance in U.S. Appl. No. 13/559,476, dated May 6, 2014.
Notice of Allowance in U.S. Appl. No. 13/559,476, dated Sep. 29, 2014.
Notice of Allowance in U.S. Appl. No. 13/905,048, dated Dec. 19, 2013, 8 pages.
Notice of Allowance in U.S. Appl. No. 13/905,053, dated Dec. 11, 2013.
Notice of Allowance in U.S. Appl. No. 14/173,219 dated Jul. 7, 2014.
Notice of Allowance in U.S. Appl. No. 14/489,269, dated Oct. 8, 2015.
Notice of Allowance, *SanDisk Corporation* v. *Netlist*, U.S. Appl. No. 13/536,173, IPR2014-00982 (PTAB), dated Jul. 2, 2013, 8 pages.
Office Action in U.S. Appl. No. 12/240,916, dated Apr. 3, 2012.
Office Action in U.S. Appl. No. 13/536,176, dated Apr. 15, 2013.
Office Action in U.S. Appl. No. 13/625,563, dated Aug. 5, 2013.
Office Action in U.S. Appl. No. 13/625,563, dated May 9, 2014.
Office Action in U.S. Appl. No. 13/905,048, dated Aug. 1, 2013.
Office Action in U.S. Appl. No. 13/905,053, dated Aug. 1, 2013.
Office Action in U.S. Appl. No. 14/173,219, dated Mar. 13, 2014.
Office Action in U.S. Appl. No. 14/302,292, dated Dec. 21, 2015.
Office Action in U.S. Appl. No. 12/240,916, dated Feb. 1, 2012, 14 pages.
Office Action in U.S. Appl. No. 14/173,242, dated Mar. 14, 2014, 7 pages.
Office Action, U.S. Appl. No. 12/240,916, IPR2014-0099-1005 (PTAB), dated Jul. 29, 2011, 8 pages.
Office Action, U.S. Appl. No. 12/240,916, IPR2014-0099-1010 (PTAB), dated Apr. 3, 2012, 11 pages.
Office Action, U.S. Appl. No. 13/536,173, IPR2014-00982 (PTAB), dated Apr. 15, 2013, 9 pages.
Patent Owner's Demonstratives, U.S. Pat. No. 8,671,243, Case No. IPR 2017-00587-2023, (PTAB), 57 pages.
Patent Owner's Listing of New Arguments and Evidence in Petitioners' Reply, U.S. Pat. No. 8,671,243, Case No. IPR2017-00587 (PTAB), filed Jan. 29, 2018, 6 pages.
Patent Owner's Opposition to Petitioners' Motion to Exclude, U.S. Pat. No. 8,874,831, Case No. IPR2017-00692 (PTAB), filed Apr. 2, 2018, 2017, 11 pages.
Patent Owner's Preliminary Response, U.S. Pat. No. 8,516,187, Case IPR2014-01371 (PTAB), filed Dec. 16, 2014, 66 pages.
Patent Owner's Preliminary Response, *SanDisk Corporation* v. *Netlist, Inc*., U.S. Pat. No. 8,301,833, IPR2014-00994 (PTAB), Paper 8, dated Oct. 2, 2014, 60 pages.
Patent Owner's Preliminary Response, *SanDisk Corporation* v. *Netlist*, U.S. Pat. No. 8,516,187, IPR2014-00982 (PTAB), dated Sep. 26, 2014, 57 pages.
Patent Owner's Preliminary Response, *SK hynix Inc., et al*., v. *Netlist*, U.S. Pat. No. 8,301,833, Case No. IPR2017-00649 (PTAB), filed May 1, 2017, 67 pages.
Patent Owner's Preliminary Response, *SK hynix Inc., et al*., v. *Netlist*, U.S. Pat. No. 8,874,831, Case No. IPR2017-00692 (PTAB), filed May 1, 2017, 48 pages.
Patent Owner's Preliminary Response, *Smart Modular Tech* v *Netlist Inc*., U.S. Pat. No. 8,301,833, Case IPR2014-01370 (PTAB), filed Dec. 16, 2014, 66 pages.
Patent Owner's Response, U.S. Pat. No. 8,671,243, Case No. IPR2017-00587-12 (PTAB), filed Oct. 13, 2017, 80 pages.
Patent Owner's Response, *SK hynix Inc., et al*., v. *Netlist*, U.S. Pat. No. 8,874,831, Case No. IPR2017-00692-12 (PTAB), filed Nov. 10, 2017, 77 pages.
Petition for Inter Partes Review of Claims 1-30 of U.S. Pat. No. 8,301,833, IPR2014-01370 (PTAB), filed Aug. 22, 2014, 68 pages.
Petition for Inter Partes Review of U.S. Pat. No. 8,301,833 (on behalf of SanDisk, Corp.), filed Jun. 20, 2014.
Petition for Inter Partes Review of U.S. Pat. No. 8,301,833, filed Sep. 29, 2008.
Petition for Inter Partes Review of U.S. Pat. No. 8,516,187 (on behalf of SanDisk, Corp.), filed Jun. 19, 2014.

* cited by examiner

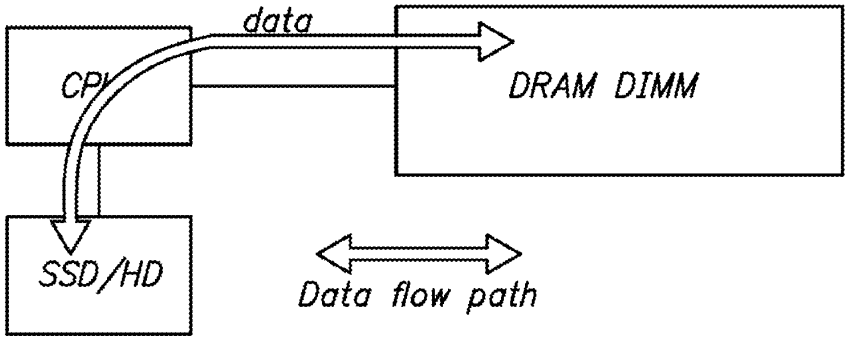


FIG. 1
(PRIOR ART)

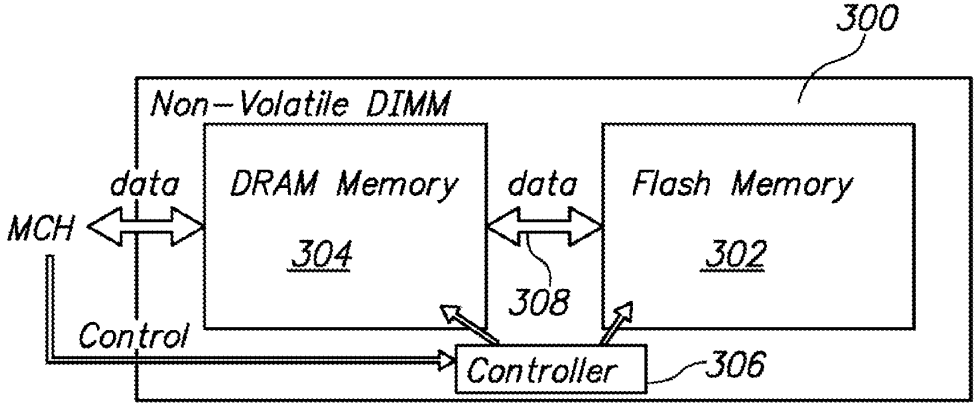


FIG. 3A

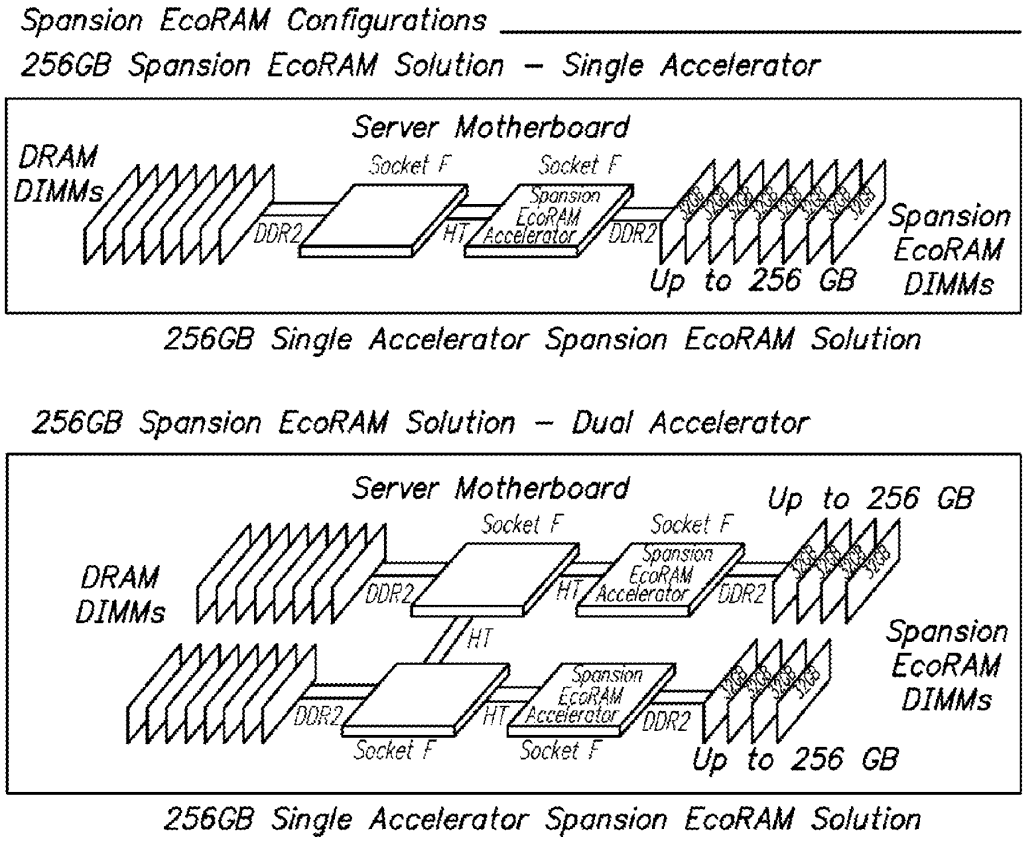


FIG. 2
(PRIOR ART)

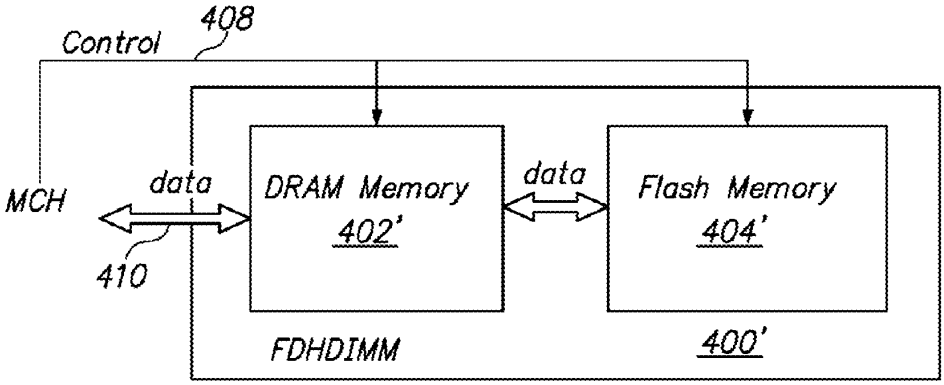


FIG. 4B

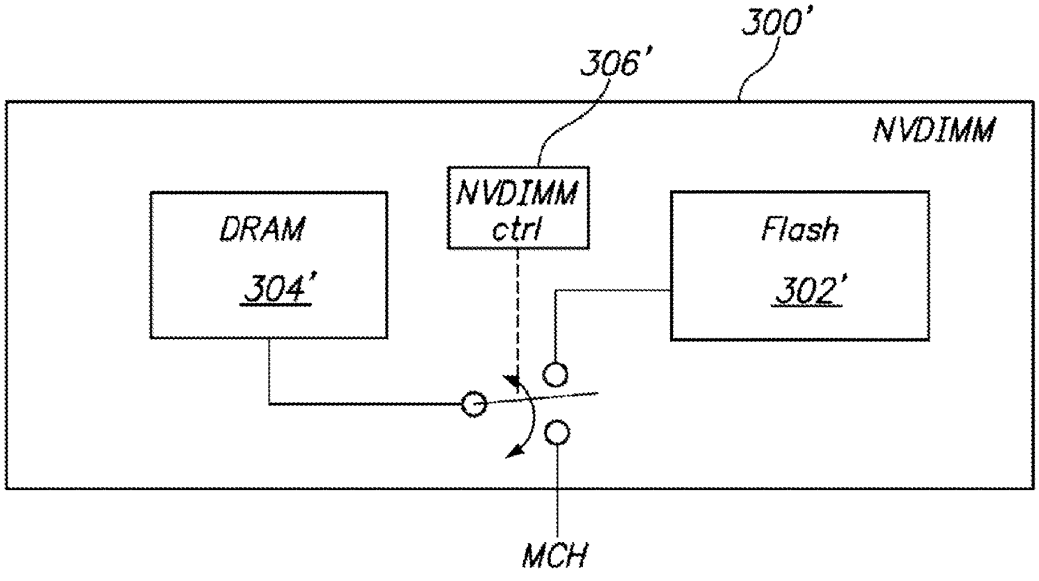

300'
306'
NVDIMM
NVDIMM ctrl
DRAM
304'
Flash
302'
MCH

FIG. 3B

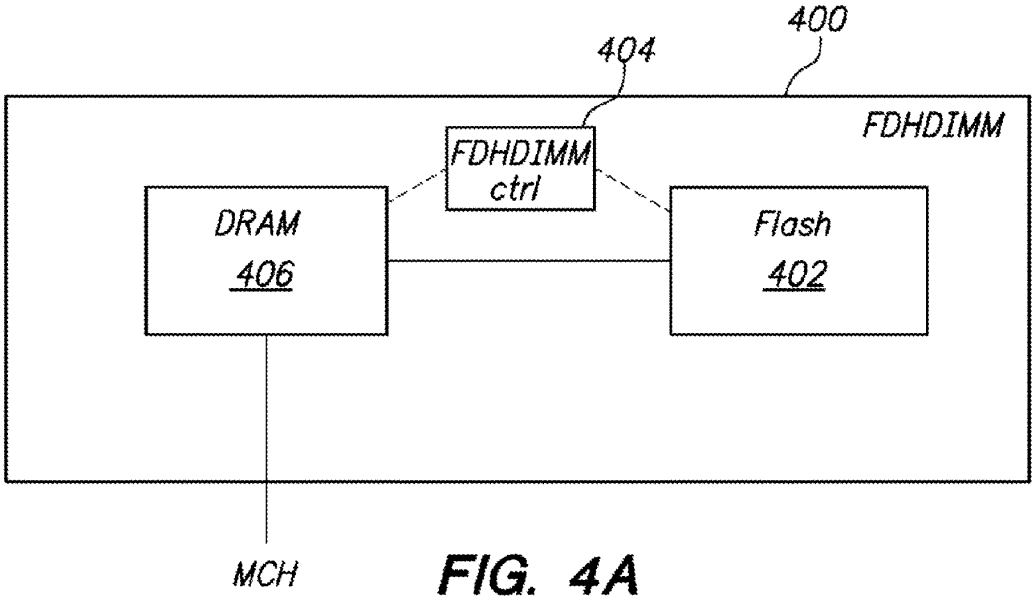

400
404
FDHDIMM
FDHDIMM ctrl
DRAM
406
Flash
402
MCH

FIG. 4A

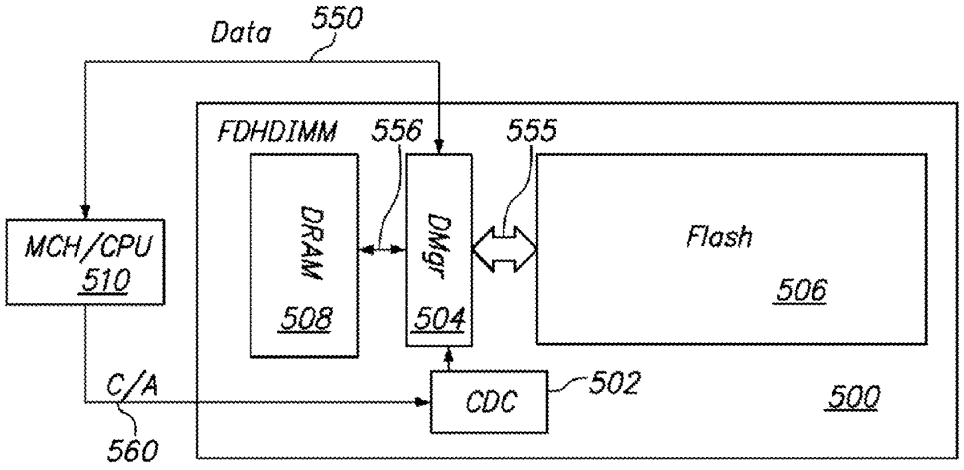

Data
550
FDHDIMM
556
555
DRAM
508
DMgr
504
Flash
506
MCH/CPU
510
C/A
560
CDC
502
500

FIG. 5A

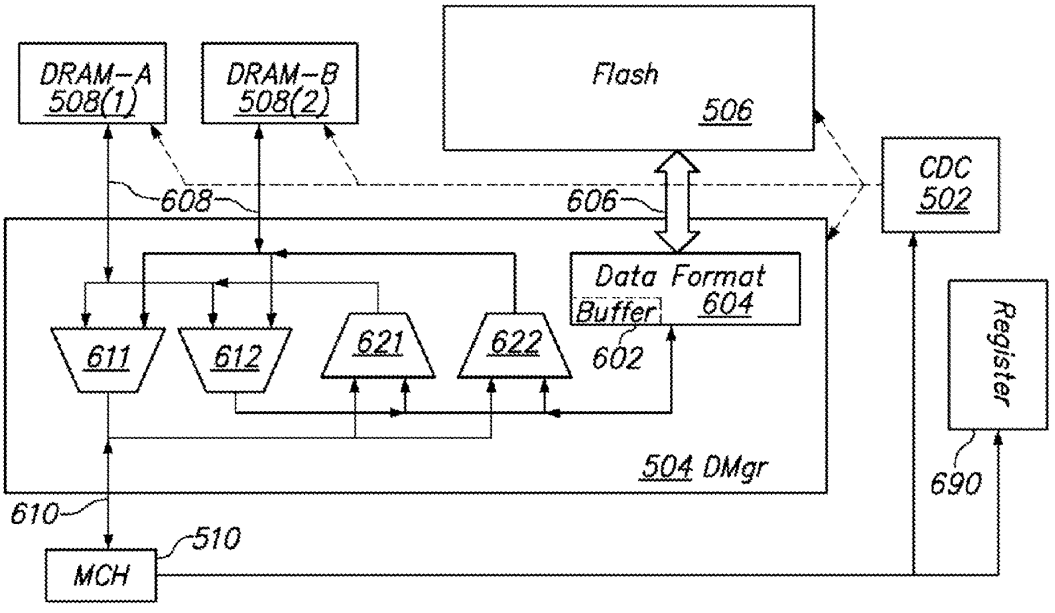

DRAM-A
508(1)
DRAM-B
508(2)
Flash
506
CDC
502
608
606
Data Format
Buffer
604
611
612
621
622
602
Register
504 DMgr
690
610
510
MCH

FIG. 6

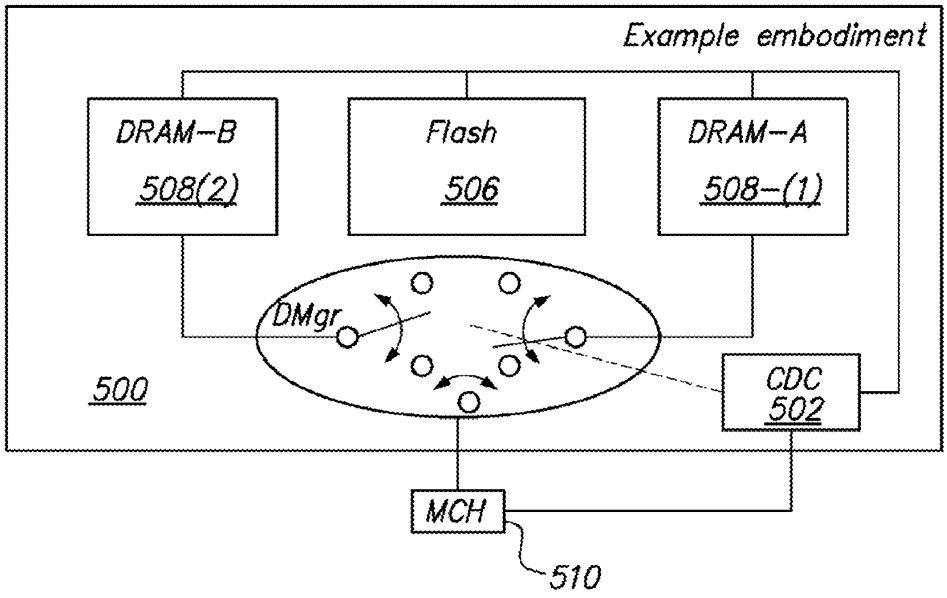

Example embodiment
DRAM–B
508(2)
Flash
506
DRAM–A
508–(1)
DMgr
500
CDC
502
MCH
510

FIG. 5B

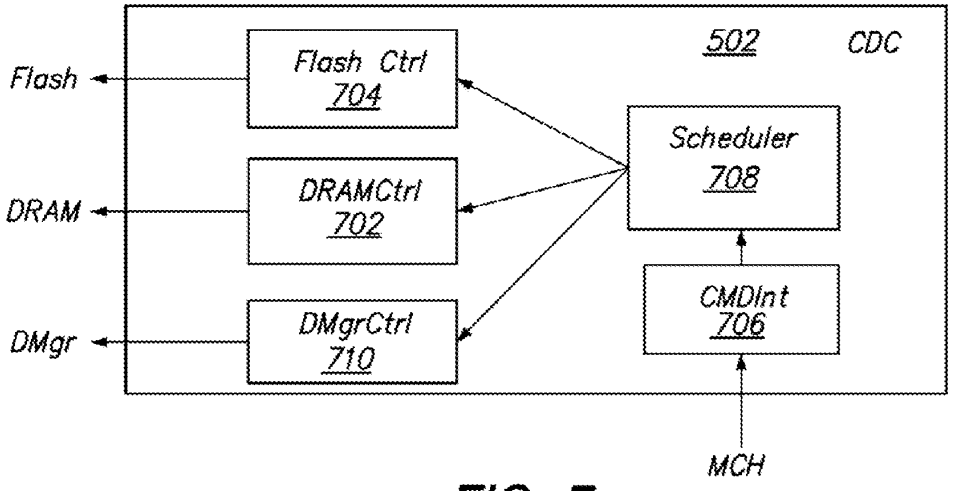

Flash
DRAM
DMgr
Flash Ctrl
704
DRAMCtrl
702
DMgrCtrl
710
502
CDC
Scheduler
708
CMDInt
706
MCH

FIG. 7

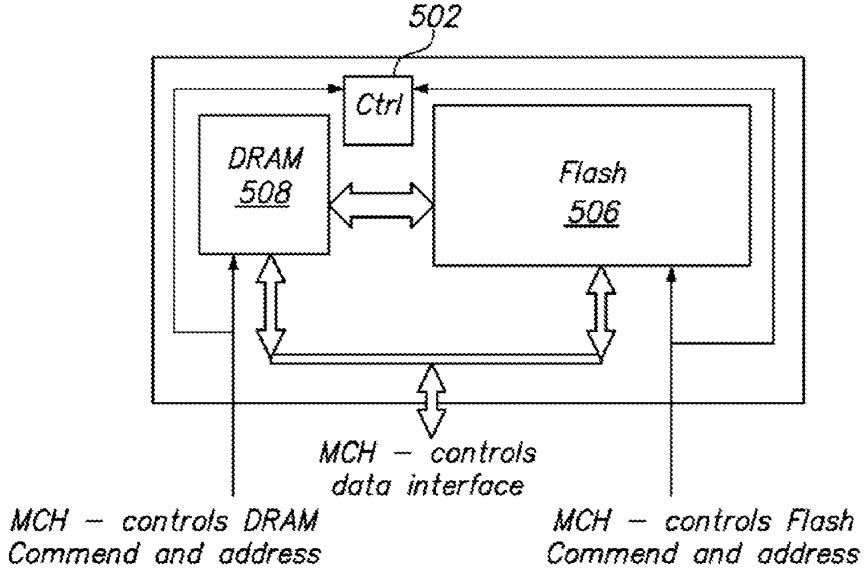

502
Ctrl
DRAM
508
Flash
506
MCH – controls
data interface
MCH – controls DRAM
Commend and address
MCH – controls Flash
Commend and address

FIG. 8A

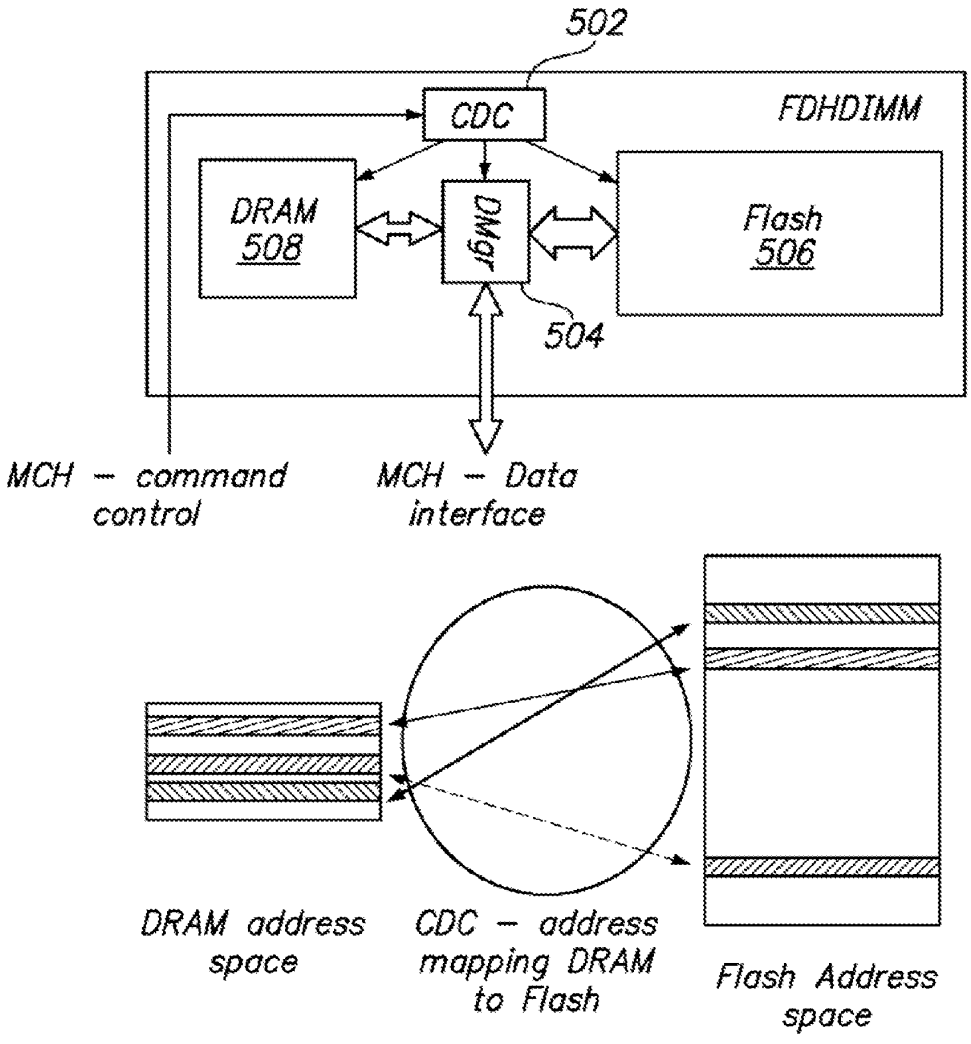

502
CDC
FDHDIMM
DRAM
508
DMgr
Flash
506
504
MCH – command control
MCH – Data interface
DRAM address space
CDC – address mapping DRAM to Flash
Flash Address space

FIG. 8B

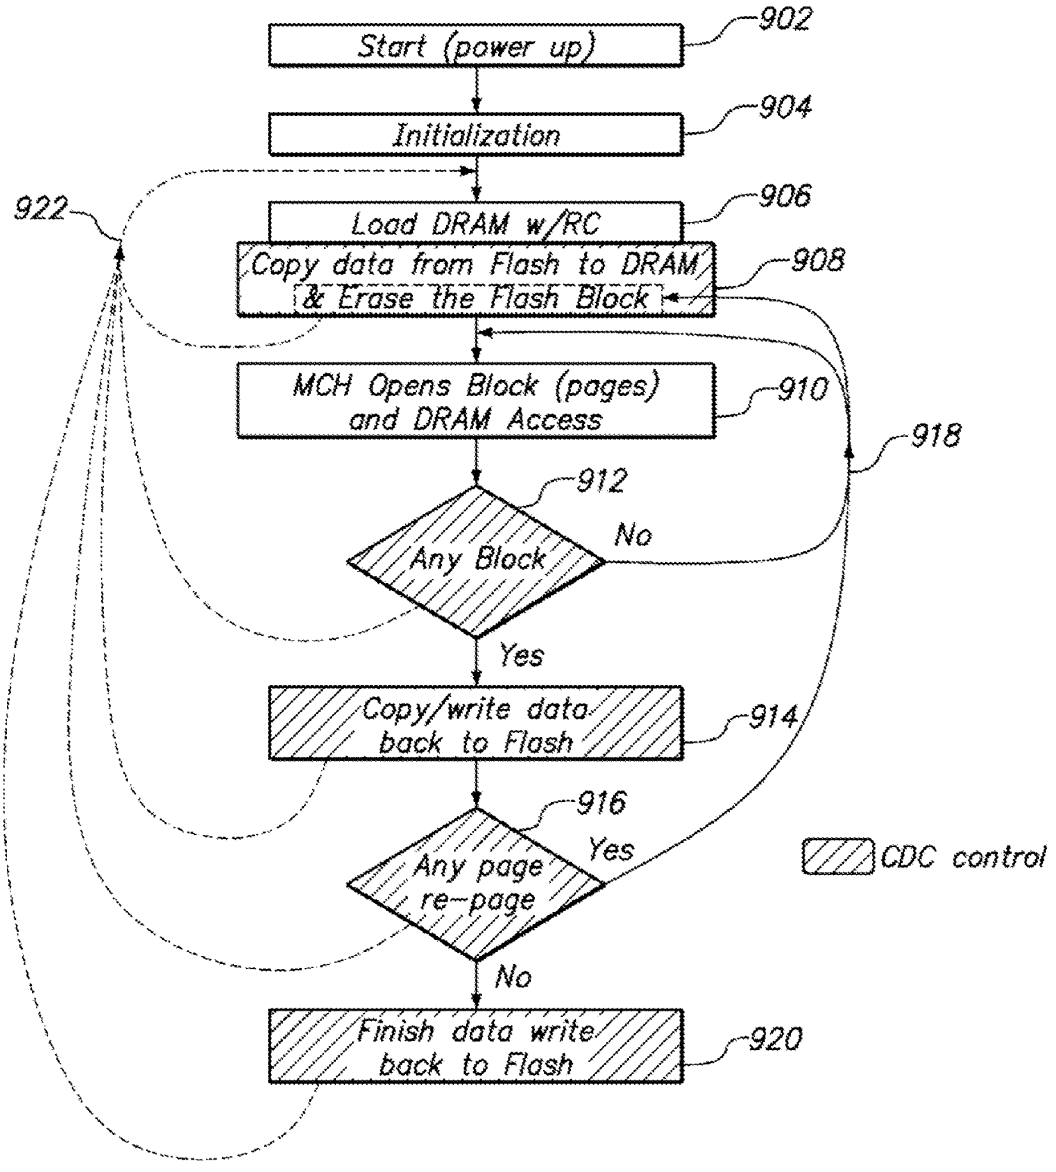

Start (power up)
902
Initialization
904
Load DRAM w/RC
906
Copy data from Flash to DRAM & Erase the Flash Block
908
MCH Opens Block (pages) and DRAM Access
910
Any Block
912
No
Yes
Copy/write data back to Flash
914
Any page re-page
916
Yes
No
Finish data write back to Flash
920
922
918
CDC control

FIG. 9

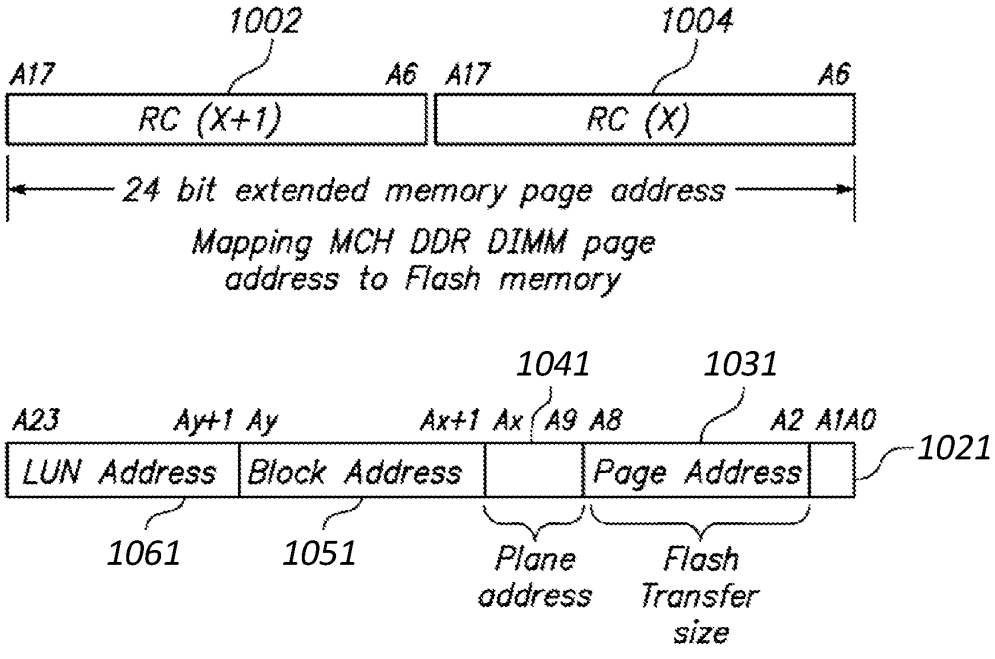

1002
1004
A17
A6
A17
A6
RC (X+1)
RC (X)
24 bit extended memory page address
Mapping MCH DDR DIMM page address to Flash memory
1041
1031
A23
Ay+1
Ay
Ax+1
Ax
A9
A8
A2
A1A0
1021
LUN Address
Block Address
Page Address
1061
1051
Plane address
Flash Transfer size

FIG. 10

DRAM density (GB)	# of blocks per bank	Flash wr-time to rd-time ratio	Avg block use time (sec)	Flash write time (sec)	Max allowed Closed Blk in queue to be written back to Flash
1	250	55	1.00E-03	2.00E-02	0
1	250	55	1.00E-02	2.00E-02	2
1	250	55	2.00E-02	2.00E-02	5
1	250	55	5.00E-02	2.00E-02	11
2	500	55	1.00E-03	2.00E-02	0
2	500	55	1.00E-02	2.00E-02	5
2	500	55	2.00E-02	2.00E-02	9
2	500	55	5.00E-02	2.00E-02	23
4	1000	55	1.00E-03	2.00E-02	1
4	1000	55	1.00E-02	2.00E-02	9
4	1000	55	2.00E-02	2.00E-02	18
4	1000	55	5.00E-02	2.00E-02	45

FIG. 11

Case: 24-2203 Document: 16 Page: 703 Filed: 12/16/2024

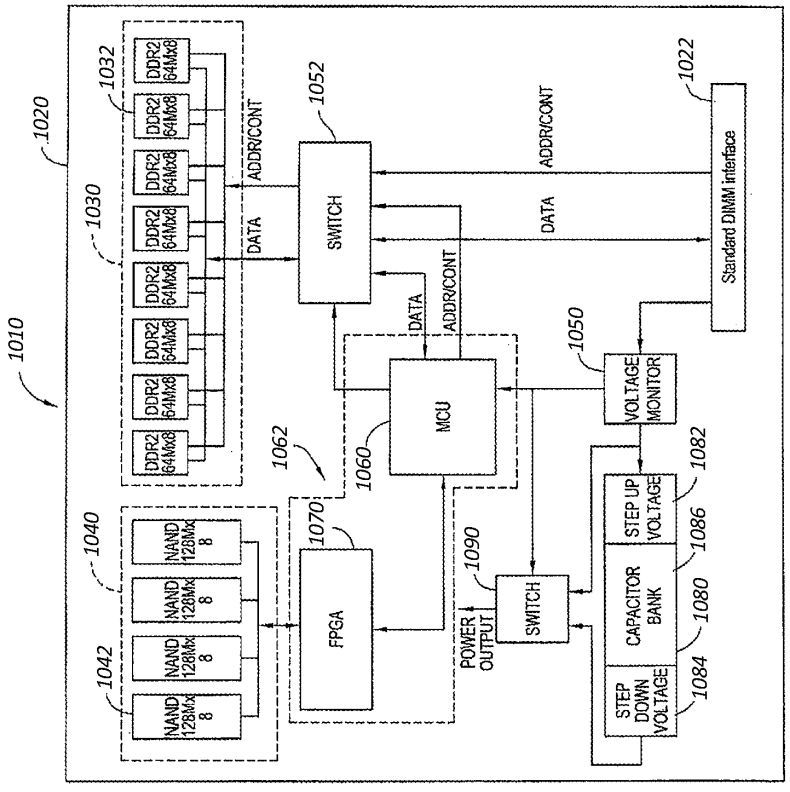


FIG. 12

JTX0003-0016
Appx615

NETLIST_SAMSUNG_EDTX00038091

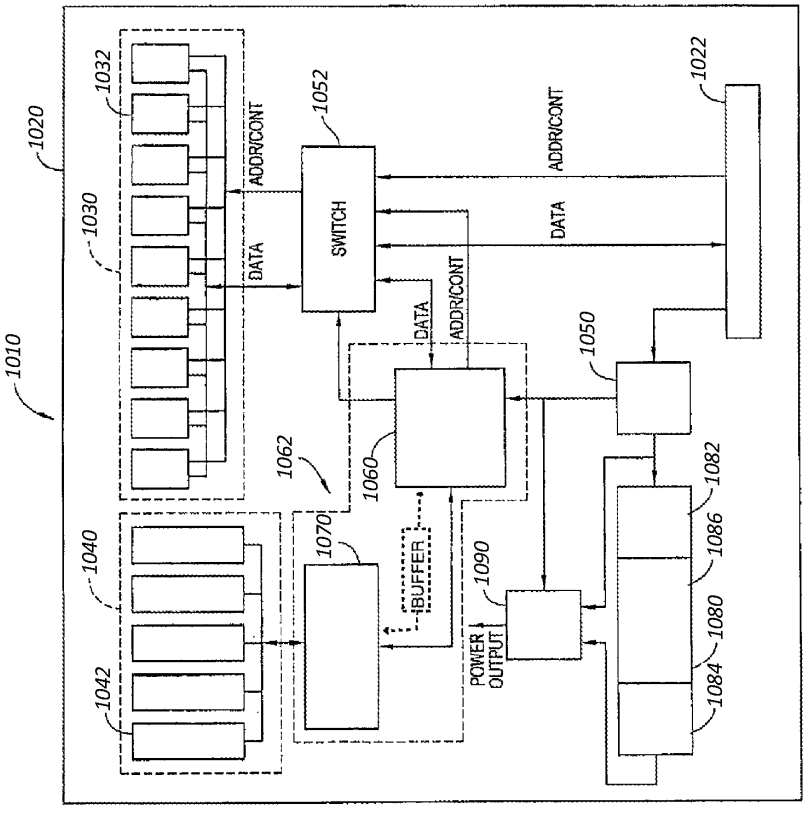


FIG. 13

Case: 24-2203 Document: 16 Page: 704 Filed: 12/16/2024

JTX0003-0017
Appx616

NETLIST_SAMSUNG_EDTX00038092

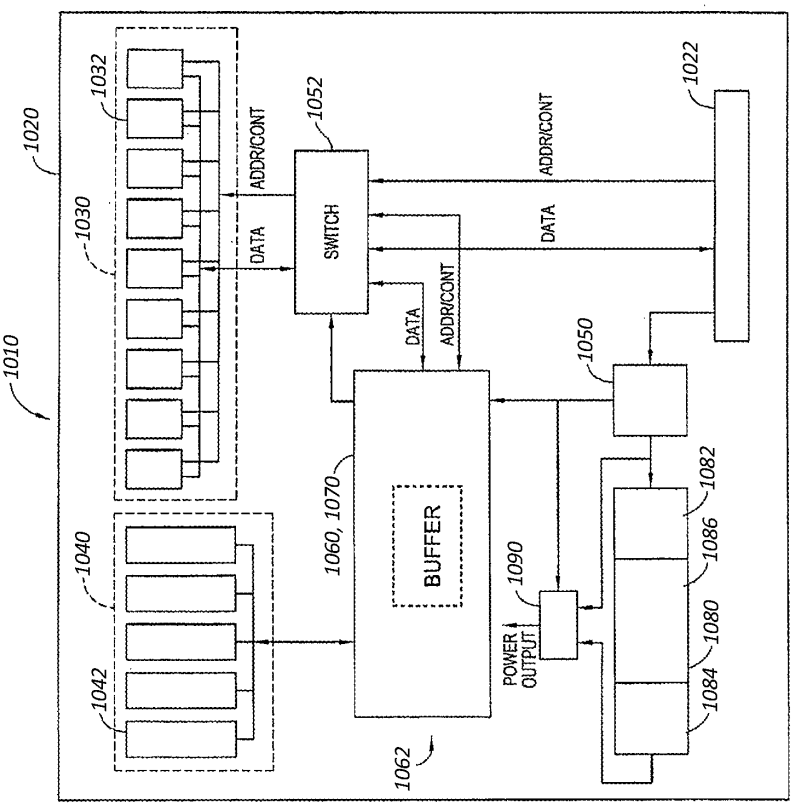


FIG. 14

Case: 24-2203 Document: 16 Page: 705 Filed: 12/16/2024

JTX0003-0018
Appx617

NETLIST_SAMSUNG_EDTX00038093

Case: 24-2203 Document: 16 Page: 706 Filed: 12/16/2024

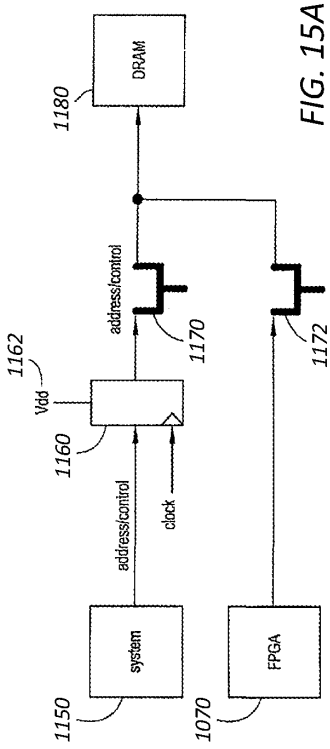


FIG. 15A

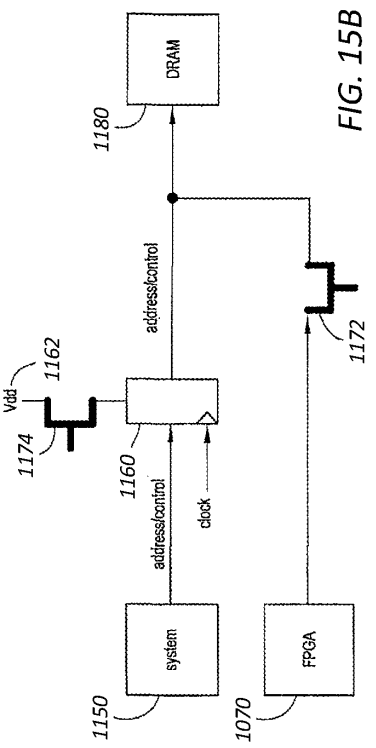


FIG. 15B

JTX0003-0019
Appx618

NETLIST_SAMSUNG_EDTX00038094

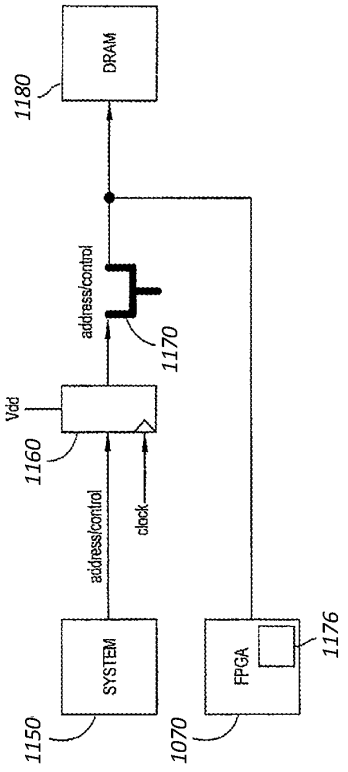

DRAM
1180
address/control
1170
Vdd
1160
address/control
clock
SYSTEM
1150
FPGA
1176
1070

FIG. 15C

Case: 24-2203 Document: 16 Page: 708 Filed: 12/16/2024

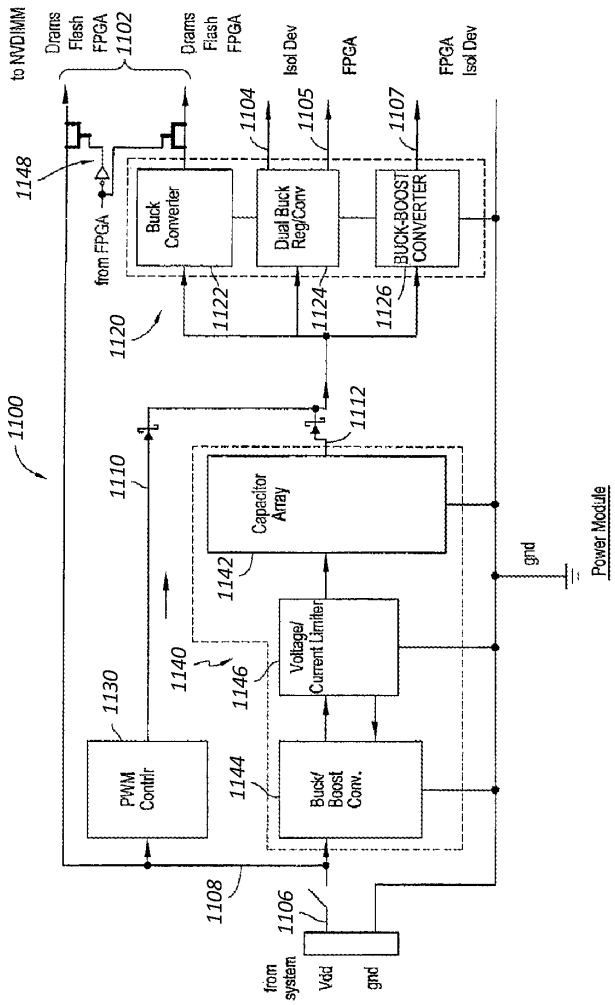


FIG. 16

JTX0003-0021
Appx620

NETLIST_SAMSUNG_EDTX00038096

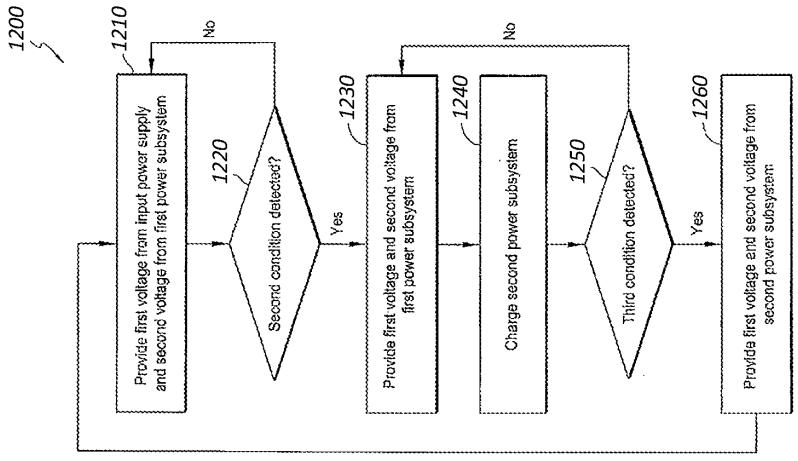


FIG. 17

Case: 24-2203 Document: 16 Page: 709 Filed: 12/16/2024

JTX0003-0022
Appx621

NETLIST_SAMSUNG_EDTX00038097

Samsung Electronics Co., Ltd.
Ex. 1077, p. 709

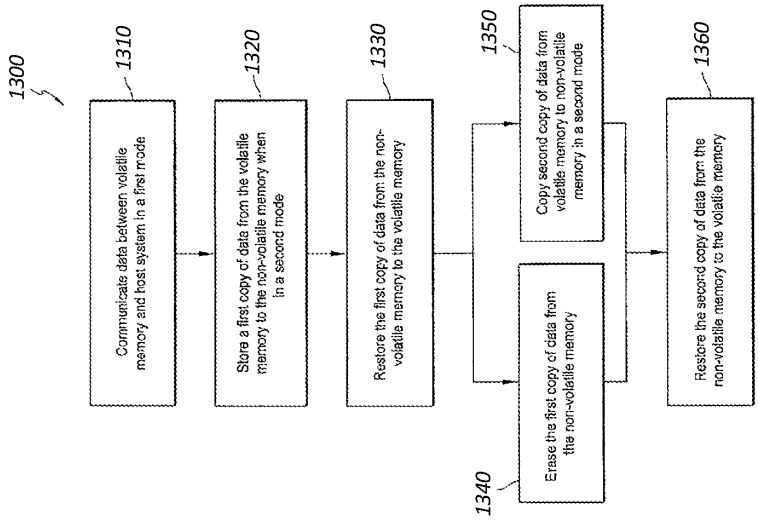


FIG. 18

Case: 24-2203 Document: 16 Page: 710 Filed: 12/16/2024

JTX0003-0023
Appx622

NETLIST_SAMSUNG_EDTX00038098

Samsung Electronics Co., Ltd.
Ex. 1077, p. 710

Case: 24-2203 Document: 16 Page: 711 Filed: 12/16/2024

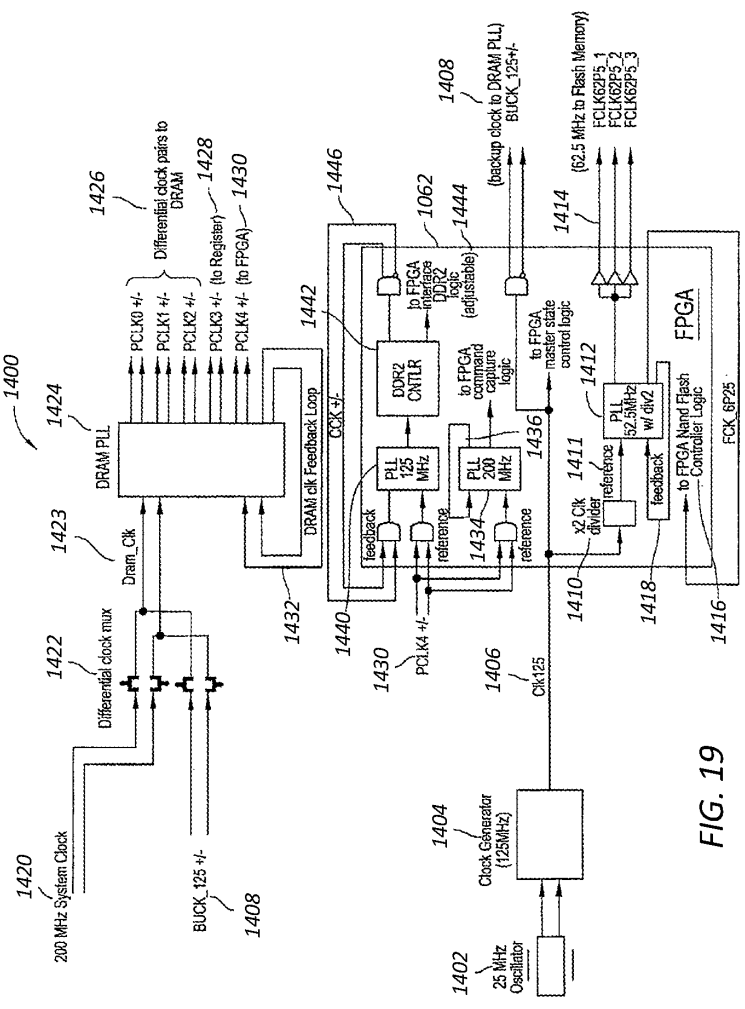


FIG. 19

JTX0003-0024
Appx623

NETLIST_SAMSUNG_EDTX00038099

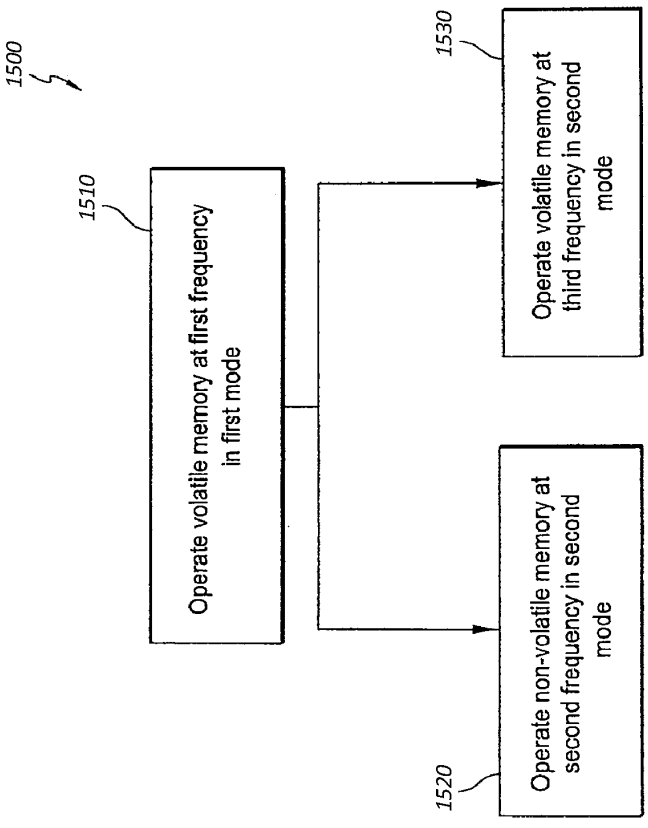

1500
1510
Operate volatile memory at first frequency in first mode
1530
Operate volatile memory at third frequency in second mode
1520
Operate non-volatile memory at second frequency in second mode

FIG. 20

Case: 24-2203 Document: 16 Page: 713 Filed: 12/16/2024

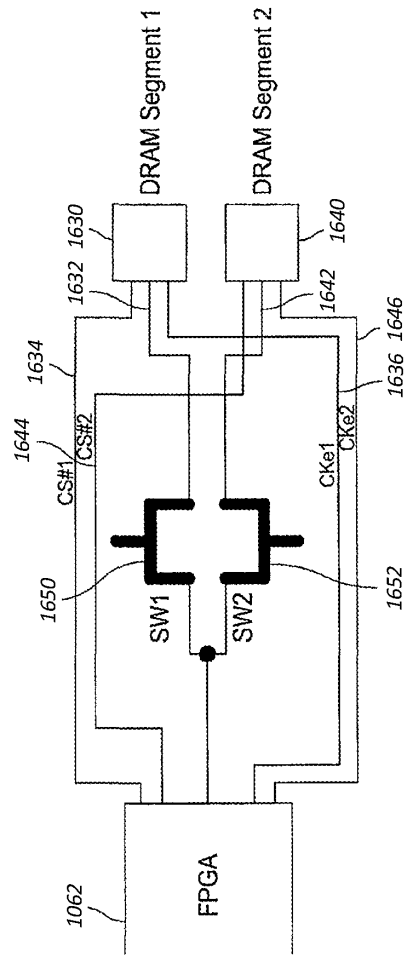


FIG. 21

JTX0003-0026
Appx625

NETLIST_SAMSUNG_EDTX00038101

Samsung Electronics Co., Ltd.
Ex. 1077, p. 713

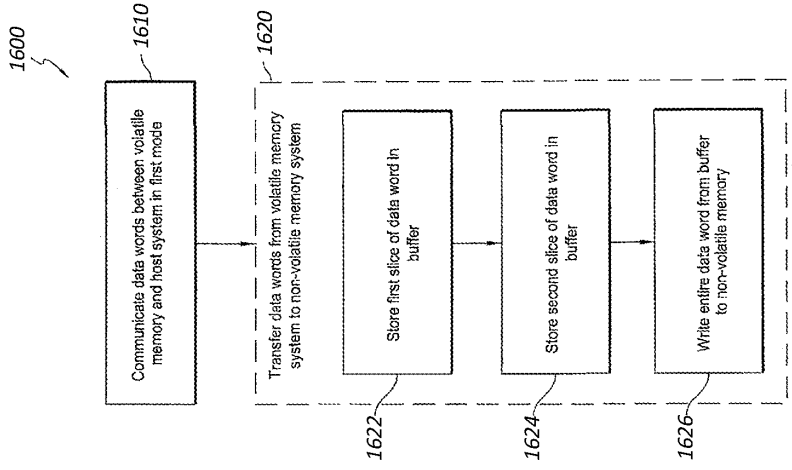


FIG. 22

Case: 24-2203 Document: 16 Page: 714 Filed: 12/16/2024

JTX0003-0027
Appx626

NETLIST_SAMSUNG_EDTX00038102

Samsung Electronics Co., Ltd.
Ex. 1077, p. 714

FLASH-DRAM HYBRID MEMORY MODULE

PRIORITY CLAIM

This application is a continuation of U.S. patent application Ser. No. 15/934,416, filed Mar. 23, 2018, titled "Flash-Dram Hybrid Memory Module," which is a continuation of U.S. patent application Ser. No. 14/840,865, filed Aug. 31, 2015, titled "Flash-Dram Hybrid Memory Module," now U.S. Pat. No. 9,928,186, which is a continuation of U.S. patent application Ser. No. 14/489,269, filed Sep. 17, 2014, titled "Flash-Dram Hybrid Memory Module," now U.S. Pat. No. 9,158,684, which is a continuation of U.S. patent application Ser. No. 13/559,476, filed Jul. 26, 2012, titled "Flash-Dram Hybrid Memory Module," now U.S. Pat. No. 8,874,831, which claims the benefit of U.S. Provisional Patent Application No. 61/512,871, filed Jul. 28, 2011, and is a continuation-in-part of U.S. patent application Ser. No. 12/240,916, filed Sep. 29, 2008, titled "Non-Volatile Memory Module," now U.S. Pat. No. 8,301,833, which is a continuation of U.S. patent application Ser. No. 12/131,873, filed Jun. 2, 2008, which claims the benefit of U.S. Provisional Patent Application No. 60/941,586, filed Jun. 1, 2007, the contents of all of which are incorporated herein by reference in their entirety.

This application may be considered related to U.S. patent application Ser. No. 14/173,242, titled "Isolation Switching For Backup Of Registered Memory," filed Feb. 5, 2014, which is a continuation of U.S. patent application Ser. No. 13/905,053, titled "Isolation Switching For Backup Of Registered Memory," filed May 29, 2013, now U.S. Pat. No. 8,677,060, issued Mar. 18, 2014, which is a continuation of U.S. patent application Ser. No. 13/536,173, titled "Data Transfer Scheme For Non-Volatile Memory Module," filed Jun. 28, 2012, now U.S. Pat. No. 8,516,187, issued Aug. 20, 2013, which is a divisional of U.S. patent application Ser. No. 12/240,916, titled "Non-Volatile Memory Module," filed Sep. 29, 2008, now U.S. Pat. No. 8,301,833, issued Oct. 30, 2012, which is a continuation of U.S. patent application Ser. No. 12/131,873, filed Jun. 2, 2008, now abandoned, which claims the benefit of U.S. Provisional Application No. 60/941,586, filed Jun. 1, 2007, the contents of which are incorporated by reference herein in their entirety.

This application may also be considered related to U.S. patent application Ser. No. 15/000,834, filed Jan. 19, 2016, which is a continuation of U.S. patent application Ser. No. 14/489,332, filed Sep. 17, 2014, now U.S. Pat. No. 9,269,437, which is a continuation of U.S. patent application Ser. No. 14/173,219, filed Feb. 5, 2014, now U.S. Pat. No. 8,904,099, which is a continuation of U.S. patent application Ser. No. 13/905,048, filed May 29, 2013, now U.S. Pat. No. 6,671,243, which is a continuation U.S. patent application Ser. No. 13/536,173 above.

This application may also be considered related to U.S. patent application Ser. No. 15/924,866, which is a continuation of U.S. patent application Ser. No. 14/489,281, filed Sep. 17, 2014, now U.S. Pat. No. 9,921,762, which is a continuation of U.S. patent application Ser. No. 13/625,563, filed Sep. 24, 2012, now U.S. Pat. No. 8,904,098, which claims the benefit of U.S. Provisional Application No. 61/583,775, filed Sep. 23, 2011.

TECHNICAL FIELD

The present disclosure relates generally to computer memory devices, and more particularly, to devices that employ different types of memory devices such as combinations of Flash and random access memories.

BACKGROUND

As technology advances and the usage of portable computing devices, such as tablet notebook computers, increases, more data needs to be transferred among data centers and to/from end users. In many cases, data centers are built by clustering multiple servers that are networked to increase performance.

Although there are many types of networked servers that are specific to the types applications envisioned, the basic concept is generally to increase server performance by dynamically allocating computing and storage resources. In recent years, server technology has evolved to be specific to particular applications such as 'finance transactions' (for example, point-of-service, inter-bank transaction, stock market transaction), 'scientific computation' (for example, fluid dynamic for automobile and ship design, weather prediction, oil and gas expeditions), 'medical diagnostics' (for example, diagnostics based on the fuzzy logic, medical data processing), 'simple information sharing and searching' (for example, web search, retail store website, company home page), 'email' (information distribution and archive), 'security service', 'entertainment' (for example, video-on-demand), and so on. However, all of these applications suffer from the same information transfer bottleneck due to the inability of a high speed CPU (central processing unit) to efficiently transfer data in and out of relatively slower speed storage or memory subsystems, particularly since data transfers typically pass through the CPU input/output (I/O) channels.

The data transfer limitations by the CPU are exemplified by the arrangement shown in FIG. **1**, and apply to data transfers between main storage (for example the hard disk (HD) or solid state drive (SSD) and the memory subsystems (for example DRAM DIMM (Dynamic Random Access Memory Dual In-line Memory Module) connected to the front side bus (FSB)). In arrangements such as that of FIG. **1**, the SSD/HD and DRAM DIMM of a conventional memory arrangement are connected to the CPU via separate memory control ports (not shown). FIG. **1** specifically shows, through the double-headed arrow, the data flow path between the computer or server main storage (SSD/HD) to the DRAM DIMMs. Since the SSD/HD data I/O and the DRAM DIMM data I/O are controlled by the CPU, the CPU needs to allocate its process cycles to control these I/Os, which may include the IRQ (Interrupt Request) service which the CPU performs periodically. As will be appreciated, the more time a CPU allocates to controlling the data transfer traffic, the less time the CPU has to perform other tasks. Therefore, the overall performance of a server will deteriorate with the increased amount of time the CPU has to expend in performing data transfer.

There have been various approaches to increase the data transfer throughput rates from/to the main storage, such as SSD/HD, to local storage, such as DRAM DIMM. In one example as illustrated in FIG. **2**, EcoRAM™ developed by Spansion provides a storage SSD based system that assumes a physical form factor of a DIMM. The EcoRAM™ is populated with Flash memories and a relatively small memory capacity using DRAMs which serve as a data buffer. This arrangement is capable of delivering higher throughput rate than a standard SSD based system since the EcoRAM™ is connected to the CPU (central processing unit) via a high speed interface, such as the HT (Hyper

Transport) interface, while an SSD/HD is typically connected via SATA (serial AT attachment), USB (universal serial bus), or PCI Express (peripheral component interface express). For example, the read random access throughput rate of EcoRAM™ is near 3 GB/s compared with 400 MB/s for a NAND SSD memory subsystem using the standard PCI Express-based. This is a 7.5× performance improvement. However, the performance improvement for write random access throughput rate is less than 2×(197 MBs for the EcoRAM vs. 104 MBs for NAND SSD). This is mainly due to the fact that the write speed is cannot be faster than the NAND Flash write access time. FIG. **2** is an example of EcoRAM™ using SSD with the form factor of a standard DIMM such that it can be connected to the FSB (front side bus). However, due to the interface protocol difference between DRAM and Flash, an interface device, EcoRAM Accelerator™), which occupies one of the server's CPU sockets is used, and hence further reducing server's performance by reducing the number of available CPU sockets available, and in turn reducing the overall computation efficiency. The server's performance will further suffer due to the limited utilization of the CPU bus due to the large difference in the data transfer throughput rate between read and write operations.

The EcoRAM™ architecture enables the CPU to view the Flash DIMM controller chip as another processor with a large size of memory available for CPU access.

In general, the access speed of a Flash based system is limited by four items: the read/write speed of the Flash memory, the CPU's FSB bus speed and efficiency, the Flash DIMM controller's inherent latency, and the HT interconnect speed and efficiency which is dependent on the HT interface controller in the CPU and Flash DIMM controller chip.

The published results indicate that these shortcomings are evident in that the maximum throughput rate is 1.56 GBs for the read operation and 104 MBs for the write operation. These access rates are 25% of the DRAM read access speed, and 1.7% of the DRAM access speed at 400 MHz operation. The disparity in the access speed (15 to 1) between the read operation and write operation highlight a major disadvantage of this architecture. The discrepancy of the access speed between this type of architecture and JEDEC standard DRAM DIMM is expected to grow wider as the DRAM memory technology advances much faster than the Flash memory.

Certain types of memory modules comprise a plurality of dynamic random-access memory (DRAM) devices mounted on a printed circuit board (PCB). These memory modules are typically mounted in a memory slot or socket of a computer system (e.g., a server system or a personal computer) and are accessed by the computer system to provide volatile memory to the computer system.

Volatile memory generally maintains stored information only when it is powered. Batteries have been used to provide power to volatile memory during power failures or interruptions. However, batteries may require maintenance, may need to be replaced, are not environmentally friendly, and the status of batteries can be difficult to monitor.

Non-volatile memory can generally maintain stored information while power is not applied to the non-volatile memory. In certain circumstances, it can therefore be useful to backup volatile memory using non-volatile memory.

OVERVIEW

Described herein is a memory module couplable to a memory controller of a host system. The memory module includes a non-volatile memory subsystem, a data manager coupled to the non-volatile memory subsystem, a volatile memory subsystem coupled to the data manager and operable to exchange data with the non-volatile memory subsystem by way of the data manager, and a controller operable to receive commands from the memory controller and to direct (i) operation of the non-volatile memory subsystem, (ii) operation of the volatile memory subsystem, and (iii) transfer of data between any two or more of the memory controller, the volatile memory subsystem, and the non-volatile memory subsystem based on at least one received command from the memory controller.

Also described herein is a method for managing a memory module by a memory controller, the memory module including volatile and non-volatile memory subsystems. The method includes receiving control information from the memory controller, wherein the control information is received using a protocol of the volatile memory subsystem. The method further includes identifying a data path to be used for transferring data to or from the memory module using the received control information, and using a data manager and a controller of the memory module to transfer data between any two or more of the memory controller, the volatile memory subsystem, and the non-volatile memory subsystem based on at least one of the received control information and the identified data path.

Also described herein is a memory module wherein the data manager is operable to control one or more of data flow rate, data transfer size, data buffer size, data error monitoring, and data error correction in response to receiving at least one of a control signal and control information from the controller.

Also described herein is a memory module wherein the data manager controls data traffic between any two or more of the memory controller, the volatile memory subsystem, and the non-volatile memory subsystem based on instructions received from the controller.

Also described herein is a memory module wherein data traffic control relates to any one or more of data flow rate, data transfer size, data buffer size, data transfer bit width, formatting information, direction of data flow, and the starting time of data transfer.

Also described herein is a memory module wherein the controller configures at least one of a first memory address space of the volatile memory subsystem and a second memory address space of the non-volatile memory subsystem in response to at least one of a received command from the memory controller and memory address space initialization information of the memory module.

Also described herein is a memory module wherein the data manager is configured as a bi-directional data transfer fabric having two or more sets of data ports coupled to any one of the volatile and non-volatile memory subsystems.

Also described herein is a memory module wherein at least one of the volatile and non-volatile memory subsystems comprises one or more memory segments.

Also described herein is a memory module wherein each memory segment comprises at least one memory circuit, memory device, or memory die.

Also described herein is a memory module wherein the volatile memory subsystem comprises DRAM memory.

Also described herein is a memory module wherein the non-volatile memory subsystem comprises flash memory.

Also described herein is a memory module wherein at least one set of data ports is operated by the data manager to

independently and/or concurrently transfer data to or from one or more memory segments of the volatile or non-volatile memory subsystems.

Also described herein is a memory module wherein the data manager and controller are configured to effect data transfer between the memory controller and the non-volatile memory subsystem in response to memory access commands received by the controller from the memory controller.

Also described herein is a memory module wherein the volatile memory subsystem is operable as a buffer for the data transfer between the memory controller and non-volatile memory.

Also described herein is a memory module wherein the data manager further includes a data format module configured to format data to be transferred between any two or more of the memory controller, the volatile memory subsystem, and the non-volatile memory subsystem based on control information received from the controller.

Also described herein is a memory module wherein the data manager further includes a data buffer for buffering data delivered to or from the non-volatile memory subsystem.

Also described herein is a memory module wherein the controller is operable to perform one or more of memory address translation, memory address mapping, address domain conversion, memory access control, data error correction, and data width modulation between the volatile and non-volatile memory subsystems.

Also described herein is a memory module wherein the controller is configured to effect operation with the host system in accordance with a prescribed protocol.

Also described herein is a memory module wherein the prescribed protocol is selected from one or more of DDR, DDR2, DDR3, and DDR4 protocols.

Also described herein is a memory module wherein the controller is operable to configure memory space in the memory module based on at least one of a command received from the memory controller, a programmable value written into a register, a value corresponding to a first portion of the volatile memory subsystem, a value corresponding to a first portion of the non-volatile memory subsystem, and a timing value.

Also described herein is a memory module wherein the controller configures the memory space of the memory module using at least a first portion of the volatile memory subsystem and a first portion of the non-volatile memory subsystem, and the controller presents a unified memory space to the memory controller.

Also described herein is a memory module wherein the controller configures the memory space in the memory module using partitioning instructions that are application-specific.

Also described herein is a memory module wherein the controller is operable to copy booting information from the non-volatile to the volatile memory subsystem during power up.

Also described herein is a memory module wherein the controller includes a volatile memory control module, a non-volatile memory control module, data manager control module, a command interpreter module, and a scheduler module.

Also described herein is a memory module wherein commands from the volatile memory control module to the volatile memory subsystem are subordinated to commands from the memory controller to the controller.

Also described herein is a memory module wherein the controller effects pre-fetching of data from the non-volatile to the volatile memory.

Also described herein is a memory module wherein the pre-fetching is initiated by the memory controller writing an address of requested data into a register of the controller.

Also described herein is a memory module wherein the controller is operable to initiate a copy operation of data of a closed block in the volatile memory subsystem to a target block in the non-volatile memory subsystem.

Also described herein is a memory module wherein, if the closed block is re-opened, the controller is operable to abort the copy operation and to erase the target block from the non-volatile memory subsystem.

Also described herein is a method for managing a memory module wherein the transfer of data includes a bidirectional transfer of data between the non-volatile and the volatile memory subsystems.

Also described herein is a method for managing a memory module further comprising operating the data manager to control one or more of data flow rate, data transfer size, data width size, data buffer size, data error monitoring, data error correction, and the starting time of the transfer of data.

Also described herein is a method for managing a memory module further comprising operating the data manager to control data traffic between the memory controller and at least one of the volatile and non-volatile memory subsystems.

Also described herein is a method for managing a memory module wherein data traffic control relates to any one or more of data transfer size, formatting information, direction of data flow, and the starting time of the transfer of data.

Also described herein is a method for managing a memory module wherein data traffic control by the data manager is based on instructions received from the controller.

Also described herein is a method for managing a memory module further comprising operating the data manager as a bi-directional data transfer fabric with two or more sets of data ports coupled to any one of the volatile and non-volatile memory subsystems.

Also described herein is a method for managing a memory module wherein at least one of the volatile and non-volatile memory subsystems comprises one or more memory segments.

Also described herein is a method for managing a memory module wherein each memory segment comprises at least one memory circuit, memory device, or memory die.

Also described herein is a method for managing a memory module wherein the volatile memory subsystem comprises DRAM memory.

Also described herein is a method for managing a memory module wherein the non-volatile memory subsystem comprises Flash memory.

Also described herein is a method for managing a memory module further comprising operating the data ports to independently and/or concurrently transfer data to or from one or more memory segments of the volatile or non-volatile memory subsystems.

Also described herein is a method for managing a memory module further comprising directing transfer of data bi-directionally between the volatile and non-volatile memory subsystems using the data manager and in response to memory access commands received by the controller from the memory controller.

Also described herein is a method for managing a memory module further comprising buffering the data transferred between the memory controller and non-volatile memory subsystem using the volatile memory subsystem.

Also described herein is a method for managing a memory module further comprising using the controller to perform one or more of memory address translation, memory address mapping, address domain conversion, memory access control, data error correction, and data width modulation between the volatile and non-volatile memory subsystems.

Also described herein is a method for managing a memory module further comprising using the controller to effect communication with a host system by the volatile memory subsystem in accordance with a prescribed protocol.

Also described herein is a method for managing a memory module wherein the prescribed protocol is selected from one or more of DDR, DDR2, DDR3, and DDR4 protocols.

Also described herein is a method for managing a memory module further comprising using the controller to configure memory space in the memory module based on at least one of a command received from the memory controller, a programmable value written into a register, a value corresponding to a first portion of the volatile memory subsystem, a value corresponding to a first portion of the non-volatile memory subsystem, and a timing value.

Also described herein is a method for managing a memory module wherein the controller configures the memory space of the memory module using at least a first portion of the volatile memory subsystem and a first portion of the non-volatile memory subsystem, and the controller presents a unified memory space to the memory controller.

Also described herein is a method for managing a memory module wherein the controller configures the memory space in the memory module using partitioning instructions that are application-specific.

Also described herein is a method for managing a memory module further comprising using the controller to copy booting information from the non-volatile to the volatile memory subsystem during power up.

Also described herein is a method for managing a memory module wherein the controller includes a volatile memory control module, the method further comprising generating commands by the volatile memory control module in response to commands from the memory controller, and transmitting the generated commands to the volatile memory subsystem.

Also described herein is a method for managing a memory module further comprising pre-fetching of data from the non-volatile memory subsystem to the volatile memory subsystem.

Also described herein is a method for managing a memory module wherein the pre-fetching is initiated by the memory controller writing an address of requested data into a register of the controller.

Also described herein is a method for managing a memory module further comprising initiating a copy operation of data of a closed block in the volatile memory subsystem to a target block in the non-volatile memory subsystem.

Also described herein is a method for managing a memory module further comprising aborting the copy operation when the closed block of the volatile memory subsystem is re-opened, and erasing the target block in the non-volatile memory subsystem.

Also described herein is a memory system having a volatile memory subsystem, a non-volatile memory subsystem, a controller coupled to the non-volatile memory subsystem, and a circuit coupled to the volatile memory subsystem, to the controller, and to a host system. In a first mode of operation, the circuit is operable to selectively isolate the controller from the volatile memory subsystem, and to selectively couple the volatile memory subsystem to the host system to allow data to be communicated between the volatile memory subsystem and the host system. In a second mode of operation, the circuit is operable to selectively couple the controller to the volatile memory subsystem to allow data to be communicated between the volatile memory subsystem and the nonvolatile memory subsystem using the controller, and the circuit is operable to selectively isolate the volatile memory subsystem from the host system.

Also described herein is a method for operating a memory system. The method includes coupling a circuit to a host system, a volatile memory subsystem, and a controller, wherein the controller is coupled to a non-volatile memory subsystem. In a first mode of operation that allows data to be communicated between the volatile memory subsystem and the host system, the circuit is used to (i) selectively isolate the controller from the volatile memory subsystem, and (ii) selectively couple the volatile memory subsystem to the host system. In a second mode of operation that allows data to be communicated between the volatile memory subsystem and the nonvolatile memory subsystem via the controller, the circuit is used to (i) selectively couple the controller to the volatile memory subsystem, and (ii) selectively isolate the volatile memory subsystem from the host system.

Also described herein is a nontransitory computer readable storage medium storing one or more programs configured to be executed by one or more computing devices. The programs, when executing on the one or more computing devices, cause a circuit that is coupled to a host system, to a volatile memory subsystem, and to a controller that is coupled to a nonvolatile memory subsystem, to perform a method in which, in a first mode of operation that allows data to be communicated between the volatile memory subsystem and the host system, operating the circuit to (i) selectively isolate the controller from the volatile memory subsystem, and (ii) selectively couple the volatile memory subsystem to the host system. In a second mode of operation that allows data to be communicated between the volatile memory subsystem and the nonvolatile memory subsystem via the controller, operating the circuit to (i) selectively couple the controller to the volatile memory subsystem, and (ii) selectively isolate the volatile memory subsystem from the host system.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated into and constitute a part of this specification, illustrate one or more examples of embodiments and, together with the description of example embodiments, serve to explain the principles and implementations of the embodiments.

In the drawings:

FIG. **1** is a block diagram illustrating the path of data transfer, via a CPU, of a conventional memory arrangement;

FIG. **2** is a block diagram of a known EcoRAM™ architecture;

FIGS. **3**A and **3**B are block diagrams of a non-volatile memory DIMM or NVDIMM;

FIGS. **4**A and **4**B are block diagrams of a Flash-DRAM hybrid DIMM or FDHDIMM;

FIG. **5**A is a block diagram of a memory module **500** in accordance with certain embodiments described herein;

FIG. **5**B is a block diagram showing some functionality of a memory module such as that shown in FIG. **5**A;

FIG. **6** is a block diagram showing some details of the data manager (DMgr);

FIG. **7** is a functional block diagram of the on-module controller (CDC);

FIG. **8**A is a block diagram showing more details of the prior art Flash-DRAM hybrid DIMM (FDHDIMM) of FIGS. **4**A and **4**B;

FIG. **8**B is a block diagram of a Flash-DRAM hybrid DIMM (FDHDIMM) in accordance with certain embodiments disclosed herein;

FIG. **9** is a flow diagram directed to the transfer of data from Flash memory to DRAM memory and vice versa in an exemplary FDHDIMM;

FIG. **10** is a block diagram showing an example of mapping of DRAM address space to Flash memory address space; and

FIG. **11** is a table showing estimates of the maximum allowed closed blocks in a queue to be written back to Flash memory for different DRAM densities using various average block use time.

FIG. **12** is a block diagram of an example memory system compatible with certain embodiments described herein.

FIG. **13** is a block diagram of an example memory module with ECC (error-correcting code) having a volatile memory subsystem with nine volatile memory elements and a non-volatile memory subsystem with five non-volatile memory elements in accordance with certain embodiments described herein.

FIG. **14** is a block diagram of an example memory module having a microcontroller unit and logic element integrated into a single device in accordance with certain embodiments described herein.

FIGS. **15**A-**15**C schematically illustrate example embodiments of memory systems having volatile memory subsystems comprising registered dual in-line memory modules in accordance with certain embodiments described herein.

FIG. **16** schematically illustrates an example power module of a memory system in accordance with certain embodiments described herein.

FIG. **17** is a flowchart of an example method of providing a first voltage and a second voltage to a memory system including volatile and non-volatile memory subsystems.

FIG. **18** is a flowchart of an example method of controlling a memory system operatively coupled to a host system and which includes at least 100 percent more storage capacity in non-volatile memory than in volatile memory.

FIG. **19** schematically illustrates an example clock distribution topology of a memory system in accordance with certain embodiments described herein.

FIG. **20** is a flowchart of an example method of controlling a memory system operatively coupled to a host system, the method including operating a volatile memory subsystem at a reduced rate in a back-up mode.

FIG. **21** schematically illustrates an example topology of a connection to transfer data slices from two DRAM segments of a volatile memory subsystem of a memory system to a controller of the memory system.

FIG. **22** is a flowchart of an example method of controlling a memory system operatively coupled to a host system, the method including backing up and/or restoring a volatile memory subsystem in slices.

DESCRIPTION OF EXAMPLE EMBODIMENTS

Example embodiments are described herein in the context of a system of computers, servers, controllers, memory modules, hard disk drives and software. Those of ordinary skill in the art will realize that the following description is illustrative only and is not intended to be in any way limiting. Other embodiments will readily suggest themselves to such skilled persons having the benefit of this disclosure. Reference will now be made in detail to implementations of the example embodiments as illustrated in the accompanying drawings. The same reference indicators will be used to the extent possible throughout the drawings and the following description to refer to the same or like items.

In the interest of clarity, not all of the routine features of the implementations described herein are shown and described. It will, of course, be appreciated that in the development of any such actual implementation, numerous implementation-specific decisions must be made in order to achieve the developer's specific goals, such as compliance with application- and business-related constraints, and that these specific goals will vary from one implementation to another and from one developer to another. Moreover, it will be appreciated that such a development effort might be complex and time-consuming, but would nevertheless be a routine undertaking of engineering for those of ordinary skill in the art having the benefit of this disclosure.

In accordance with this disclosure, the components, process steps, and/or data structures described herein may be implemented using various types of operating systems, computing platforms, computer programs, and/or general purpose machines. In addition, those of ordinary skill in the art will recognize that devices of a less general purpose nature, such as hardwired devices, field programmable gate arrays (FPGAs), application specific integrated circuits (ASICs), or the like, may also be used without departing from the scope and spirit of the inventive concepts disclosed herein. Where a method comprising a series of process steps is implemented by a computer or a machine and those process steps can be stored as a series of instructions readable by the machine, they may be stored on a tangible medium such as a computer memory device (e.g., ROM (Read Only Memory), PROM (Programmable Read Only Memory), EEPROM (Electrically Eraseable Programmable Read Only Memory), Flash memory, Jump Drive, and the like), magnetic storage medium (e.g., tape, magnetic disk drive, and the like), optical storage medium (e.g., CD-ROM, DVD-ROM, paper card, paper tape and the like) and other types of program memory.

The term "exemplary" where used herein is intended to mean "serving as an example, instance or illustration." Any embodiment described herein as "exemplary" is not necessarily to be construed as preferred or advantageous over other embodiments.

Disclosed herein are arrangements for improving memory access rates and addressing the high disparity (15 to 1 ratio) between the read and write data throughput rates. In one arrangement, a Flash-DRAM-hybrid DIMM (FDHDIMM) with integrated Flash and DRAM is used. Methods for controlling such an arrangement are described.

In certain embodiments, the actual memory density (size or capacity) of the DIMM and/or the ratio of DRAM memory to Flash memory are configurable for optimal use with a particular application (for example, POS, inter-bank transaction, stock market transaction, scientific computation such as fluid dynamics for automobile and ship design, weather prediction, oil and gas expeditions, medical diagnostics such as diagnostics based on the fuzzy logic, medical data processing, simple information sharing and searching such as web search, retail store website, company home

page, email or information distribution and archive, security service, and entertainment such as video-on-demand).

In certain embodiments, the device contains a high density Flash memory with a low density DRAM, wherein the DRAM is used as a data buffer for read/write operation. The Flash serves as the main memory. Certain embodiments described herein overcome the needs of having a long separation period between an Activate command (may be referred to as RAS) and a corresponding read or write command (may be referred to as first CAS command).

In accordance with one embodiment, described with reference to FIGS. **3**A and **3**B, a memory system **300** includes a non-volatile (for example Flash) memory subsystem **302** and a volatile (for example DRAM) memory subsystem **304**. The examples of FIGS. **3**A and **3**B are directed to architectures of a non-volatile DIMM (NVDIMM) NVDIMM system that may use a power subsystem (not shown) that can include a battery or a capacitor as a means for energy storage to copy DRAM memory data into Flash memory when power loss occurs, is detected, or is anticipated to occur during operation. When normal power is restored, a restore NVDIMM operation is initiated and the data stored in the Flash memory is properly restored to the DRAM memory. In this architecture, the density of the Flash is about the same as the DRAM memory size or within a few multiples, although in some applications it may be higher. This type of architecture may also be used to provide non-volatile storage that is connected to the FSB (front side bus) to support RAID (Redundant Array of Independent Disks) based systems or other type of operations. An NVDIMM controller **306** receives and interprets commands from the system memory controller hub (MCH). The NVDIMM controller **306** control the NVDIMM DRAM and Flash memory operations. In FIG. **3**A, the DRAM **304** communicates data with the MCH, while an internal bus **308** is used for data transfer between the DRAM and Flash memory subsystems. In FIG. **3**B, the NVDIMM controller **306'** of NVDIMM **300'** monitors events or commands and enables data transfer to occur in a first mode between the DRAM **304'** and Flash **302'** or in a second mode between the DRAM and the MCH.

In accordance with one embodiment, a general architecture for a Flash and DRAM hybrid DIMM (FDHDIMM) system **400** is shown in FIG. **4**A. The FDHDIMM interfaces with an MCH (memory controller hub) to operate and behave as a high density DIMM, wherein the MCH interfaces with the non-volatile memory subsystem (for example Flash) **402** is controlled by an FDHDIMM controller **404**. Although the MCH interfaces with the Flash via the FDHDIMM controller, the FDHDIMM overall performance is governed by the Flash access time. The volatile memory subsystem (for example DRAM) **406** is primarily used as a data buffer or a temporary storage location such that data from the Flash memory **402** is transferred to the DRAM **406** at the Flash access speed, and buffered or collected into the DRAM **406**, which then transfers the buffered data to the MCH based on the access time of DRAM. Similarly, when the MCH transfers data to the DRAM **406**, the FDHDIMM controller **404** manages the data transfer from the DRAM **406** to the Flash **402**. Since the Flash memory access speed (both read and write) is relatively slower than DRAM, (e.g. for example a few hundred microseconds for read access), the average data throughput rate of FDHDIMM **400** is limited by the Flash access speed. The DRAM **406** serves as a data buffer stage that buffers the MCH read or write data. Thus, the DRAM **406** serves as a temporary storage for the data to be transferred from/to the Flash **402**. Furthermore, in accordance with one embodiment, the MCH recognizes the physical density of an FDHDIMM operating as a high density DIMM as the density of Flash alone.

In accordance with one embodiment, a read operation can be performed by the MCH by sending an activate command (may be simply referred to as RAS, or row address strobe) to the FDHDIMM **400** to conduct a pre-fetch read data operation from the Flash **402** to the DRAM **406**, with the pre-fetch data size being for example a page (1 KB or 2 KB, or may be programmable to any size). The MCH then sends a read command (may be simply referred to as CAS, or column address strobe) to read the data out input of the DRAM. In this embodiment, the data transfer from Flash to DRAM occurs at Flash access speed rates, while data transfer from DRAM to MCH occurs at DRAM access speed rates. In this example, data latency and throughput rates are the same as any DRAM operation as long as the read operations are executed onto the pages that were opened with the activate command previously sent to pre-fetch data from the Flash to DRAM. Thus, a longer separation time period between the RAS (e.g. Activate command) and the first CAS (column address strobe e.g. read or write command) is required to account for the time it takes to pre-fetch data from the Flash to DRAM.

An example of FDHDIMM operating as a DDR DIMM with SSD is shown in FIG. **4**B, wherein the FDHDIMM **400'** supports two different interface interpretations to the MCH. In the first interface interpretation, the MCH views the FDHDIMM **400'** as a combination of DRAM DIMM and SSD (not illustrated). In this mode the MCH needs to manage two address spaces, one for the DRAMs **402'** and one for the Flash **404'**. The MCH is coupled to, and controls, both of the DRAM and Flash memory subsystems. One advantage of this mode is that the CPU does not need to be in the data path when data is moved from DRAM to Flash or from Flash to DRAM. In the second interface interpretation, the MCH views the FDHDIMM **400'** as an on-DIMM Flash with the SSD in an extended memory space that is behind the DRAM space. Thus, in this mode, the MCH physically fetches data from the SSD to the DDR DRAM and then the DRAM sends the data to the MCH. Since all data movement occurs on the FDHDIMM, this mode will provide better performance than if the data were to be moved through or via the CPU.

In accordance with one embodiment and as shown in FIG. **4**B, the FDHDIMM **400'** receives control signals **408** from the MCH, where the control signals may include one or more control signals specifically for the DRAM **402'** operation and one or more control signals specifically for the Flash **404'** operation. In this embodiment, the MCH or CPU is coupled to the FDHDIMM via a single data bus interface **410** which couples the MCH to the DRAM.

FIGS. **5**A and **5**B are block diagrams of a memory module **500** that is couplable to a host system (not shown). The host system may be a server or any other system comprising a memory system controller or an MCH for providing and controlling the read/write access to one or more memory systems, wherein each memory system may include a plurality of memory subsystems, a plurality of memory devices, or at least one memory module. The term "read/write access" means the ability of the MCH to interface with a memory system or subsystem in order to write data into it or read data from it, depending on the particular requirement at a particular time.

In certain embodiments, memory module **500** is a Flash-DRAM hybrid memory subsystem which may be integrated with other components of a host system. In certain embodiments, memory module **500** is a Flash-DRAM hybrid

memory module that has the DIMM (dual-inline memory module) form factor, and may be referred to as a FDHDIMM, although it is to be understood that in both structure and operation it may be different from the FDHDIMM discussed above and described with reference to FIGS. **4**A and **4**B. Memory module **500** includes two on-module intermediary components: a controller and a data manager. These on-module intermediary components may be physically separate components, circuits, or modules, or they may be integrated onto a single integrated circuit or device, or integrated with other memory devices, for example in a three dimensional stack, or in any one of several other possible expedients for integration known to those skilled in the art to achieve a specific design, application, or economic goal. In the case of a DIMM, these on-module intermediary components are an on-DIMM Controller (CDC) **502** and an on-DIMM data manager (DMgr) **504**. While the DIMM form factor will predominate the discussion herein, it should be understood that this is for illustrative purposes only and memory systems using other form factors are contemplated as well. CDC **502** and data manager DMgr **504** are operative to manage the interface between a non-volatile memory subsystem such as a Flash **506**, a volatile memory subsystem such as a DRAM **508**, and a host system represented by MCH **510**.

In certain embodiments, CDC **502** controls the read/write access to/from Flash memory **506** from/to DRAM memory **508**, and to/from DRAM memory from/to MCH **510**.

Read/write access between DRAM **508**, Flash **506** and MCH **510** may be referred to herein generally as communication, wherein control and address information C/A **560** is sent from MCH **510** to CDC **502**, and possible data transfers follow as indicated by Data **550**, Data **555**, and/or Data **556**. In certain embodiments, the CDC **502** performs specific functions for memory address transformation, such as address translation, mapping, or address domain conversion, Flash access control, data error correction, manipulation of data width or data formatting or data modulation between the Flash memory and DRAM, and so on. In certain embodiments, the CDC **502** ensures that memory module **500** provides transparent operation to the MCH in accordance with certain industry standards, such as DDR, DDR2, DDR3, DDR4 protocols. In the arrangement shown in FIGS. **5**A and **5**B, there is no direct access from the MCH **510** to the Flash **506** memory subsystem. Thus in accordance with certain embodiments, the Flash access speed has minimal impact on the overall FDHDIMM access speed. In the schematic illustration of FIG. **5**B and in accordance with one embodiment, the CDC controller **502** receives standard DDR commands from the MCH, interprets, and produces commands and/or control signals to control the operation of the Data manager (DMgr), the Flash memory and the DRAM memory. The DMgr controls the data path routing amongst DRAMs, Flash and MCH, as detailed below. The data path routing control signals are independently operated without any exclusivity.

An exemplary role of DMgr **504** is described with reference to FIG. **6**. In certain embodiments and in response to communication from CDC **502**, DMgr **504** provides a variety of functions to control data flow rate, data transfer size, data buffer size, data error monitoring or data error correction. For example, these functions or operations can be performed on-the-fly (while data is being transferred via the DMgr **504**) or performed on buffered or stored data in DRAM or a buffer. In addition, one role of DMgr **504** is to provide interoperability among various memory subsystems or components and/or MCH **510**.

In one embodiment, an exemplary host system operation begins with initialization. The CDC **502** receives a first command from the MCH **510** to initialize FDHDIMM **500** using a certain memory space. The memory space as would be controlled by MCH **510** can be configured or programmed during initialization or after initialization has completed. The MCH **510** can partition or parse the memory space in various ways that are optimized for a particular application that the host system needs to run or execute. In one embodiment, the CDC **502** maps the actual physical Flash **506** and DRAM **508** memory space using the information sent by MCH **510** via the first command. In one embodiment, the CDC **502** maps the memory address space of any one of the Flash **506** and DRAM **508** memory subsystems using memory address space information that is received from the host system, stored in a register within FDHDIMM **500**, or stored in a memory location of a non-volatile memory subsystem, for example a portion of Flash **506** or a separate non-volatile memory subsystem. In one embodiment, the memory address space information corresponds to a portion of initialization information of the FDHDIMM **500**.

In one embodiment, MCH **510** may send a command to restore a certain amount of data information from Flash **506** to DRAM **508**. The CDC **502** provides control information to DMgr **504** to appropriately copy the necessary information from Flash **506** to the DRAM **508**. This operation can provide support for various host system booting operations and/or a special host system power up operation.

In one embodiment, MCH **510** sends a command which may include various fields comprising control information regarding data transfer size, data format options, and/or startup time. CDC **502** receives and interprets the command and provides control signals to DMgr **504** to control the data traffic between the Flash **506**, the DRAM **508**, and the MCH **510**. For example, DMgr **504** receives the data transfer size, formatting information, direction of data flow (via one or more multiplexers such as **611**, **612**, **621**, **622** as detailed below), and the starting time of the actual data transfer from CDC **502**. DMgr **504** may also receive additional control information from the CDC **502** to establish a data flow path and/or to correctly establish the data transfer fabric. In certain embodiments, DMgr **504** also functions as a bi-directional data transfer fabric. For example, DMgr **504** may have more than 2 sets of data ports facing the Flash **506** and the DRAM **508**. Multiplexers **611** and **612** provide controllable data paths from any one of the DRAMs **508**(**1**) and **508**(**2**) (DRAM-A and DRAM-B) to any one of the MCH **510** and the Flash **506**. Similarly multiplexers **621** and **622** provide controllable data paths from any one of the MCH and the Flash memory to any one of the DRAMs **508**(**1**) and **508**(**2**) (DRAM-A and DRAM-B). In one embodiment, DRAM **508**(**1**) is a segment of DRAM **508**, while in other embodiments, DRAM **508**(**1**) is a separate DRAM memory subsystem. It will be understood that each memory segment can comprise one or more memory circuits, a memory devices, and/or memory integrated circuits. Of course other configurations for DRAM **508** are possible, and other data transfer fabrics using complex data paths and suitable types of multiplexing logic are contemplated.

In accordance with one embodiment, the two sets of multiplexors **611**, **612** and **621**, **622** allow independent data transfer to Flash **506** from DRAM-A **508**(**1**) and DRAM-B **508**(**2**). For example, in response to one or more control signals or a command from CDC **502**, DMgr **504** can transfer data from DRAM-A **508**(**1**) to MCH **510**, via multiplexer **611**, at the same time as from DRAM-B **508**(**2**)

to the Flash **506**, via multiplexer **612**; or data is transferred from DRAM-B **508**(**2**) to MCH **510**, via multiplexer **611**, and simultaneously data is transferred from the Flash **506** to DRAM-A **508**(**1**), via multiplexer **621**. Further, in the same way that data can be transferred to or from the DRAM in both device-wide or segment-by-segment fashion, data can be transferred to or from the flash memory in device-wide or segment-by-segment fashion, and the flash memory can be addressed and accessed accordingly.

In accordance with one embodiment the illustrated arrangement of data transfer fabric of DMgr **504** also allows the CDC **502** to control data transfer from the Flash memory to the MCH by buffering the data from the Flash **506** using a buffer **602**, and matching the data rate and/or data format of MCH **510**. The buffer **602** is shown in FIG. **6** as a portion of a data format module **604**; however, buffer **602** may also be a distributed buffer such that one buffer is used for each one of the set of multiplexer logic elements shown as multiplexers **611**, **612**, **621**, and **622**. Various buffer arrangements may be used, such as a programmable size buffer to meet the requirement of a given system design requirement, for example the disparity between read/write access time; or overall system performance, for example latency. In certain embodiments, the buffer **604** may introduce one or more clock cycle delays into a data communication path between MCH **510**, DRAM **508**, and Flash **506**.

In certain embodiments, data format module **604** contains a data formatting subsystem (not shown) to enable DMgr **504** to format and perform data transfer in accordance with control information received from CDC**502**. Data buffer **604** of data format module **602**, discussed above, also supports a wide data bus **606** coupled to the Flash memory **506** operating at a first frequency, while receiving data from DRAM **508** using a relatively smaller width data bus **608** operating at a second frequency, the second frequency being larger than the first frequency in certain embodiments. The buffer **602** is designed to match the data flow rate between the DRAM **508** and the Flash **506**.

A register **690** provides the ability to register commands received from MCH **510** via C/A **560** (FIG. **5**A). The register **690** may communicate these commands to CDC **502** and/or to the DRAM **508** and/or Flash **506**. The register **690** communicates these registered commands to CDC **502** for processing. The register **690** may also include multiple registers (not shown), such that it can provide the ability to register multiple commands, a sequence of commands, or provide a pipeline delay stage for buffering and providing a controlled execution of certain commands received form MCH **510**.

In certain embodiments, the register **690** may register commands from MCH **510** and transmit the registered commands to DRAM **508** and/or Flash **506** memory subsystems. In certain embodiments, the CDC **502** monitors commands received from MCH **510**, via control and address bus C/A **560**, and provides appropriate control information to DMgr **504**, DRAM **508**, or Flash **506** to execute these commands and perform data transfer operations between MCH **510** and FDHDIMM **500** via MCH data bus **610**.

FIG. **7** illustrates a functional block diagram of the CDC **502**. In certain embodiments, the major functional blocks of the CDC **502** are a DRAM control block DRAMCtrl **702**, Flash control block FlashCtrl **704**, MCH command interpreter Cmdlnt **706**, DRAM-Flash interface scheduler Scheduler **708**, and DMgr control block (DMgrCtrl) **710**.

In accordance with one embodiment, DRAMCtrl **702** generates DRAM commands that are independent from the commands issued by the MCH **510**. In accordance with one embodiment, when the MCH **510** initiates a read/write operation from/to the same DRAM **508** that is currently executing a command from the DRAMCtrl **702**, then the CDC **502** may choose to instruct DRAMCtrl **702** to abort its operation in order to execute the operation initiated by the MCH. However, the CDC **502** may also pipeline the operation so that it causes DRAMCtrl **702** to either halt or complete its current operation prior to executing that of the MCH. The CDC **502** may also instruct DRAMCtrl **702** to resume its operation once the command from MCH **510** is completed.

In accordance with one embodiment, the FlashCtrl **704** generates appropriate Flash commands for the proper read/write operations. The Cmdlnt **706** intercepts commands received from MCH **510** and generates the appropriate control information and control signals and transmit them to the appropriate FDHDIMM functional block. For example, Cmdlnt **706** issues an interrupt signal to the DRAMCtrl **702** when the MCH issues a command that collides (conflicts) with the currently executing or pending commands that DRAMCtrl **702** has initiated independently from MCH **510**, thus subordinating these commands to those from the MCH. The Scheduler **708** schedules the Flash-DRAM interface operation such that there is no resource conflict in the DMgr **504**. In accordance with one embodiment, the Scheduler **708** assigns time slots for the DRAMCtrl **702** and FlashCtrl **704** operation based on the current status and the pending command received or to be received from the MCH. The DMgrCtrl **710** generates and sends appropriate control information and control signals for the proper operation and control of the data transfer fabric to enable or disable data paths between Flash **506**, DRAM **508**, and the MCH **510**.

FIG. **8**A is a block diagram showing a Flash-DRAM hybrid DIMM (FDHDIMM). As seen from FIG. **8**A, this Flash-DRAM hybrid DIMM requires two separate and independent address buses to separately control the address spaces: one for the Flash memory Flash **506** and the other for the DRAM memory DRAM **508**. The MCH treats the DRAM **508** and Flash **506** as separate memory subsystems, for example DRAM and SSD/HD memory subsystems. The memory in each address space is controlled directly by the MCH. However, the on-DIMM data path between Flash **506** and DRAM **508** allows for direct data transfer to occur between the Flash **506** and the DRAM **508** in response to control information from Ctrl **502**. In this embodiment, this data transfer mechanism provides direct support for executing commands from the MCH without having the MCH directly controlling the data transfer, and thus improving data transfer performance from Flash **506** to the DRAM **508**. However, the MCH needs to manage two address spaces and two different memory protocols simultaneously. Moreover, the MCH needs to map the DRAM memory space into the Flash memory space, and the data interface time suffers due to the difference in the data access time between the Flash memory and the DRAM memory.

In accordance with one embodiment, a memory space mapping of a Flash-DRAM hybrid DIMM is shown in FIG. **8**B. A memory controller of a host system (not shown) controls both of the DRAM **508** address space and the Flash **506** address space using a single unified address space. The CDC **502** receives memory access commands from the MCH and generates control information for appropriate mapping and data transfer between Flash and DRAM memory subsystem to properly carry out the memory access commands. In one embodiment, the memory controller of the host system views the large Flash memory space as a DRAM memory space, and accesses this unified memory

space with a standard DDR (double data rate) protocol used for accessing DRAM. The unified memory space in this case can exhibit overlapping memory address space between the Flash **506** and the DRAM **508**. The overlapping memory address space may be used as a temporary storage or buffer for data transfer between the Flash **506** and the DRAM **508**. For example, the DRAM memory space may hold a copy of data from the selected Flash memory space such that the MCH can access this data normally via DDR memory access commands. The CDC **502** controls the operation of the Flash **506** and DRAM **508** memory subsystems in response to commands received from a memory controller of a host system.

In one embodiment, the unified memory space corresponds to a contiguous address space comprising a first portion of the address space of the Flash **506** and a first portion of the address space of the DRAM **508**. The first portion of the address space of the Flash **506** can be determined via a first programmable register holding a first value corresponding to the desired Flash memory size to be used. Similarly, the first portion of the address space of the DRAM **508** can be determined via a second programmable register holding a second value corresponding to the desired DRAM memory size to be used. In one embodiment, any one of the first portion of the address space of the Flash **506** and the first portion of the address space of the DRAM **508** is determined via a first value corresponding to a desired performance or memory size, the first value being received by the CDC **502** via a command sent by memory controller of the host system.

In accordance with one embodiment, a flow diagram directed to the transfer of data from Flash memory to DRAM memory and vice versa in an exemplary FDHDIMM is shown in FIG. **9**. In certain embodiments, data transfer from the Flash **506** to the DRAM **508** occurs in accordance with memory access commands which the CDC **502** receives from the memory controller of the host system. In certain embodiments, the CDC **502** controls the data transfer from the DRAM **508** to the Flash **506** so as to avoid conflict with any memory operation that is currently being executed. For example, when all the pages in a particular DRAM memory block are closed. The CDC **502** partitions the DRAM memory space into a number of blocks for the purpose of optimally supporting the desired application. The controller can configure memory space in the memory module based on at least one of one or more commands received from the MCH, instructions received from the MCH, a programmable value written into a register, a value corresponding to a first portion of the volatile memory subsystem, a value corresponding to a first portion of the non-volatile memory subsystem, and a timing value. Furthermore, the block size can be configurable by the memory controller of the host system, such that the number pages in a block can be optimized to support a particular application or a task. Furthermore, the block size may be configured on-the-fly, e.g. CDC **502** can receive instruction regarding a desired block size from the memory controller via a memory command, or via a programmable value.

In certain embodiments, a memory controller can access the memory module using a standard access protocol, such as JEDEC's DDR DRAM, by sending a memory access command to the CDC **502** which in turn determines what type of a data transfer operation it is and the corresponding target address where the data information is stored, e.g. data information is stored in the DRAM **508** or Flash **506** memory subsystems. In response to a read operation, if the CDC **502** determines that data information, e.g. a page (or block), does not reside in the DRAM **508** but resides in Flash **506**, then the CDC **502** initiates and controls all necessary data transfer operations from Flash **506** to DRAM **508** and subsequently to the memory controller. In one embodiment, once the CDC **502** completes the data transfer operation of the requested data information from the Flash **506** to the DRAM **508**, the CDC **502** alerts the memory controller to retrieve the data information from the DRAM **508**. In on embodiment, the memory controller initiates the copying of data information from Flash **506** to DRAM **508** by writing, into a register in the CDC **502**, the target Flash address along with a valid block size. The CDC **502** in turn, executes appropriate operations and generates control information to copy the data information to the DRAM **508**. Consequently, the memory controller can access or retrieve the data information using standard memory access commands or protocol.

An exemplary flow chart is shown in FIG. **9**, a starting step or power up **902**, is followed by an initialization step **904**, the memory controller initiates, at step **906**, a data move from the Flash **506** to the DRAM **508** by writing target address and size, to a control register in the CDC **502**, which then copies, at **908**, data information from the Flash **506** to the DRAM **508** and erases the block in the Flash. Erasing the data information from Flash may be accomplished independently from (or concurrently with) other steps that CDC **502** performs in this flow chart, i.e. other steps can be executed concurrently with the Erase the Flash block step. Once the data information or a block of data information is thus moved to the DRAM **508**, the memory controller can operate on this data block using standard memory access protocol or commands at **910**. The CDC **502** checks, at **912**, if any of the DRAM **508** blocks, or copied blocks, are closed. If the memory controller closed any open blocks in DRAM **508**, then the CDC **502** initiate a Flash write to write the closed block from the DRAM **508** to the Flash **506**, at **914**. In addition, the memory controller, at **916**, reopens the closed block that is currently being written into the Flash **506**, then the CDC **502** stops the Flash write operation and erases the Flash block which was being written to, as shown at **918**. Otherwise, the CDC **502** continues and completes the writing operation to the Flash at **920**.

The dashed lines in FIG. **9** indicate independent or parallel activities that can be performed by the CDC **502**. At any time the CDC **502** receives a DRAM load command from a memory controller which writes a Flash target address and/or block size information into the RC register(s) at **922**, as described above, then the CDC **502** executes a load DRAM w/RC step **906** and initiates another branch (or a thread) of activities that includes steps **908**-**922**. In one embodiment, the CDC **502** controls the data transfer operations between DRAM **508** and Flash **506** such that the Flash **506** is completely hidden from the memory controller. The CDC **502** monitors all memory access commands sent by the memory controller using standard DRAM protocol and appropriately configures and manipulate both Flash **506** and DRAM **508** memory subsystems to perform the requested memory access operation and thus achieve the desired results. The memory controller does not interface directly with the Flash memory subsystem. Instead, the memory controller interfaces with the CDC **502** and/or DMgr **504** as shown in FIG. **5** and FIG. **6**. Moreover, the memory controller may use one or more protocol, such as DDR, DDR2, DDR3, DDR4 protocols or the like.

In accordance with one embodiment, an example of mapping a DRAM address space to Flash memory address space is shown in FIG. **10**. Two sets (**1002**, **1004**) of address

bits AD**6** to AD**17**, forming a 24 bit extended memory page address, are allocated for the block address. For example, assuming a Block size of 256K Bytes, then a 24-bit block address space (using the two sets of AD**6** to AD**17 1002** and **1004**) would enable access to 4 TB of Flash memory storage space. If a memory module has 1 GB of DRAM storage capacity, then it can hold approximately 4K Blocks of data in the DRAM memory, each Block comprise 256 K Bytes of data. The DRAM address space, corresponding to the 4K blocks, can be assigned to different virtual ranks and banks, where the number of virtual ranks and banks is configurable and can be manipulated to meet a specific design or performance needs. For example, if a 1G Bytes memory module is configured to comprise two ranks with eight banks per rank, then each bank would hold two hundred fifty (**250**) blocks or the equivalent of 62 M Bytes or 62K pages, where each page correspond to a 1K Bytes. Other configurations using different page, block, banks, or ranks numbers may also be used. Furthermore, an exemplary mapping of 24-bit DDR DIMM block address to Flash memory address, using Block addressing as described above, is shown in FIG. **10**. The 24-bit can be decomposed into fields, such as a logical unit number LUN address **1061** field, a Block address **1051** field, a Plane address **1041**, a Page address **1031**, and a group of least significant address bits A_0A_1 **1021**. The Plane address **1041** is a sub address of the block address, and it may be used to support multiple page IO so as to improve Flash memory subsystem operation. In this example, it is understood that different number of bits may be allocated to each field of the 24-bit

The CDC **502** manages the block write-back operation by queuing the blocks that are ready to be written back to the Flash memory. As described above, if any page in a queued block for a write operation is reopened, then the CDC **502** will stop the queued block write operation, and remove the block from the queue. Once all the pages in a block are closed, then the CDC **502** restarts the write-back operation and queue the block for a write operation.

In accordance with one embodiment, an exemplary read operation from Flash **506** to DRAM **508** can be performed in approximately 400 μs, while a write operation from DRAM **508** to Flash **506** can be performed in approximately 22 ms resulting in a read to write ratio of 55 to 1. Therefore, if the average time a host system's memory controller spends accessing data information in a Block of DRAM is about 22 ms (that is the duration that a Block comprises one or more pages that are open), then the block write-back operation from DRAM to Flash would not impact performance and hence the disparity between read and write access may be completely hidden from the memory controller. If the block usage time is 11 ms instead of 22 ms, then the CDC **502** control the data transfer operation between DRAM **508** and Flash **506** such that there are no more than 9 closed blocks in the queue to be written-back to the Flash memory, hence approximately an average of 100 ms can be maintained for a standard DDR DRAM operation. Moreover, the number of closed Blocks in the queue to be written-back to the Flash memory subsystem varies with the average block usage time and the desired performance for a specific host system or for a specific application running using the host system resources.

Consequently, the maximum number of closed Blocks to be written-back to Flash can be approximated to be

((# of blocks per bank)/(ratio of 'Flash_block_write-_time' to 'Flash_read_time'))*((Block usage time)/('Flash_block_write_time'))

In order to maintain less than 100 ms time period for queued write-back Blocks, then using a Flash memory subsystem having 22 ms write access time per Block would results in a maximum number of four Blocks to be queued for write operation to Flash **506**. Therefore, on average approximately 88 ms (=22 ms*4) for blocks means that each bank should not have more than four Blocks that need to be written back to the Flash **506**.

The above equation also indicates that bigger DRAM memory space can support shorter block usage times. For example, 2 GB of DRAM memory allows the 8 closed blocks to be written-back to Flash. The table in FIG. **11** provides an estimation of the maximum allowed closed blocks in the queue to be written back to the Flash memory for different DRAM density using various average block use time.

Certain embodiments described herein include a memory system which can communicate with a host system such as a disk controller of a computer system. The memory system can include volatile and non-volatile memory, and a controller. The controller backs up the volatile memory using the non-volatile memory in the event of a trigger condition. Trigger conditions can include, for example, a power failure, power reduction, request by the host system, etc. In order to power the system in the event of a power failure or reduction, the memory system can include a secondary power source which does not comprise a battery and may include, for example, a capacitor or capacitor array.

In certain embodiments, the memory system can be configured such that the operation of the volatile memory is not adversely affected by the non-volatile memory or by th controller when the volatile memory is interacting with the host system. For example, one or more isolation devices may isolate the non-volatile memory and the controller from the volatile memory when the volatile memory is interacting with the host system and may allow communication between the volatile memory and the non-volatile memory when the data of the volatile memory is being restored or backed-up. This configuration generally protects the operation of the volatile memory when isolated while providing backup and restore capability in the event of a trigger condition, such as a power failure.

In certain embodiments described herein, the memory system includes a power module which provides power to the various components of the memory system from different sources based on a state of the memory system in relation to a trigger condition (e.g., a power failure). The power module may switch the source of the power to the various components in order to efficiently provide power in the event of the power failure. For example, when no power failure is detected, the power module may provide power to certain components, such as the volatile memory, from system power while charging a secondary power source (e.g., a capacitor array). In the event of a power failure or other trigger condition, the power module may power the volatile memory elements using the previously charged secondary power source.

In certain embodiments, the power module. transitions relatively smoothly from powering the volatile memory with system power to powering it with the secondary power source. For example, the power system may power volatile memory with a third power source from the time the memory system detects that power failure is likely to occur until the time the memory system detects that the power failure has actually occurred.

In certain embodiments, the volatile memory system can be operated at a reduced frequency during backup and/or

restore operations which can improve the efficiency of the system and save power. In some embodiments, during backup and/or restore operations, the volatile memory communicates with the non-volatile memory by writing and/or. reading data words in bit-wise slices instead of by writing entire words at once. In certain embodiments, when each slice is being written to or read from the volatile memory the unused slice(s) of volatile memory is not active, which can reduce the power consumption of the system.

In yet other embodiments, the non-volatile memory can include at least 100 percent more storage capacity than the volatile memory. This configuration can allow the memory system to efficiently handle subsequent trigger conditions.

FIG. **12** is a block diagram of an example memory system **1010** compatible with certain embodiments described herein. The memory system **1010** can be coupled to a host computer system and can include a volatile memory subsystem **1030**, a non-volatile memory subsystem **1040**, and a controller **1062** operatively coupled to the non-volatile memory subsystem **1040**. In certain embodiments, the memory system **1010** includes at least one circuit **1052** configured to selectively operatively decouple the controller **1062** from the volatile memory subsystem **1030**.

In certain embodiments, the memory system **1010** comprises a memory module. The memory system **1010** may comprise a printed-circuit board (PCB) **1020**. In certain embodiments, the memory system **1010** has a memory capacity of 512-MB, 1-GB, 2-GB, 4-GB, or 8-GB. Other volatile memory capacities are also compatible with certain embodiments described herein. In certain embodiments, the memory system **10** has a non-volatile memory capacity of 512-MB, 1-GB, 2-GB, 4-GB, 8-GB, 16-GB, or 32-GB. Other non-volatile memory capacities are also compatible with certain embodiments described herein. In addition, memory systems **1010** having widths of 4 bytes, 8 bytes, 16 bytes, 32 bytes, or 32 bits, 64 bits, 128 bits, 256 bits, as well as other widths (in bytes or in bits), are compatible with embodiments described herein. In certain embodiments, the PCB **1020** has an industry-standard form factor. For example, the PCB **1020** can have a low profile (LP) form factor with a height of 30 millimeters and a width of 133.35 millimeters. In certain other embodiments, the PCB **1020** has a very high profile (VHP) form factor with a height of 50 millimeters or more. In certain other embodiments, the PCB **1020** has a very low profile (VLP) form factor with a height of 18.3 millimeters. Other form factors including, but not limited to, small-outline (SO-DIMM), unbuffered (UDIMM), registered (RDIMM), fully-buffered (FB-DIMM), miniDIMM, mini-RDIMM, VLP mini-DIMM, micro-DIMM, and SRAM DIMM are also compatible with certain embodiments described herein. For example, in other embodiments, certain non-DIMM form factors are possible such as, for example, single in-line memory module (SIMM), multi-media card (MMC), and small computer system interface (SCSI).

In certain preferred embodiments, the memory system **1010** is in electrical communication with the host system. In other embodiments, the memory system **1010** may communicate with a host system using some other type of communication, such as, for example, optical communication. Examples of host systems include, but are not limited to, blade servers, 1U servers, personal computers (PCs), and other applications in which space is constrained or limited. The memory system **1010** can be in communication with a disk controller of a computer system, for example. The PCB **1020** can comprise an interface **1022** that is configured to be in electrical communication with the host system (not shown). For example, the interface **1022** can comprise a plurality of edge connections which fit into a corresponding slot connector of the host system. The interface **1022** of certain embodiments provides a conduit for power voltage as well as data, address, and control signals between the memory system **1010** and the host system. For example, the interface **1022** can comprise a standard 240-pin DDR2 edge connector.

The volatile memory subsystem **1030** comprises a plurality of volatile memory elements **1032** and the non-volatile memory subsystem **1040** comprises a plurality of non-volatile memory elements **1042**. Certain embodiments described herein advantageously provide nonvolatile storage via the non-volatile memory subsystem **1040** in addition to high-performance (e.g., high speed) storage via the volatile memory subsystem **1030**. In certain embodiments, the first plurality of volatile memory elements **1032** comprises two or more dynamic random-access memory (DRAM) elements. Types of DRAM elements **1032** compatible with certain embodiments described herein include, but are not limited to, DDR, DDR2, DDR3, and synchronous DRAM (SDRAM). For example, in the block diagram of FIG. **12**, the first memory bank **1030** comprises eight 64M×8 DDR2 SDRAM elements **1032**. The volatile memory elements **1032** may comprise other types of memory elements such as static random-access memory (SRAM). In addition, volatile memory elements **1032** having bit widths of 4, 8, 16, 32, as well as other bit widths, are compatible with certain embodiments described herein. Volatile memory elements **1032** compatible with certain embodiments described herein have packaging which include, but are not limited to, thin small-outline package (TSOP), ball-grid-array (BGA), fine-pitch BGA (FBOA), micro-BOA (1.1,BGA), mini-BGA (mBGA), and chip-scale packaging (CSP).

In certain embodiments, the second plurality of non-volatile memory elements **1042** comprises one or more flash memory elements. Types of flash memory elements **1042** compatible with certain embodiments described herein include, but are not limited to, NOR flash, NAND flash, ONE-NAND flash, and multi-level cell (MLC). For example, in the block diagram of FIG. **12**, the second memory bank **1040** comprises 512 MB of flash memory organized as four 128 Mb×8 NAND flash memory elements **1042**. In addition, nonvolatile memory elements **1042** having bit widths of 4, 8, 16, 32, as well as other bit widths, are compatible with certain embodiments described herein. Non-volatile memory elements **1042** compatible with certain embodiments described herein have packaging which include, but are not limited to, thin small-outline package (TSOP), ball-grid-array (BOA), fine-pitch BOA (FBGA), micro-BOA (POA), mini-BGA (mBGA), and chip-scale packaging (CSP).

FIG. **13** is a block diagram of an example memory module **10** with ECC (error-correcting code) having a volatile memory subsystem **1030** with nine volatile memory elements **1032** and a non-volatile memory subsystem **1040** with five non-volatile memory elements **1042** in accordance with certain embodiments described herein. The additional memory element **1032** of the first memory bank **1030** and the additional memory element **1042** of the second memory bank **1040** provide the ECC capability. In certain other embodiments, the volatile memory subsystem **1030** comprises other numbers of volatile memory elements **1032** (e.g., 2, 3, 4, 5, 6, 7, more than 9). In certain embodiments, the non-volatile memory subsystem **1040** comprises other numbers of nonvolatile memory elements **1042** (e.g., 2, 3, more than 5).

Referring to FIG. **12**, in certain embodiments, the logic element **1070** comprises a field-programmable gate array (FPGA). In certain embodiments, the logic element **1070** comprises an FPGA available from Lattice Semiconductor Corporation which includes an internal flash. In certain other embodiments, the logic element **1070** comprises an FPOA available from another vendor. The internal flash can improve the speed of the memory system **1010** and save physical space. Other types of logic elements **1070** compatible with certain embodiments described herein include, but are not limited to, a programmable-logic device (PLD), an application-specific integrated circuit (ASIC), a custom-designed semiconductor device, a complex programmable logic device (CPLD). In certain embodiments, the logic element **1070** is a custom device. In certain embodiments, the logic element **1070** comprises various discrete electrical elements, while in certain other embodiments, the logic element **1070** comprises one or more integrated circuits. FIG. **14** is a block diagram of an example memory module **1010** having a microcontroller unit **1060** and logic element **1070** integrated into a single controller **1062** in accordance with certain embodiments described herein. In certain embodiments, the controller **1062** includes one or more other components. For example, in one embodiment, an FPGA without an internal flash is used and the controller **1062** includes a separate flash memory component which stores configuration information to program the FPGA.

In certain embodiments, the at least one circuit **1052** comprises one or more switches coupled to the volatile memory subsystem **1030**, to the controller **1062**, and to the host computer (e.g., via the interface **1022**, as schematically illustrated by FIGS. **12-14**). The one or more switches are responsive to signals (e.g., from the controller **1062**) to selectively operatively decouple the controller **1062** from the volatile memory subsystem **1030** and to selectively operatively couple the controller **1062** to the volatile memory subsystem **1030**. In addition, in certain embodiments, the at least one circuit **1052** selectively operatively couples and decouples the volatile memory subsystem **1030** and the host system.

In certain embodiments, the volatile memory subsystem **1030** can comprise a registered DIMM subsystem comprising one or more registers **1160** and a plurality of DRAM elements **1180**, as schematically illustrated by FIG. **15**A. In certain such embodiments, the at least one circuit **1052** can comprise one or more switches **1172** coupled to the controller **1062** (e.g., logic element **1070**) and to the volatile memory subsystem **1030** which can be actuated to couple and decouple the controller **1062** to and from the volatile memory subsystem **1030**, respectively. The memory system **1010** further comprises one or more switches **1170** coupled to the one or more registers **1160** and to the plurality of DRAM elements **1180** as schematically illustrated by FIG. **15**A. The one or more switches **1170** can be selectively switched, thereby selectively operatively coupling the volatile memory subsystem **1030** to the host system **1150**. In certain other embodiments, as schematically illustrated by FIG. **15**B, the one or more switches **1174** are also coupled to the one or more registers **1160** and to a power source **1162** for the one or more registers **1160**. The one or more switches **1174** can be selectively switched to turn power on or off to the one or more registers **1160**, thereby selectively operatively coupling the volatile memory subsystem **1030** to the host system **1150**. As schematically illustrated by FIG. **15**C, in certain embodiments the at least one circuit **1052** comprises a dynamic on-die termination (ODT) **1176** circuit of the logic element **1070**. For example, the logic element **1070** can comprise a dynamic ODT circuit **1176** which selectively operatively couples and decouples the logic element **1070** to and from the volatile memory subsystem **1030**, respectively. In addition, and similar to the example embodiment of FIG. **15**A described above, the one or more switches **1170** can be selectively switched, thereby selectively operatively coupling the volatile memory subsystem **1030** to the host system **1150**.

Certain embodiments described herein utilize the non-volatile memory subsystem **1040** as a flash "mirror" to provide backup of the volatile memory subsystem **1030** in the event of certain system conditions. For example, the non-volatile memory subsystem **1040** may backup the volatile memory subsystem **1030** in the event of a trigger condition, such as, for example, a power failure or power reduction or a request from the host system. In one embodiment, the nonvolatile memory subsystem **1040** holds intermediate data results in a noisy system environment when the host computer system is engaged in a long computation. In certain embodiments, a backup may be performed on a regular basis. For example, in one embodiment, the backup may occur every millisecond in response to a trigger condition. In certain embodiments, the trigger condition occurs when the memory system **1010** detects that the system voltage is below a certain threshold voltage. For example, in one embodiment, the threshold voltage is **1010** percent below a specified operating voltage. In certain embodiments, a trigger condition occurs when the voltage goes above a certain threshold value, such as, for example, **1010** percent above a specified operating voltage. In some embodiments, a trigger condition occurs when the voltage goes below a threshold or above another threshold. In various embodiments, a backup and/or restore operation may occur in reboot and/or non-reboot trigger conditions.

As schematically illustrated by FIGS. **12** and **13**, in certain embodiments, the controller **1062** may comprise a microcontroller unit (MCU) **1060** and a logic element **1070**. In certain embodiments, the MCU **1060** provides memory management for the non-volatile memory subsystem **1040** and controls data transfer between the volatile memory subsystem **30** and the nonvolatile memory subsystem **1040**. The MCU **1060** of certain embodiments comprises a 16-bit microcontroller, although other types of microcontrollers are also compatible with certain embodiments described herein. As schematically illustrated by FIGS. **12** and **13**, the logic element **1070** of certain embodiments is in electrical communication with the non-volatile memory subsystem **1040** and the MCU **1060**. The logic element **1070** can provide signal level translation between the volatile memory elements **1032** (e.g., 1.8V SSTL-2 for DDR2 SDRAM elements) and the non-volatile memory elements **1042** (e.g., 3V TTL for NAND flash memory elements). In certain embodiments, the logic element **1070** is also programmed to perform address/address translation between the volatile memory subsystem **1030** and the non-volatile memory subsystem **1040**. In certain preferred embodiments, 1-NAND type flash are used for the non-volatile memory elements **1042** because of their superior read speed and compact structure.

The memory system **1010** of certain embodiments is configured to be operated in at least two states. The at least two states can comprise a first state in which the controller **1062** and the non-volatile memory subsystem **1040** are operatively decoupled (e.g., isolated) from the volatile memory subsystem **1030** by the at least one circuit **1052** and a second state in which the volatile memory subsystem **1030** is operatively coupled to the controller **1062** to allow data to

be communicated between the volatile memory subsystem **1030** and the nonvolatile memory subsystem **1040** via the controller **1062**. The memory system **1010** may transition from the first state to the second state in response to a trigger condition, such as when the memory system **1010** detects that there is a power interruption (e.g., power failure or reduction) or a system hang-up.

The memory system **1010** may further comprise a voltage monitor **1050**. The voltage monitor circuit **1050** monitors the voltage supplied by the host system via the interface **1022**. Upon detecting a low voltage condition (e.g., due to a power interruption to the host system), the voltage monitor circuit **1050** may transmit a signal to the controller **1062** indicative of the detected condition. The controller **1062** of certain embodiments responds to the signal from the voltage monitor circuit **1050** by transmitting a signal to the at least one circuit **1052** to operatively couple the controller to the volatile memory system **1030**, such that the memory system **1010** enters the second state. For example, the voltage monitor **1050** may send a signal to the MCU **1060** which responds by accessing the data on the volatile memory system **1030** and by executing a write cycle on the nonvolatile memory subsystem **1040**. During this write cycle, data is read from the volatile memory subsystem **1030** and is transferred to the non-volatile memory subsystem **1040** via the MCU **1060**. In certain embodiments, the voltage monitor circuit **1050** is part of the controller **1062** (e.g., part of the MCU **1060**) and the voltage monitor circuit **1050** transmits a signal to the other portions of the controller **1062** upon detecting a power threshold condition.

The isolation or operational decoupling of the volatile memory subsystem **1030** from the non-volatile memory subsystem in the first state can preserve the integrity of the operation of the memory system **1010** during periods of operation in which signals (e.g., data) are transmitted between the host system and the volatile memory subsystem **1030**. For example, in one embodiment during such periods of operation, the controller **1062** and the nonvolatile memory subsystem **1040** do not add a significant capacitive load to the volatile memory system **1030** when the memory system **1010** is in the first state. In certain such embodiments, the capacitive load of the controller **1062** and the non-volatile memory subsystem **1040** do not significantly affect the signals propagating between the volatile memory subsystem **1030** and the host system. This can be particularly advantageous in relatively high-speed memory systems where loading effects can be significant. In one preferred embodiment, the at least one circuit **1052** comprises an FSA1208 Low-Power, Eight-Port, Hi-Speed Isolation Switch from Fairchild Semiconductor. In other embodiments, the at least one circuit **1052** comprises other types of isolation devices.

Power may be supplied to the volatile memory subsystem **1030** from a first power supply (e.g., a system power supply) when the memory system **1010** is in the first state and from a second power supply **1080** when the memory system **1010** is in the second state. In certain embodiments, the memory system **1010** is in the first state when no trigger condition (e.g., a power failure) is present and the memory system **1010** enters the second state in response to a trigger condition. In certain embodiments, the memory system **1010** has a third state in which the controller **1062** is operatively decoupled from the volatile memory subsystem **1030** and power is supplied to the volatile memory subsystem **1030** from a third power supply (not shown). For example, in one embodiment the third power supply may provide power to the volatile memory subsystem **1030** when the memory system **1010** detects that a trigger condition is likely to occur but has not yet occurred.

In certain embodiments, the second power supply **1080** does not comprise a battery. Because a battery is not used, the second power supply **1080** of certain embodiments may be relatively easy to maintain, does not generally need to be replaced, and is relatively environmentally friendly. In certain embodiments, as schematically illustrated by FIGS. **12-14**, the second power supply **1080** comprises a step-up transformer **1082**, a step-down transformer **1084**, and a capacitor bank **1086** comprising one or more capacitors (e.g., double-layer capacitors). In one example embodiment, capacitors may take about three to four minutes to charge and about two minutes to discharge. In other embodiments, the one or more capacitors may take a longer time or a shorter time to charge and/or discharge. For example, in certain embodiments, the second power supply **1080** is configured to power the volatile memory subsystem **1030** for less than thirty minutes. In certain embodiments, the second power supply **1080** may comprise a battery. For example, in certain embodiments, the second power supply **1080** comprises a battery and one or more capacitors and is configured to power the volatile memory subsystem **1030** for no more than thirty minutes.

In certain embodiments, the capacitor bank **1086** of the second power supply **1080** is charged by the first power supply while the memory system **1010** is in the first state. As a result, the second power supply **1080** is fully charged when the memory system **1010** enters the second state. The memory system **1010** and the second power supply **1080** may be located on the same printed circuit board **1020**. In other embodiments, the second power supply **1080** may not be on the same printed circuit board **1020** and may be tethered to the printed circuit board **1020**, for example.

When operating in the first state, in certain embodiments, the step-up transformer **1082** keeps the capacitor bank **1086** charged at a peak value. In certain embodiments, the step-down transformer **1084** acts as a voltage regulator to ensure that regulated voltages are supplied to the memory elements (e.g., 1.8V to the volatile DRAM elements **1032** and 3.0V to the non-volatile flash memory elements **1042**) when operating in the second state (e.g., during power down). In certain embodiments, as schematically illustrated by FIGS. **12-14**, the memory module **1010** further comprises a switch **1090** (e.g., FET switch) that switches power provided to the controller **1062**, the volatile memory subsystem **1030**, and the non-volatile memory subsystem **1040**, between the power from the second power supply **1080** and the power from the first power supply (e.g., system power) received via the interface **1022**. For example, the switch **1090** may switch from the first power supply to the second power supply **1080** when the voltage monitor **1050** detects a low voltage condition. The switch **1090** of certain embodiments advantageously ensures that the volatile memory elements **1032** and non-volatile memory elements **1042** are powered long enough for the data to be transferred from the volatile memory elements **1032** and stored in the non-volatile memory elements **1042**. In certain embodiments, after the data transfer is complete, the switch **1090** then switches back to the first power supply and the controller **1062** transmits a signal to the at least one circuit **1052** to operatively decouple the controller **1062** from the volatile memory subsystem **1030**, such that the memory system **1010** reenters the first state.

When the memory system **1010** re-enters the first state, data may be transferred back from the non-volatile memory

subsystem **1040** to the volatile memory subsystem **1030** via the controller **1062**. The host system can then resume accessing the volatile memory subsystem **1030** of the memory module **1010**. In certain embodiments, after the memory system **1010** enters or re-enters the first state (e.g., after power is restored), the host system accesses the volatile memory subsystem **1030** rather than the non-volatile memory subsystem **1040** because the volatile memory elements **1032** have superior read/write characteristics. In certain embodiments, the transfer of data from the volatile memory bank **1030** to the nonvolatile memory bank **1040**, or from the non-volatile memory bank **1040** to the volatile. memory bank **1030**, takes less than one minute per GB.

In certain embodiments, the memory system **1010** protects the operation of the volatile memory when communicating with the host-system and provides backup and restore capability in the event of a trigger condition such as a power failure. In certain embodiments, the memory system **1010** copies the entire contents of the volatile memory subsystem **1030** into the nonvolatile memory subsystem **1040** on each backup operation. Moreover, in certain embodiments, the entire contents of the non-volatile memory subsystem **1040** are copied back into the volatile memory subsystem **1030** on each restore operation. In certain embodiments, the entire contents of the non-volatile memory subsystem **1040** are accessed for each backup and/or restore operation, such that the non-volatile memory subsystem **1040** (e.g., flash memory subsystem) is used generally uniformly across its memory space and wear-leveling is not performed by the memory system **1010**. In certain embodiments, avoiding wear-leveling can decrease cost and complexity of the memory system **1010** and can improve the performance of the memory system **1010**. In certain other embodiments, the entire contents of the volatile memory subsystem **1030** are not copied into the non-volatile memory subsystem **1040** on each backup operation, but only a partial copy is performed. In certain embodiments, other management capabilities such as bad-block management and error management for the flash memory elements of the non-volatile memory subsystem **1040** are performed in the controller **1062**.

The memory system **1010** generally operates as a write-back cache in certain embodiments. For example, in one embodiment, the host system (e.g., a disk controller) writes data to the volatile memory subsystem **1030** which then writes the data to non-volatile storage which is not part of the memory system **1010**, such as, for example, a hard disk. The disk controller may wait for an acknowledgment signal from the memory system **1010** indicating that the data has been written to the hard disk or is otherwise secure. The memory system **1010** of certain embodiments can decrease delays in the system operation by indicating that the data has been written to the hard disk before it has actually done so. In certain embodiments, the memory system **1010** will still be able to recover the data efficiently in the event of a power outage because of the backup and restore capabilities described herein. In certain other embodiments, the memory system **1010** may be operated as a write-through cache or as some other type of cache.

FIG. **16** schematically illustrates an example power module **1100** of the memory system **1010** in accordance with certain embodiments described herein. The power module **1100** provides power to the various components of the memory system **1010** using different elements based on a state of the memory system **1010** in relation to a trigger condition. In certain embodiments, the power module **1100** comprises one or more of the components described above with respect to FIG. **12**. For example, in certain embodiments, the power module **1100** includes the second power supply **1080** and the switch **1090**.

The power module **1100** provides a plurality of voltages to the memory system **1010** comprising non-volatile and volatile memory subsystems **1030**, **1040**. The plurality of voltages comprises at least a first voltage **1102** and a second voltage **1104**. The power module **1100** comprises an input **1106** providing a third voltage **1108** to the power module **1100** and a voltage conversion element **1120** configured to provide the second voltage **1104** to the memory system **1010**. The power module **1100** further comprises a first power element **1130** configured to selectively provide a fourth voltage **1110** to the conversion element **1120**. In certain embodiments, the first power element **1130** comprises a pulse-width modulation power controller. For example, in one example embodiment, the first power element **1130** is configured to receive a 1.8V input system voltage as the third voltage **1108** and to output a modulated 5V output as the fourth voltage **1110**.

The power module **1100** further comprises a second power element **1140** can be configured to selectively provide a fifth voltage **1112** to the conversion element **1120**. The power module **1100** can be configured to selectively provide the first voltage **1102** to the memory system **1010** either from the conversion element **1120** or from the input **1106**.

The power module **1100** can be configured to be operated in at least three states in certain embodiments. In a first state, the first voltage **1102** is provided to the memory system **1010** from the input **1106** and the fourth voltage **1110** is provided to the conversion element **1120** from the first power element **1130**. In a second state, the fourth voltage **1110** is provided to the conversion element **1120** from the first power element **1130** and the first voltage **1102** is provided to the memory system **1010** from the conversion element **1120**. In the third state, the fifth voltage **1112** is provided to the conversion element **1120** from the second power element **1140** and the first voltage **1104** is provided to the memory system **1010** from the conversion element **1120**.

In certain embodiments, the power module **1100** transitions from the first state to the second state upon detecting that a trigger condition is likely to occur and transitions from the second state to the third state upon detecting that the trigger condition has occurred. For example, the power module **1100** may transition to the second state when it detects that a power failure is about to occur and transitions to the third state when it detects that the power failure has occurred. In certain embodiments, providing the first voltage **1102** in the second state from the first power element **1130** rather than from the input **1106** allows a smoother transition from the first state to the third state. For example, in certain embodiments, providing the first voltage **1102** from the first power element **1130** has capacitive and other smoothing effects. In addition, switching the point of power transition to be between the conversion element **1120** and the first and second power elements **1130**, **1140** (e.g., the sources of the pre-regulated fourth voltage **1110** in the second state and the pre-regulated fifth voltage **1112** in the third state) can smooth out potential voltage spikes.

In certain embodiments, the second power element **1140** does not comprise a battery and may comprise one or more capacitors. For example, as schematically illustrated in FIG. **16**, the second power element **1140** comprises a capacitor array **1142**, a buck-boost converter **1144** which adjusts the voltage for charging the capacitor array and a voltage/current limiter **1146** which limits the charge current to the capacitor array **1142** and stops charging the capacitor array **1142** when it has reached a certain charge voltage. In one

example embodiment, the capacitor array **1142** comprises two 50 farad capacitors capable of holding a total charge of 4.6V. For example, in one example embodiment, the buck-boost converter **1144** receives a 1.8V system voltage (first voltage **1108**) and boosts the voltage to 4.3V which is outputted to the voltage current limiter **1146**. The voltage/current limiter **1146** limits the current going to the capacitor array **1142** to 1A and stops charging the array **1142** when it is charged to 4.3V. Although described with respect to certain example embodiments, one of ordinary skill will recognize from the disclosure herein that the second power element **1140** may include alternative embodiments. For example, different components and/or different value components may be used. For example, in other embodiments, a pure boost converter may be used instead of a buck-boost converter. In another embodiment, only one capacitor may be used instead of a capacitor array **1142**.

The conversion element **1120** can comprise one or more buck converters and/or one or more buck-boost converters. The conversion element **1120** may comprise a plurality of sub-blocks **1122**, **1124**, **1126** as schematically illustrated by FIG. **16**, which can provide more voltages in addition to the second voltage **1104** to the memory system **1010**. The sub-blocks may comprise various converter circuits such as buck-converters, boost converters, and buck-boost converter circuits for providing various voltage values to the memory system **1010**. For example, in one embodiment, sub-block **1122** comprises a buck converter, sub-block **1124** comprises a dual buck converter, and sub-block **1126** comprises a buck-boost converter as schematically illustrated by FIG. **16**. Various other components for the sub-blocks **1122**, **1124**, **1126** of the conversion element **1120** are also compatible with certain embodiments described herein. In certain embodiments, the conversion element **1120** receives as input either the fourth voltage **1110** from the first power element **1130** or the fifth voltage **1112** from the second power element **1140**, depending on the state of the power module **1100**, and reduces the input to an appropriate amount for powering various components of the memory system. For example, the buck-converter of sub-block **1122** can provide 1.8V at 2A for about 60 seconds to the volatile memory elements **1032** (e.g., DRAM), the non-volatile memory elements **1042** (e.g., flash), and the controller **1062** (e.g., an FPGA) in one embodiment. The sub-block **1124** can provide the second voltage **1104** as well as another reduced voltage **1105** to the memory system **1010**. In one example embodiment, the second voltage **1104** is 2.5V and is used to power the at least one circuit **1052** (e.g., isolation device) and the other reduced voltage **1105** is 1.2V and is used to power the controller **1062** (e.g., FPGA). The subblock **1126** can provide yet another voltage **1107** to the memory system **1010**. For example, the voltage **1107** may be 3.3V and may be used to power both the controller **1062** and the at least one circuit **1052**.

Although described with respect to certain example embodiments, one of ordinary skill will recognize from the disclosure herein that the conversion element **1120** may include alternative embodiments. For example, there may be more or less sub-blocks which may comprise other types of converters (e.g., pure boost converters) or which may produce different voltage values. In one embodiment, the volatile memory elements **1032** and nonvolatile memory elements **1042** are powered using independent voltages and are not both powered using the first voltage **1102**.

FIG. **17** is a flowchart of an example method **1200** of providing a first voltage **1102** and a second voltage **1104** to a memory system **1010** including volatile and nonvolatile memory subsystems **1030**, **1040**. While the method **1200** is described herein by reference to the memory system **1010** schematically illustrated by FIGS. **12-15**, other memory systems are also compatible with embodiments of the method **1200**. During a first condition, the method **1200** comprises providing the first voltage **1102** to the memory system **1010** from an input power supply **1106** and providing the second voltage **1104** to the memory system **1010** from a first power subsystem in operational block **1210**. For example, in one embodiment, the first power subsystem comprises the first power element **1130** and the voltage conversion element **1120** described above with respect to FIG. **16**. In other embodiments, other first power subsystems are used.

The method **1200** further comprises detecting a second condition in operational block **1220**. In certain embodiments, detecting the second condition comprises detecting that a trigger condition is likely to occur. During the second condition, the method **1200** comprises providing the first voltage **1102** and the second voltage **1104** to the memory system **1010** from the first power subsystem in an operational block **1230**. For example, referring to FIG. **16**, a switch **1148** can be toggled to provide the first voltage **1102** from the conversion element **1120** rather than from the input power supply.

The method **1200** further comprises charging a second power subsystem in operational block **1240**. In certain embodiments, the second power subsystem comprises the second power element **1140** or another power supply that does not comprise a battery. For example, in one embodiment, the second power subsystem comprises the second power element **1140** and the voltage conversion element **1120** described above with respect to FIG. **16**. In other embodiments, some other second power subsystem is used.

The method **1200** further comprises detecting a third condition in an operational block **1250** and during the third condition, providing the first voltage **1102** and the second voltage **1104** to the memory system **1010** from the second power subsystem **1140** in an operational block **1260**. In certain embodiments, detecting the third condition comprises detecting that the trigger condition has occurred. The trigger condition may comprise various conditions described herein. In various embodiments, for example, the trigger condition comprises a power reduction, power failure, or system hang-up. The operational blocks of the method **1200** may be performed in different orders in various embodiments. For example, in certain embodiments, the second power subsystem **1140** is charged before detecting the second condition.

In certain embodiments, the memory system **1010** comprises a volatile memory subsystem **1030** and a non-volatile memory subsystem **1040** comprising at least 100 percent more storage capacity than does the volatile memory subsystem. The memory system **1010** also comprises a controller **1062** operatively coupled to the volatile memory subsystem **1030** and operatively coupled to the non-volatile memory subsystem **1040**. The controller **1062** can be configured to allow data to be communicated between the volatile memory subsystem **1030** and the host system when the memory system **1010** is operating in a first state and to allow data to be communicated between the volatile memory subsystem **1030** and the non-volatile memory subsystem **1040** when the memory system **1010** is operating in a second state.

Although the memory system **1010** having extra storage capacity of the non-volatile memory subsystem **1040** has been described with respect to certain embodiments, alter-

native configurations exist. For example, in certain embodiments, there may be more than 100 percent more storage capacity in the non-volatile memory subsystem **1040** than in the volatile memory subsystem **1030**. In various embodiments, there may be at least 200, 300, or 400 percent more storage capacity in the non-volatile memory subsystem **1040** than in the volatile memory subsystem **1030**. In other embodiments, the non-volatile memory subsystem **1040** includes at least some other integer multiples of the storage capacity of the volatile memory subsystem **1030**. In some embodiments, the non-volatile memory subsystem **1040** includes a non-integer multiple of the storage capacity of the volatile memory subsystem **1030**. In one embodiment, the non-volatile memory subsystem **1040** includes less than 100 percent more storage capacity than does the volatile memory subsystem **1030**.

The extra storage capacity of the non-volatile memory subsystem **1040** can be used to improve the backup capability of the memory system **1010**. In certain embodiments in which data can only be written to portions of the non-volatile memory subsystem **1040** which do not contain data (e.g., portions which have been erased), the extra storage capacity of the nonvolatile memory subsystem **1040** allows the volatile memory subsystem **1030** to be backed up in the event of a subsequent power failure or other trigger event. For example, the extra storage capacity of the non-volatile memory subsystem **1040** may allow the memory system **1010** to backup the volatile memory subsystem **1030** efficiently in the event of multiple trigger conditions (e.g., power failures). In the event of a first power failure, for example, the data in the volatile memory system **1030** is copied to a first, previously erased portion of the nonvolatile memory subsystem **1040** via the controller **1062**. Since the non-volatile memory subsystem **1040** has more storage capacity than does the volatile memory subsystem **1030**, there is a second portion of the non-volatile memory subsystem **1040** which does not have data from the volatile memory subsystem **1030** copied to it and which remains free of data (e.g., erased). Once system power is restored, the controller **1062** of the memory system **1010** restores the data to the volatile memory subsystem **1030** by copying the backed-up data from the non-volatile memory subsystem **1040** back to the volatile memory subsystem **1030**. After the data is restored, the memory system **1010** erases the non-volatile memory subsystem **1040**. While the first portion of the non-volatile memory subsystem **1040** is being erased, it may be temporarily unaccessible.

If a subsequent power failure occurs before the first portion of the non-volatile memory subsystem **1040** is completely erased, the volatile memory subsystem **1030** can be backed-up or stored again in the second portion of the non-volatile memory subsystem **1040** as described herein. In certain embodiments, the extra storage capacity of the non-volatile memory subsystem **1040** may allow the memory system **1010** to operate more efficiently. For example, because of the extra storage capacity of the non-volatile memory subsystem **1040**, the memory system **1010** can handle a higher frequency of trigger events that is not limited by the erase time of the non-volatile memory subsystem **1040**.

FIG. **18** is a flowchart of an example method **1300** of controlling a memory system **1010** operatively coupled to a host system and which includes a volatile memory subsystem **1030** and a non-volatile memory subsystem **1040**. In certain embodiments, the non-volatile memory subsystem **1040** comprises at least 100 percent more storage capacity than does the volatile memory subsystem **30** as described herein. While the method **1300** is described herein by reference to the memory system **1010** schematically illustrated by FIGS. **12-14**, the method **1300** can be practiced using other memory systems in accordance with certain embodiments described herein. In an operational block **1310**, the method **1300** comprises communicating data between the volatile memory subsystem **1030** and the host system when the memory system **1010** is in a first mode of operation. The method **1300** further comprises storing a first copy of data from the volatile memory subsystem **1030** to the non-volatile memory subsystem **1040** at a first time when the memory system **1010** is in a second mode of operation in an operational block **1320**.

In an operational block **1330**, the method **1300** comprises restoring the first copy of data from the non-volatile memory subsystem **1040** to the volatile memory subsystem **1030**. The method **1300** further comprises erasing the first copy of data from the non-volatile memory subsystem **1040** in an operational block **1340**. The method further comprises storing a second copy of data from the volatile memory subsystem **1030** to the non-volatile memory subsystem **1040** at a second time when the memory system **1010** is in the second mode of operation in an operational block **1350**. Storing the second copy begins before the first copy is completely erased from the non-volatile memory subsystem **1040**.

In some embodiments, the memory system **1010** enters the second mode of operation in response to a trigger condition, such as a power failure. In certain embodiments, the first copy of data and the second copy of data are stored in separate portions of the nonvolatile memory subsystem **1040**. The method **1300** can also include restoring the second copy of data from the non-volatile memory subsystem **1040** to the volatile memory subsystem **1030** in an operational block **1360**. The operational blocks of method **1300** referred to herein may be performed in different orders in various embodiments. For example, in some embodiments, the second copy of data is restored to the volatile memory subsystem **1030** at operational block **1360** before the first copy of data is completely erased in the operational block **1340**.

FIG. **19** schematically illustrates an example clock distribution topology **1400** of a memory system **1010** in accordance with certain embodiments described herein. The clock distribution topology **1400** generally illustrates the creation and routing of the clock signals provided to the various components of the memory system **1010**. A clock source **1402** such as, for example, a 25 MHz oscillator, generates a clock signal. The clock source **1402** may feed a clock generator **1404** which provides a clock signal **1406** to the controller **1062**, which may be an FPGA. In one embodiment, the clock generator **1404** generates a 125 MHz clock signal **1406**. The controller **1062** receives the clock signal **1406** and uses it to clock the controller **1062** master state control logic. For example, the master state control logic may control the general operation of an FPGA controller **1062**.

The clock signal **1406** can also be input into a clock divider **1410** which produces a frequency-divided version of the clock signal **1406**. In an example embodiment, the clock divider **1410** is a divide by two clock divider and produces a 62.5 MHz clock signal in response to the 125 MHz clock signal **1406**. A non-volatile memory phase-locked loop (PLL) block **1412** can be included (e.g., in the controller **1062**) which distributes a series of clock signals to the non-volatile memory subsystem **1040** and to associated control logic. For example, a series of clock signals **1414** can

be sent from the controller **1062** to the non-volatile memory subsystem **1040**. Another clock signal **1416** can be used by the controller logic which is dedicated to controlling the non-volatile memory subsystem **1040**. For example, the clock signal **1416** may clock the portion of the controller **1062** which is dedicated to generating address and/or control lines for the non-volatile memory subsystem **1040**. A feedback clock signal **1418** is fed back into the non-volatile memory PLL block **1412**. In one embodiment, the PLL block **1412** compares the feedback clock **1418** to the reference clock **1411** and varies the phase and frequency of its output until the reference **1411** and feedback **1418** clocks are phase and frequency matched.

A version of the clock signal **1406** such as the backup clock signal **1408** may be sent from the controller to the volatile memory subsystem **1030**. The clock signal **1408** may be, for example, a differential version of the clock signal **1406**. As described herein, the backup clock signal **1408** may be used to clock the volatile memory subsystem **1030** when the memory system **1010** is backing up the data from the volatile memory subsystem **1030** into the non-volatile memory subsystem **1040**. In certain embodiments, the backup clock signal **1408** may also be used to clock the volatile memory subsystem **1030** when the memory system **1010** is copying the backed-up data back into the volatile memory subsystem **1030** from the nonvolatile memory subsystem **1040** (also referred to as restoring the volatile memory subsystem **1030**). The volatile memory subsystem **1030** may normally be run at a higher frequency (e.g., DRAM running at 400 MHz) than the nonvolatile memory subsystem **1040** (e.g., flash memory running at 62.5 MHz) when communicating with the host system (e.g., when no trigger condition is present). However, in certain embodiments the volatile memory subsystem **1030** may be operated at a reduced frequency (e.g., at twice the frequency of the non-volatile memory subsystem **1040**) without introducing significant delay into the system during backup operation and/or restore operations. Running the volatile memory subsystem **1030** at the reduced frequency during a backup and/or restore operation may advantageously reduce overall power consumption of the memory system **1010**.

In one embodiment, the backup clock **1408** and the volatile memory system clock signal **1420** are received by a multiplexer **1422**, as schematically illustrated by FIG. **19**. The multiplexer **1422** can output either the volatile memory system clock signal **1420** or the backup clock signal **1408** depending on the backup state of the memory system **1010**. For example, when the memory system **1010** is not performing a backup or restore operation and is communicating with the host system (e.g., normal operation), the volatile memory system clock signal **1420** may be provided by the multiplexer **422** to the volatile memory PLL block **1424**. When the memory system **1010** is performing a backup (or restore) operation, the backup clock signal **1408** may be provided.

The volatile memory PLL block **1424** receives the volatile memory reference clock signal **1423** from the multiplexer **1422** and can generate a series of clock signals which are distributed to the volatile memory subsystem **1030** and associated control logic. For example, in one embodiment, the PLL block **1424** generates a series of clock signals **1426** which clock the volatile memory elements **1032**. A clock signal **1428** may be used to clock control logic associated with the volatile memory elements, such as one or more registers (e.g., the one or more registers of a registered DIMM). Another clock signal **1430** may be sent to the controller **1062**. A feedback clock signal **1432** is fed back into the volatile memory PLL block **1424**. In one embodiment, the PLL block **1424** compares the feedback clock signal **1432** to the reference clock signal **1423** and varies the phase and frequency of its output until the reference clock signal **1423** and the feedback clock signal **1432** clocks are phase and frequency matched.

The clock signal **1430** may be used by the controller **1062** to generate and distribute clock signals which will be used by controller logic which is configured to control the volatile memory subsystem **1030**. For example, control logic in the controller **1062** may be used to control the volatile memory subsystem **1030** during a backup or restore operation. The clock signal **1430** may be used as a reference clock signal for the PLL block **1434** which can generate one or more clocks **1438** used by logic in the controller **1062**. For example, the PLL block **1434** may generate one or more clock signals **1438** used to drive logic circuitry associated with controlling the volatile memory subsystem **1030**. In certain embodiments, the PLL block **1434** includes a feedback clock signal **1436** and operates in a similar manner to other PLL blocks described herein.

The clock signal **1430** may be used as a reference clock signal for the PLL block **1440** which may generate one or more clock signals used by a sub-block **1442** to generate one or more other clock signals **1444**. In one embodiment, for example, the volatile memory subsystem **1030** comprises DDR2 SDRAM elements and the sub-block **1442** generates one or more DDR2 compatible clock signals **1444**. A feedback clock signal **1446** is fed back into the PLL block **1440**. In certain embodiments, the PLL block **1440** operates in a similar manner to other PLL blocks described herein.

While described with respect to the example embodiment of FIG. **19**, various alternative clock distribution topologies are possible. For example, one or more of the clock signals have a different frequency in various other embodiments. In some embodiments, one or more of the clocks shown as differential signals are single ended signals. In one embodiment, the volatile memory subsystem **1030** operates on the volatile memory clock signal **1420** and there is no backup clock signal **1408**. In some embodiments, the volatile memory subsystem **1030** is operated at a reduced frequency during a backup operation and not during a restore operation. In other embodiments, the volatile memory subsystem **1030** is operated at a reduced frequency during a restore operation and not during a backup operation.

FIG. **20** is a flowchart of an example method **1500** of controlling a memory system **1010** operatively coupled to a host system. Although described with respect to the memory system **1010** described herein, the method **1500** is compatible with other memory systems. The memory system **1010** may include a clock distribution topology **1400** similar to the one described above with respect to FIG. **19** or another clock distribution topology. The memory system **1010** can include a volatile memory subsystem **1030** and a non-volatile memory subsystem **1040**.

In an operational block **1510**, the method **1500** comprises operating the volatile memory subsystem **1030** at a first frequency when the memory system **1010** is in a first mode of operation in which data is communicated between the volatile memory subsystem **1030** and the host system. In an operational block **1520**, the method **1500** comprises operating the non-volatile memory subsystem **1040** at a second frequency when the memory system **1010** is in a second mode of operation in which data is communicated between the volatile memory subsystem **1030** and the non-volatile memory subsystem **1040**. The method **1500** further comprises operating the volatile memory subsystem **1030** at a

third frequency in an operational block **1530** when the memory system **1010** is in the second mode of operation. In certain embodiments, the memory system **1010** is not powered by a battery when it is in the second mode of operation. The memory system **1010** may switch from the first mode of operation to the second mode of operation in response to a trigger condition. The trigger condition may be any trigger condition described herein such as, for example, a power failure condition. In certain embodiments, the second mode of operation includes both backup and restore operations as described herein. In other embodiments, the second mode of operation includes backup operations but not restore operations. In yet other embodiments, the second mode of operation includes restore operations but not backup operations.

The third frequency can be less than the first frequency. For example, the third frequency can be approximately equal to the second frequency. In certain embodiments, the reduced frequency operation is an optional mode. In yet other embodiments, the first, second and/or third frequencies are configurable by a user or by the memory system **1010**.

FIG. **21** schematically illustrates an example topology of a connection to transfer data slices from two DRAM segments **1630**, **1640** of a volatile memory subsystem **1030** of a memory system **1010** to a controller **1062** of the memory system **1010**. While the example of FIG. **21** shows a topology including two DRAM segments **1630**, **1640** for the purposes of illustration, each address location of the volatile memory subsystem **1030** comprises more than the two segments in certain embodiments. The data lines **1632**, **1642** from the first DRAM segment **1630** and the second DRAM segment **1640** of the volatile memory subsystem **1030** are coupled to switches **1650**, **1652** which are coupled to the controller **1062** (e.g., logic element **1070**) of the memory system **1010**. The chip select lines **1634**, **1644** and the self-refresh lines **1636**, **1646** (e.g., CKe signals) of the first and second DRAM segments **1630**, **1640**, respectively, are coupled to the controller **1062**. In certain embodiments, the controller **1062** comprises a buffer (not shown) which is configured to store data from the volatile memory subsystem **1030**. In certain embodiments, the buffer is a first-in, first out buffer (FIFO). In certain embodiments, data slices from each DRAM segment **1630**, **1640** comprise a portion of the volatile memory subsystem data bus. In one embodiment, for example, the volatile memory subsystem **1030** comprises a 72-bit data bus (e.g., each data word at each addressable location is 72 bits wide and includes, for example, 64 bits of accessible SDRAM and 8 bits of ECC), the first data slice from the first DRAM segment **1630** may comprise 40 bits of the data word, and the second data slice from the second DRAM segment **1640** may comprise the remaining 32 bits of the data word. Certain other embodiments comprise data buses and/or data slices of different sizes.

In certain embodiments, the switches **1650**, **1652** can each be selectively switched to selectively operatively couple the data lines **1632**, **1642**, respectively from the first and second DRAM segments **1630**, **1640** to the controller **1062**. The chip select lines **1634**, **1644** enable the first and second DRAM segments **1630**, **1640**, respectively, of the volatile memory subsystem **1030**, and the self-refresh lines **1636**, **1646** toggle the first and second DRAM segments **1630**, **1640**, respectively, from self-refresh mode to active mode. In certain embodiments, the first and second DRAM segments **1630**, **1640** maintain stored information but are not accessible when they are in self-refresh mode, and maintain stored information and are accessible when they are in active mode.

In certain embodiments, when the memory system **1010** is backing up the volatile memory system **1030**, data slices from only one of the two DRAM segments **1630**, **1640** at a time are sent to the controller **1062**. For example, when the first slice is being written to the controller **1062** during a back-up, the controller **1062** sends a signal via the CKe line **1636** to the first DRAM segment **1630** to put the first DRAM segment **1630** in active mode. In certain embodiments, the data slice from the first DRAM segment **1630** for multiple words (e.g., a block of words) is written to the controller **1062** before writing the second data slice from the second DRAM segment **1640** to the controller **1062**. While the first data slice is being written to the controller **1062**, the controller **1062** also sends a signal via the CKe line **1646** to put the second DRAM segment **1640** in self-refresh mode. Once the first data slice for one word or for a block of words is written to the controller **1062**, the controller **1062** puts the first DRAM segment **1630** into self-refresh mode by sending a signal via the CKe line **1636** to the first DRAM segment **1640**. The controller **1062** also puts the second DRAM segment **1640** into active mode by sending a signal via the CKe line **1646** to the DRAM segment **1640**. The second slice for a word or for a block of words is written to the controller **1062**. In certain embodiments, when the first and second data slices are written to the buffer in the controller **1062**, the controller **1062** combines the first and second data slices **1630**, **1640** into complete words or blocks of words and then writes each complete word or block of words to the non-volatile memory subsystem **1040**. In certain embodiments, this process is called "slicing" the volatile memory subsystem **1030**.

In certain embodiments, the data may be sliced in a restore operation as well as, or instead of, during a backup operation. For example, in one embodiment, the nonvolatile memory elements **1042** write each backed-up data word to the controller **1062** which writes a first slice of the data word to the volatile memory subsystem **1030** and then a second slice of the data word to the volatile memory subsystem **1030**. In certain embodiments, slicing the volatile memory subsystem **1030** during a restore operation may be performed in a manner generally inverse to slicing the volatile memory subsystem **1030** during a backup operation.

FIG. **22** is a flowchart of an example method **1600** of controlling a memory system **1010** operatively coupled to a host system and which includes a volatile memory subsystem **1030** and a non-volatile memory subsystem **1040**. Although described with respect to the memory system **1010** described herein with respect to FIGS. **12-14** and **21**, the method **1600** is compatible with other memory systems. The method **1600** comprises communicating data words between the volatile memory subsystem **1030** and the host system when the memory system **1010** is in a first mode of operation in an operational block **1610**. For example, the memory system **1010** may be in the first mode of operation when no trigger condition has occurred and the memory system is not performing a backup and/or restore operation or is not being powered by a secondary power supply.

In an operational block **1620**, the method further comprises transferring data words from the volatile memory subsystem **1030** to the non-volatile memory subsystem **1040** when the memory system **1010** is in a second mode of operation. In certain embodiments, each data word comprises the data stored in a particular address of the memory system **1010**. The memory system **1010** may enter the second mode of operation, for example, when a trigger condition (e.g., a power failure) occurs. In certain embodiments, transferring each data word comprises storing a first

portion (also referred to as a slice) of the data word in a buffer in an operational block **1622**, storing a second portion of the data word in the buffer in an operational block **1624**, and writing the entire data word from the buffer to the non-volatile memory subsystem **1040** in an operational block **1626**.

In one example embodiment, the data word may be a 72 bit data word (e.g., 64 bits of accessible SDRAM and 8 bits of ECC), the first portion (or "slice") may comprise 40 bits of the data word, and the second portion (or "slice") may comprise the remaining 32 bits of the data word. In certain embodiments, the buffer is included in the controller **1062**. For example, in one embodiment, the buffer is a first-in, first-out buffer implemented in the controller **1062** which comprises an FPGA. The method **1600** may generally be referred to as "slicing" the volatile memory during a backup operation. In the example embodiment, the process of "slicing" the volatile memory during a backup includes bringing the 32-bit slice out of self-refresh, reading a 32-bit block from the slice into the buffer, and putting the 32-bit slice back into self-refresh. The 40-bit slice is then brought out of self-refresh and a 40-bit block from the slice is read into a buffer. Each block may comprise a portion of multiple words. For example, each 32-bit block may comprise 32-bit portions of multiple 72-bit words. In other embodiments, each block comprises a portion of a single word. The 40-bit slice is then put back into self-refresh in the example embodiment. The 32-bit and 40-bit slices are then combined into a 72-bit block by the controller **1062** and ECC detection/correction is performed on each 72-bit word as it is read from the buffer and written into the non-volatile memory subsystem (e.g., flash).

In some embodiments, the entire data word may comprise more than two portions. For example, the entire data word may comprise three portions instead of two and transferring each data word further comprises storing a third portion of each data word in the buffer. In certain other embodiments, the data word may comprise more than three portions.

In certain embodiments, the data may be sliced in a restore operation as well as, or instead of, during a backup operation. For example, in one embodiment, the nonvolatile memory elements **1040** write each backed-up data word to the controller **1062** which writes a first portion of the data word to the volatile memory subsystem **1030** and then a second portion of the data word to the volatile memory **1030**. In certain embodiments, slicing the volatile memory subsystem **1030** during a restore operation may be performed in a manner generally inverse to slicing the volatile memory subsystem **1030** during a backup operation.

The method **1600** can advantageously provide significant power savings and can lead to other advantages. For example, in one embodiment where the volatile memory subsystem **1030** comprises DRAM elements, only the slice of the DRAM which is currently being accessed (e.g., written to the buffer) during a backup is configured in full-operational mode. The slice or slices that are not being accessed may be put in self-refresh mode. Because DRAM in self-refresh mode uses significantly less power than DRAM in full-operational mode, the method **1600** can allow significant power savings. In certain embodiments, each slice of the DRAM includes a separate self-refresh enable (e.g., CKe) signal which allows each slice to be accessed independently.

In addition, the connection between the DRAM elements and the controller **1062** may be as large as the largest slice instead of as large as the data bus. In the example embodiment, the connection between the controller **1062** and the DRAM may be 40 bits instead of 72 bits. As a result, pins on the controller **1062** may be used for other purposes or a smaller controller may be used due to the relatively low number of pin-outs used to connect to the volatile memory subsystem **1030**. In certain other embodiments, the full width of the data bus is connected between the volatile memory subsystem **1030** and the controller **1062** but only a portion of it is used during slicing operations. For example, in some embodiments, memory slicing is an optional mode.

While embodiments and applications have been shown and described, it would be apparent to those skilled in the art having the benefit of this disclosure that many more modifications than mentioned above are possible without departing from the inventive concepts disclosed herein. The invention, therefore, is not to be restricted except in the spirit of the appended claims.

What is claimed is:

1. A memory module comprising:

a printed circuit board (PCB) having an interface configured to fit into a corresponding slot connector of a host system, the interface including a plurality of edge connections configured to couple power, data, address and control signals between the memory module and the host system;

a first buck converter configured to provide a first regulated voltage having a first voltage amplitude;

a second buck converter configured to provide a second regulated voltage having a second voltage amplitude;

a third buck converter configured to provide a third regulated voltage having a third voltage amplitude;

a converter circuit configured to provide a fourth regulated voltage having a fourth voltage amplitude; and

a plurality of components coupled to the PCB, each component of the plurality of components coupled to one or more regulated voltages of the first, second, third and fourth regulated voltages, the plurality of components comprising:

a plurality of synchronous dynamic random access memory (SDRAM) devices coupled to the first regulated voltage, and

at least one circuit coupled between a first portion of the plurality of edge connections and the plurality of SDRAM devices, the at least one circuit operable to (i) receive a first plurality of address and control signals via the first portion of the plurality of edge connections, and (ii) output a second plurality of address and control signals to the plurality of SDRAM devices, the at least one circuit coupled to both the second regulated voltage and the fourth regulated voltage, wherein a first one of the second and fourth voltage amplitudes is less than a second one of the second and fourth voltage amplitudes.

2. The memory module of claim **1**, wherein the first and third buck converters are further configured to operate as a dual buck converter.

3. The memory module of claim **1**, wherein the first voltage amplitude is 1.8 volts.

4. The memory module of claim **1**, wherein the second, third, and fourth voltage amplitudes are 2.5 volts, 1.2 volts, and 3.3 volts, respectively.

5. The memory module of claim **1**, further comprising:

a voltage monitor circuit configured to monitor a power input voltage received via a second portion of the plurality of edge connections, the voltage monitor circuit configured to produce a trigger signal in response to the power input voltage having a voltage amplitude that is greater than a first threshold voltage.

6. The memory module of claim 5, wherein the voltage monitor circuit is further configured to produce the trigger signal in response to the power input voltage having a voltage amplitude that is less than a second threshold voltage.

7. The memory module of claim 6, wherein the second threshold voltage corresponds to a voltage level that is ten percent less than a specified operating voltage.

8. The memory module of claim 1, the plurality of components further comprising:

one or more registers coupled to one of the first, second, third and fourth regulated voltages, the one or more registers configured to register, in response to a clock, the first plurality of address and control signals, wherein the one of the first, second, third and fourth regulated voltages is selectively switched off to turn power off to the one or more registers while one or more components of the plurality of components are powered on.

9. The memory module of claim 5, wherein the first threshold voltage corresponds to a voltage level that is ten percent greater than a specified operating voltage.

10. The memory module of claim 5, the plurality of components further comprising:

a logic element including a non-volatile memory, the non-volatile memory is configured to store configuration information.

11. The memory module of claim 10, wherein, in response to the trigger signal, the logic element writes information into the non-volatile memory.

12. The memory module of claim 5, the plurality of components further comprising:

a non-volatile memory; and

a controller configured to receive the trigger signal, wherein, in response to the trigger signal, the controller performs a write operation to the non-volatile memory.

13. The memory module of claim 5, wherein the power input voltage is coupled to the first, second, and third buck converters and the converter circuit.

14. The memory module of claim 8, wherein, in response to selectively switching on the one of the first, second, third and fourth regulated voltages to the one or more registers, the one or more registers is configured to output the registered first plurality of address and control signals to the plurality of SDRAM devices.

15. The memory module of claim 1, the plurality of components further comprising:

a logic element including one or more integrated circuits and discrete electrical elements, the one or more integrated circuit including an internal non-volatile memory, wherein the non-volatile memory is configured to store configuration information.

16. A memory module comprising:

a printed circuit board (PCB) having an interface configured to fit into a corresponding slot connector of a host system, the interface including a plurality of edge connections configured to couple power, data, address and control signals between the memory module and the host system;

first, second, and third buck converters configured to receive a pre-regulated input voltage and to produce first, second and third regulated voltages, respectively;

a converter circuit configured to reduce the pre-regulated input voltage to provide a fourth regulated voltage, wherein the first, second, third and fourth regulated voltages have first, second, third, and fourth voltage amplitudes, respectively;

a plurality of components coupled to the PCB, the plurality of components including a plurality of synchronous dynamic random access memory (SDRAM) devices, each component of the plurality of components coupled to one or more regulated voltages of the first, second, third and fourth regulated voltages; and

a voltage monitor circuit configured to monitor an input voltage received via a first portion of the plurality of edge connections, the voltage monitor circuit configured to produce a signal in response to the input voltage having a voltage amplitude that is greater than a first threshold voltage.

17. The memory module of claim 16, wherein the second and third buck converters are configured to operate as a dual buck converter.

18. The memory module of claim 16, the plurality of components further including:

a controller coupled to the voltage monitor circuit and configured to receive the signal, wherein the controller executes a write operation in response to the signal.

19. The memory module of claim 18, wherein the write operation includes writing data information into non-volatile memory.

20. The memory module of claim 16, wherein the plurality of SDRAM devices are configured to receive at least one of the first, second, third and fourth regulated voltages having a voltage amplitude of 1.8 volts.

21. The memory module of claim 16, the plurality of components further including:

at least one circuit coupled between the interface and the plurality of SDRAM devices, the at least one circuit operable to receive a first plurality of address and control signals via a second portion of the plurality of edge connections and to output a second plurality of address and control signals to the plurality of SDRAM devices, the at least one circuit coupled to both the second regulated voltage and the fourth regulated voltage, wherein a first one of the second and fourth voltage amplitudes is less than a second one of the second and fourth voltage amplitudes.

22. The memory module of claim 16, the plurality of components further including:

a logic element including an internal non-volatile memory, wherein the non-volatile memory is configured to store configuration information, wherein the configuration information is used to program the logic element.

23. A memory module comprising:

a printed circuit board (PCB) having an interface configured to fit into a corresponding slot connector of a host system, the interface including a plurality of edge connections configured to couple power, data, address and control signals between the memory module and the host system;

a plurality of components coupled to the PCB, each component of the plurality of components coupled to one or more regulated voltages of first, second, third and fourth regulated voltages, the plurality of components including a plurality of synchronous dynamic random access memory (SDRAM) devices and one or more registers, the plurality of SDRAM devices coupled to the first regulated voltage, the one or more registers coupled to (i) the second regulated voltage, (ii) a portion of the plurality of edge connections, and (iii) the plurality of SDRAM devices, wherein a plu-

rality of address and control signals are coupled to the one or more registers via the portion of the plurality of edge connections;

first, second, and third buck converters configured to provide the first, second and third regulated voltages, respectively; and

a converter circuit configured to provide the fourth regulated voltage,

wherein the second regulated voltage is configured to be selectively switched on or off to the one or more registers while at least the plurality of SDRAM devices are powered on,

wherein if the second regulated voltage is switched on while at least the plurality of SDRAM devices are powered on, the one or more registers are configured to couple the first plurality of address and control signals to the plurality of SDRAM devices, and

wherein if the second regulated voltage is switched off while the plurality of SDRAM devices are powered on, the one or more registers are configured to decouple the plurality of SDRAM devices from the first plurality of address and control signals.

24. The memory module of claim **23**, further comprising:

a voltage monitor circuit configured to monitor an input voltage received from the host system via the interface, the voltage monitor circuit configured to produce a signal in response to the input voltage having a voltage amplitude that is greater than a first threshold voltage.

25. The memory module of claim **24**, the plurality of components further including:

a controller coupled to the voltage monitor circuit and configured to receive the signal, wherein, in response to the signal, the controller executes a write operation.

26. The memory module of claim **25**, wherein the write operation includes writing data information to non-volatile memory.

27. The memory module of claim **24**, wherein the voltage monitor circuit is further configured to produce the signal in response to the input voltage having a voltage amplitude that is less than a second threshold voltage.

28. The memory module of claim **23**, wherein the second and third buck converters are configured to operate as a dual buck converter.

29. The memory module of claim **23**, wherein the plurality of SDRAM devices are configured to receive at least one of the first, second, third and fourth regulated voltages having a voltage amplitude of 1.8 volts.

30. The memory module of claim **23**, wherein the first, second, and third buck converters are configured to receive a pre-regulated input voltage and to provide the first, second and third regulated voltages, respectively, and wherein the converter circuit is configured to reduce the pre-regulated voltage input to provide the fourth regulated voltage.

* * * * *

US011232054B2

(12) **United States Patent**
Chen et al.

(10) **Patent No.: US 11,232,054 B2**
(45) **Date of Patent: *Jan. 25, 2022**

(54) **FLASH-DRAM HYBRID MEMORY MODULE**

(71) Applicant: **Netlist, Inc.**, Irvine, CA (US)

(72) Inventors: **Chi-She Chen**, Walnut, CA (US); **Jeffrey C. Solomon**, Irvine, CA (US); **Scott H. Milton**, Irvine, CA (US); **Jayesh Bhakta**, Cerritos, CA (US)

(73) Assignee: **NETLIST, INC.**, Irvine, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **17/328,019**

(22) Filed: **May 24, 2021**

(65) **Prior Publication Data**

US 2021/0279194 A1 Sep. 9, 2021

Related U.S. Application Data

(63) Continuation of application No. 17/138,766, filed on Dec. 30, 2020, now Pat. No. 11,016,918, which is a (Continued)

(51) **Int. Cl.**
G06F 13/28 (2006.01)
G06F 12/02 (2006.01)
(Continued)

(52) **U.S. Cl.**
CPC ***G06F 13/28*** (2013.01); ***G06F 1/185*** (2013.01); ***G06F 3/0613*** (2013.01);
(Continued)

(58) **Field of Classification Search**
CPC G11C 7/1072; G11C 14/0018; G06F 12/0246; G06F 13/4243; G06F 13/1694;
(Continued)

(56) **References Cited**

U.S. PATENT DOCUMENTS

2,043,099 A 6/1936 Hanna
3,562,555 A 2/1971 Ahrons
(Continued)

FOREIGN PATENT DOCUMENTS

KR 0130873 B1 4/1999
KR 100606242 B1 7/2006
WO 2013016723 A2 1/2013

OTHER PUBLICATIONS

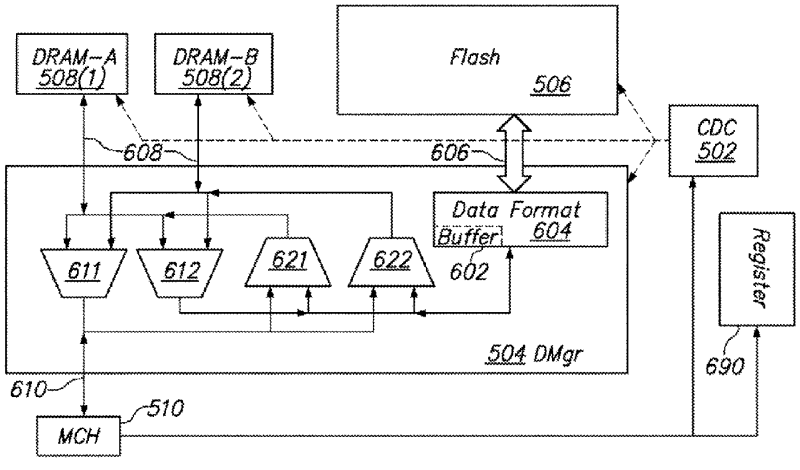
Microsoft Computer Dictionary Fifth Edition, 2002, 9 pages.
(Continued)

Primary Examiner — Hashem Farrokh
(74) *Attorney, Agent, or Firm* — Vorys, Sater, Seymour and Pease LLP

(57) **ABSTRACT**

In certain embodiments, a memory module includes a printed circuit board (PCB) having an interface that couples it to a host system for provision of power, data, address and control signals. First, second, and third buck converters receive a pre-regulated input voltage and produce first, second and third regulated voltages. A converter circuit reduces the pre-regulated input voltage to provide a fourth regulated voltage. Synchronous dynamic random access memory (SDRAM) devices are coupled to one or more regulated voltages of the first, second, third and fourth regulated voltages, and a voltage monitor circuit monitors an input voltage and produces a signal in response to the input voltage having a voltage amplitude that is greater than a threshold voltage.

30 Claims, 22 Drawing Sheets



Related U.S. Application Data

continuation of application No. 15/934,416, filed on Mar. 23, 2018, now abandoned, which is a continuation of application No. 14/840,865, filed on Aug. 31, 2015, now Pat. No. 9,928,186, which is a continuation of application No. 14/489,269, filed on Sep. 17, 2014, now Pat. No. 9,158,684, which is a continuation of application No. 13/559,476, filed on Jul. 26, 2012, now Pat. No. 8,874,831, which is a continuation-in-part of application No. 12/240,916, filed on Sep. 29, 2008, now Pat. No. 8,301,833, which is a continuation of application No. 12/131,873, filed on Jun. 2, 2008, now abandoned.

(60) Provisional application No. 61/512,871, filed on Jul. 28, 2011, provisional application No. 60/941,586, filed on Jun. 1, 2007.

(51) **Int. Cl.**

G06F 13/16	(2006.01)
G06F 1/18	(2006.01)
G06F 12/06	(2006.01)
G06F 13/42	(2006.01)
G11C 7/10	(2006.01)
G11C 14/00	(2006.01)
G06F 3/06	(2006.01)
G06F 13/40	(2006.01)

(52) **U.S. Cl.**

CPC ***G06F 3/0659*** (2013.01); ***G06F 3/0685*** (2013.01); ***G06F 12/0246*** (2013.01); ***G06F 12/0638*** (2013.01); ***G06F 13/1694*** (2013.01); ***G06F 13/4027*** (2013.01); ***G06F 13/4243*** (2013.01); ***G11C 7/1072*** (2013.01); ***G11C 14/0018*** (2013.01); *G06F 2212/205* (2013.01); *G06F 2212/7208* (2013.01)

(58) **Field of Classification Search**

CPC G06F 3/0659; G06F 13/4027; G06F 12/0638; G06F 13/28; G06F 1/185; G06F 2212/205; G06F 2212/7208; G06F 3/0685; G06F 3/0613; Y02D 10/00

USPC 711/110

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

3,916,390	A	10/1975	Chang et al.
4,234,920	A	11/1980	Ness et al.
4,420,821	A	12/1983	Hoffman
4,449,205	A	5/1984	Hoffman
4,607,332	A	8/1986	Goldberg
4,658,204	A	4/1987	Goodwin
4,882,709	A	11/1989	Wyland
4,884,242	A	11/1989	Lacy et al.
4,965,828	A	10/1990	Ergott, Jr. et al.
5,430,742	A	7/1995	Jeddeloh et al.
5,444,664	A	8/1995	Kuroda et al.
5,490,155	A	2/1996	Abdoo et al.
5,519,663	A	5/1996	Harper, Jr. et al.
5,519,831	A	5/1996	Holzhammer
5,563,839	A	10/1996	Herdt et al.
5,577,213	A	11/1996	Avery et al.
5,619,644	A	4/1997	Crockett et al.
5,630,096	A	5/1997	Zuravleff et al.
5,675,725	A	10/1997	Malcolm
5,757,712	A	5/1998	Nagel et al.
5,799,200	A	8/1998	Brant et al.
5,813,029	A	9/1998	Klein
5,870,350	A	2/1999	Bertin et al.
5,874,995	A	2/1999	Naimpally et al.
5,890,192	A	3/1999	Lee et al.
5,953,215	A	9/1999	Karabatsos
5,991,885	A	11/1999	Chang et al.
6,023,421	A	2/2000	Iadanza et al.
6,026,465	A	2/2000	Mills et al.
6,065,092	A	5/2000	Roy
6,112,310	A	8/2000	Jun et al.
6,145,068	A	11/2000	Lewis
6,158,015	A	12/2000	Klein
6,199,142	B1	3/2001	Saulsbury et al.
6,216,247	B1	4/2001	Creta et al.
6,269,382	B1	7/2001	Cabrera et al.
6,336,174	B1	1/2002	Li et al.
6,336,176	B1	1/2002	Leyda et al.
6,363,450	B1	3/2002	Lash et al.
6,421,279	B1	7/2002	Tobita et al.
6,459,647	B1	10/2002	Kengeri
6,487,102	B1	11/2002	Halbert et al.
6,487,623	B1	11/2002	Emerson et al.
6,571,244	B1	5/2003	Larson
6,614,685	B2	9/2003	Wong
6,658,507	B1	12/2003	Chan
6,691,209	B1	2/2004	O'connell
6,693,840	B2	2/2004	Shimada et al.
6,721,212	B2	4/2004	Sasaki
6,721,860	B2	4/2004	Klein
6,769,081	B1	7/2004	Parulkar
6,799,241	B2	9/2004	Kahn et al.
6,799,244	B2	9/2004	Tanaka et al.
6,810,513	B1	10/2004	Vest
6,816,982	B2	11/2004	Ravid
6,839,774	B1	1/2005	Ahn et al.
6,910,635	B1	6/2005	Miks et al.
6,944,042	B2	9/2005	Komatsuzaki
6,948,029	B2	9/2005	Yano
6,952,368	B2	10/2005	Miura et al.
7,053,470	B1	5/2006	Sellers et al.
7,062,618	B2	6/2006	Tsunoda et al.
7,089,412	B2	8/2006	Chen
7,102,391	B1	9/2006	Sun et al.
7,107,480	B1	9/2006	Moshayedi et al.
7,111,142	B2	9/2006	Spencer et al.
7,136,978	B2	11/2006	Miura et al.
7,155,627	B2	12/2006	Matsui
7,200,021	B2	4/2007	Raghuram
7,234,099	B2	6/2007	Gower et al.
7,353,325	B2	4/2008	Lofgren et al.
7,409,491	B2	8/2008	Doblar et al.
7,409,590	B2	8/2008	Moshayedi et al.
7,411,859	B2	8/2008	Sohn et al.
7,421,552	B2	9/2008	Long
7,467,251	B2	12/2008	Park et al.
7,519,754	B2	4/2009	Wang et al.
7,600,142	B2	10/2009	Ichikawa
7,613,877	B2	11/2009	Shimozono et al.
7,716,411	B2	5/2010	Panabaker et al.
7,818,488	B2	10/2010	Park et al.
7,873,750	B2	1/2011	Yabuta et al.
7,881,150	B2	2/2011	Solomon et al.
7,952,179	B2	5/2011	Chiu et al.
8,001,434	B1	8/2011	Lee et al.
8,081,536	B1	12/2011	Solomon et al.
8,086,955	B2	12/2011	Zhou et al.
8,102,614	B2	1/2012	Song et al.
8,214,616	B2	7/2012	Ware et al.
8,233,303	B2	7/2012	Best et al.
8,301,833	B1	10/2012	Chen et al.
8,407,395	B2	3/2013	Kim et al.
8,412,879	B2	4/2013	Chang et al.
8,516,187	B2	8/2013	Chen et al.
8,671,243	B2	3/2014	Chen et al.
8,677,060	B2	3/2014	Chen et al.
8,874,831	B2	10/2014	Lee et al.
8,880,791	B2	11/2014	Chen et al.
8,904,098	B2	12/2014	Amidi et al.
8,904,099	B2	12/2014	Chen et al.
9,043,677	B2	5/2015	Kong et al.
9,158,684	B2	10/2015	Lee et al.

(56) **References Cited**

U.S. PATENT DOCUMENTS

9,361,250	B2	6/2016	Shan et al.	
9,436,600	B2	9/2016	Lee	
9,921,762	B2	3/2018	Amidi et al.	
9,928,186	B2	3/2018	Lee et al.	
2002/0053944	A1	5/2002	Brass et al.	
2002/0083368	A1	6/2002	Abe et al.	
2002/0199061	A1	12/2002	Friedman et al.	
2003/0028733	A1	2/2003	Tsunoda et al.	
2003/0076726	A1	4/2003	Cowles et al.	
2003/0137881	A1	7/2003	Sasaki	
2003/0147297	A1	8/2003	Shiota et al.	
2003/0158995	A1	8/2003	Lee et al.	
2003/0204776	A1	10/2003	Testin	
2003/0206478	A1	11/2003	Ayukawa et al.	
2003/0210601	A1	11/2003	Lin et al.	
2004/0088508	A1	5/2004	Ballard et al.	
2004/0163027	A1	8/2004	MacLaren et al.	
2004/0190210	A1	9/2004	Leete	
2005/0044302	A1	2/2005	Pauley et al.	
2005/0060488	A1	3/2005	Poechmueller	
2005/0132250	A1	6/2005	Hansen et al.	
2005/0141273	A1	6/2005	Park et al.	
2005/0144418	A1	6/2005	Kita	
2005/0185472	A1	8/2005	Randell et al.	
2005/0204091	A1	9/2005	Kilbuck et al.	
2005/0249011	A1	11/2005	Maeda	
2005/0273548	A1	12/2005	Roohparvar	
2006/0039197	A1	2/2006	Khouri et al.	
2006/0069896	A1	3/2006	Sanders	
2006/0080515	A1	4/2006	Spiers et al.	
2006/0126369	A1	6/2006	Raghuram	
2006/0212651	A1	9/2006	Ashmore et al.	
2006/0294295	A1	12/2006	Fukuzo	
2007/0070669	A1	3/2007	Tsern	
2007/0136523	A1*	6/2007	Bonella	G06F 9/4401 711/113
2007/0147115	A1	6/2007	Lin et al.	
2007/0192627	A1	8/2007	Oshikiri	
2007/0255898	A1	11/2007	Nishide et al.	
2007/0276995	A1	11/2007	Caulkins et al.	
2007/0288683	A1	12/2007	Panabaker et al.	
2008/0104344	A1	5/2008	Shimozono et al.	
2008/0126624	A1*	5/2008	Prete	G11C 7/106 710/53
2008/0126690	A1	5/2008	Rajan et al.	
2008/0147968	A1	6/2008	Lee et al.	
2008/0189479	A1	8/2008	Cope et al.	
2008/0195806	A1	8/2008	Cope	
2008/0235443	A1	9/2008	Chow et al.	
2008/0291727	A1	11/2008	Seo et al.	
2009/0031099	A1	1/2009	Sartore	
2009/0235038	A1	9/2009	Sartore	
2010/0110748	A1	5/2010	Best	
2010/0122200	A1	5/2010	Merry, Jr. et al.	
2010/0274953	A1	10/2010	Lee et al.	
2010/0322020	A1	12/2010	Kim	
2011/0078496	A1	3/2011	Jeddeloh	
2011/0161569	A1	6/2011	Shan et al.	
2011/0320804	A1	12/2011	Chan et al.	
2012/0110417	A1	5/2012	Abreu et al.	
2012/0117402	A1	5/2012	Machnicki et al.	
2012/0204079	A1	8/2012	Takefman et al.	
2012/0265952	A1	10/2012	Kurita	
2012/0271990	A1*	10/2012	Chen	G06F 12/00 711/103
2012/0317433	A1	12/2012	Ellis et al.	
2013/0019076	A1	1/2013	Amidi et al.	
2013/0086309	A1	4/2013	Lee et al.	
2013/0254456	A1	9/2013	Chen et al.	
2013/0254497	A1	9/2013	Chen et al.	
2014/0032820	A1	1/2014	Harasawa et al.	
2014/0059170	A1	2/2014	Gasparakis et al.	
2014/0156919	A1	6/2014	Chen et al.	
2014/0156920	A1	6/2014	Chen et al.	
2015/0058701	A1	2/2015	Xing et al.	

OTHER PUBLICATIONS

Mutnuary, B. et al., "Analysis of Fully Buffered DIMM Interface in High-speed Server Applications", IBM Corp, xSeries eServer Development, 2006 Electronic Components and Technology Conferencepp. 203-208.

Notice of Allowance in U.S. Appl. No. 13/536,173, dated Jul. 2, 2013.

Notice of Allowance in U.S. Appl. No. 13/559,476, dated May 6, 2014.

Notice of Allowance in U.S. Appl. No. 13/559,476, dated Sep. 29, 2014.

Notice of Allowance in U.S. Appl. No. 13/905,048, dated Dec. 19, 2013, 8 pages.

Notice of Allowance in U.S. Appl. No. 13/905,053, dated Dec. 11, 2013.

Notice of Allowance in U.S. Appl. No. 14/173,219 dated Jul. 7, 2014.

Notice of Allowance in U.S. Appl. No. 14/489,269, dated Oct. 8, 2015.

Notice of Allowance in U.S. Appl. No. 17/138,766, dated Mar. 24, 2021, 10 pages.

Notice of Allowance, *SanDisk Corporation* v. *Netlist*, U.S. Appl. No. 13/536,173, IPR2014-00982 (PTAB), dated Jul. 2, 2013, 8 pages.

Office Action dated Aug. 19, 2016 of the Chinese Patent Application No. 201280047758.X, 9 pages.

Office Action in U.S. Appl. No. 12/240,916, dated Apr. 3, 2012, pp. 1-12.

Office Action in U.S. Appl. No. 12/240,916, dated Feb. 1, 2012.

Office Action in U.S. Appl. No. 12/240,916, dated Apr. 3, 2012.

Office Action in U.S. Appl. No. 12/240,916, dated Jul. 29, 2011.

Office Action in U.S. Appl. No. 13/536,176, dated Apr. 15, 2013.

Office Action in U.S. Appl. No. 13/625,563, dated Aug. 5, 2013.

Office Action in U.S. Appl. No. 13/625,563, dated May 9, 2014.

Office Action in U.S. Appl. No. 13/905,048, dated Aug. 1, 2013.

Office Action in U.S. Appl. No. 13/905,053, dated Aug. 1, 2013.

Office Action in U.S. Appl. No. 14/173,219, dated Mar. 13, 2014.

Office Action in U.S. Appl. No. 14/173,242, dated Mar. 14, 2014.

Office Action in U.S. Appl. No. 14/302,292, dated Dec. 21, 2015.

Office Action in U.S. Appl. No. 15/934,416, dated Apr. 2, 2020, 21 pages.

Office Action in U.S. Appl. No. 15/934,416, dated Jan. 1, 2021, 17 pages.

Office Action in U.S. Appl. No. 15/934,416, dated Mar. 21, 2019, 8 pages.

Office Action in U.S. Appl. No. 15/934,416, dated Oct. 15, 2019, 13 pages.

Office Action in U.S. Appl. No. 17/138,766, dated Mar. 2, 2021, 7 pages.

Office Action in U.S. Appl. No. 12/240,916, dated Feb. 1, 2012, 14 pages.

Office Action in U.S. Appl. No. 14/173,242, dated Mar. 14, 2014, 7 pages.

Office Action, U.S. Appl. No. 12/240,916, IPR2014-0099-1005 (PTAB), dated Jul. 29, 2011, 8 pages.

Office Action, U.S. Appl. No. 12/240,916, IPR2014-0099-1010 (PTAB), dated Apr. 3, 2012, 11 pages.

Office Action, U.S. Appl. No. 13/536,173, IPR2014-00982 (PTAB), dated Apr. 15, 2013, 9 pages.

Patent Owner's Demonstratives, U.S. Pat. No. 8,671,243, Case No. IPR 2017-00587-2023, (PTAB), 57 pages.

Patent Owner's Listing of New Arguments and Evidence in Petitioners' Reply, U.S. Pat. No. 8,671,243, Case No. IPR2017-00587 (PTAB), filed Jan. 29, 2018, 6 pages.

Patent Owner's Opposition to Petitioners' Motion to Exclude, U.S. Pat. No. 8,874,831, Case No. IPR2017-00692 (PTAB), filed Apr. 2, 2018, 2017, 11 pages.

Patent Owner's Preliminary Response, U.S. Pat. No. 8,516,187, Case IPR2014-01371 (PTAB), filed Dec. 16, 2014, 66 pages.

(56) **References Cited**

OTHER PUBLICATIONS

Patent Owner's Preliminary Response, *SanDisk Corporation* v. *Netlist, Inc.*, U.S. Pat. No. 8,301,833, IPR2014-00994 (PTAB), Paper 8, dated Oct. 2, 2014, 60 pages.
Wong, A., The BIOS Optimization Guide, Adrian's Rojak Pot, Rev. 6.2, 1998-2001, 67 pages.
240pin DDR2 MetaSDRAM Registered DIMM based on 1 GB version C, Hynix Semiconductor, Product Description Rev. 0.2, Sep. 2008, 32 pages.
Advisory Action in U.S. Appl. No. 12/240,916, dated Mar. 13, 2012.
Amazon.com, Search for "memory module", downloaded Oct. 12, 2017, 6 pages.
Amendment and Reply to Office Action, *SanDisk Corporation* v. *Netlist*, U.S. Appl. No. 13/536,173, IPR2014-00982 (PTAB), dated May 21, 2013, 24 pages.
Amendment and Response to Election Restriction, *SanDisk Corporation* v. *Netlist, Inc.*, U.S. Pat. No. 8,301,833, IPR2014-00994-1004 (PTAB), dated May 20, 2011, 9 pages.
American Heritage Dictionary of the English Language, Third Ed., Houghton Mifflin Company, Boston, MA, 1996, 7 pages.
American National Standard Dictionary of Electrical and Electrical Terms, IEEE, Fourth Edition, Revised, ANS/IEEE Std 100-1988, Institute of Electrical Engineers, Nov. 3, 1988, 6 pages.
Annotated added to Russel Jacob Baker Deposition, Exhibit 1026, p. 78, U.S. Pat. No. 8,671,243, Case No. IPR-2017-00587 (PTAB), filed Dec. 18, 2017, 1 page.
Annotated added to Russel Jacob Baker Deposition, Exhibit 1027, p. 79, U.S. Pat. No. 8,671,243, Case No. IPR-2017-00587 (PTAB), filed Dec. 18, 2017, 1 page.
Annotated added to Russel Jacob Baker Deposition, Exhibit 1028, p. 79, U.S. Pat. No. 8,671,243, Case No. IPR-2017-00587 (PTAB), filed Dec. 18, 2017, 1 page.
Appeals from the USPTO, PTAB in Nos. IPR2014-00882, IPR2014-00883, IPR2014-01011, US Court of Appeals for the Federal Circuit, decided Jul. 25, 2017, 8 pages.
Bonella, R., Provisional Application for Advanced Dynamic Disk Memory Module, 53 pages.
Bruce, J., Synchronous DRAM Architectures, Organizations, and Alternate Technologies, Electrical and Computer Engineering Dept., Univ. of Maryland, Dec. 10, 2002, 22 pages.
Catsoulis, J., Designing Embedded Hardware: Create New Computers and Devices, O'Reilly Media, Inc. (2005), 67 pages.
Charles C., Multiplexers and Demultiplexers, Ex. 1024, IPR No. 2017-00587-1024(PTAB), 2003, 8 pages.
Charles C., Multiplexers and Demultiplexers, Ex. 1033, IPR No. 2017-00587-1033 (PTAB), 2003, 8 pages.
Charles C., Multiplexers and Demultiplexers, Ex. 1034, IPR No. 2017-00587-1033 (PTAB), 2003, 8 pages.
Charles C., Multiplexers and Demultiplexers, Ex. 1035, IPR No. 2017-00587-1035 (PTAB), 2003, 8 pages.
Chiappetta, M., Kingston Launches 667-MHz DDR2 SO-DIMM Memory, Oct. 11, 2005, 3 pages.
Notice of Final Rejection, South Korea Patent Office with English Translation, 4 pages.
Notice of Submission of Opinion, South Korea Patent Office with English Translation, 6 pages.
Corrected Petition for Inter Partes Review of Claims 1-30 of U.S. Pat. No. 8,301,833, Declaration of Michael F. Heafey, IPR2014-01370 (PTAB), filed Sep. 22, 2014, 4 pages.
Corrected Petition for Inter Partes Review of Claims 1-32 of U.S. Pat. No. 8,516,187, IPR2014-01371 (PTAB), Sep. 22, 2014, 67 pages.
Corrected Petition for Inter Partes Review of Claims 1-32 of U.S. Pat. No. 8,516,187, Declaration of Dr. Nader Bagherzadeh, IPR2014-01371 (PTAB), filed Sep. 22, 2014, 306 pages.
Data Sheet, 74F257A Quad 2-line to 1-line selector/multiplexer, non-inverting (3-State), Product specification, IC15 Data Handbook Mar. 31, 1995, 10 pages.
David, H. et al., Fully Buffered DIMM (FB-DIMM) Design Considerations, Intel Developer Forum, Intel Corp., Feb. 18, 2004, 36 pages.
Decision Denying Institution of Inter Partes Review, U.S. Pat. No. 8,301,833, Case IPR2017-00649-7 (PTAB), Paper 7, entered Jul. 24, 2017, 17 pages.
Decision Denying Institution of Inter Partes Review, U.S. Pat. No. 8,516,187, IPR2014-01371 (PTAB), Paper 12, entered Mar. 13, 2015, 22 pages.
Decision Denying Institution of Inter Partes Review, *SanDisk Corporation* v. *Netlist, Inc.*, U.S. Pat. No. 8,301,833, IPR2014-00994 (PTAB), Paper 8, dated Dec. 16, 2014, 16 pages.
Decision Denying Institution of Inter Partes Review, *SanDisk Corporation* v. *Netlist*, U.S. Pat. No. 8,516,187, IPR2014-00982 (PTAB), Paper 9, dated Dec. 22, 2014, 16 pages.
Decision Denying Institution of Inter Partes Review, *Smart Modular Tech* v *Netlist Inc.*, U.S. Pat. No. 8,301,833, IPR2014-01370 (PTAB), Paper 13, entered Mar. 13, 2015, 19 pages.
Decision Instituting Inter Partes Review, U.S. Pat. No. 8,671,243, Case No. IPR2017-00587-7(PTAB), Paper No. 7, entered Jun. 22, 2017, 40 pages.
Decision Instituting Inter Partes Review, U.S. Pat. No. 8,874,831, Case No. IPR2017-00692-7(PTAB), Paper No. 7, entered Jul. 21, 2017, 29 pages.
Declaration of R. Jacob Baker, U.S. Pat. No. 8,874,831, Case IPR2017-00692-2016 (PTAB), filed Nov. 10, 2017, 72 pages.
Declaration of R. Jacob Baker, U.S. Pat. No. 8,671,243, Case IPR2017-00587-2016 (PTAB), filed Oct. 13, 2019, 109 pages.
Declaration of Dr. Nader Bagherzadeh, U.S. Pat. No. 8,516,187, IPR2014-01371 (PTAB), filed Sep. 22, 2014, 306 pages.
Declaration of Jeff McMullen, *Netlist, Inc.* v. *Diablo Technologies, Inc.*, Case No. 4: I 3-CV-05889-YGR (NDCA), Document 362-1, filed Sep. 5, 2018, 3 pages.
Declaration of Paul Min, In Inter Partes Review of U.S. Pat. No. 8,301,833, IPR2014-00994-1020 (PTAB), filed Sep. 29, 2008, 215 pages.
Declaration of Paul Min, In Inter Partes Review of U.S. Pat. No. 8,516,187, IPR2014-00982-1013 (PTAB), filed Jun. 28, 2012, pp. 240.
Declaration of Ron Maltiel, U.S. Pat. No. 8,671,243, No. IPR2017-00587-1003 (PTAB), filed May 29, 2013, 131 pages.
Declaration of Ron Maltiel, U.S. Pat. No. 8,874,831, No. IPR2017-00692-1003 (PTAB), filed Jul. 26, 2012 29, 172 pages.
Declaration of Steven J. Corr, *Netlist, Inc.* v. *SMART Storage Systems, Inc., et al.*, Case 4:13-cv-05889-YGR (NDCA), Document 305-8, filed Mar. 10, 2015, 3 pages.
Deposition of Baker, Exhibit 1030, p. 78, U.S. Pat. No. 8,671,243, *SK hynix Inc., et al.* v. *Netlist, Inc.*, Case No. IPR-2017-00587 (PTAB), filed Dec. 18, 2017, 268 pages.
Deposition of Ron Maltiel, U.S. Pat. No. 8,874,831, Case No. IPR2017-00692-2010 (PTAB), dated Sep. 27, 2017, 155 pages.
Netlist v *Smart Storage System and Diablo Technologies*, Inc.'s Invalidity Contentions, Case No. 13-CV-05889 YGR, dated Jun. 6, 2014, 109 pages.
Drawing by R. Jacob Bake regarding Address of Dram and Flash, U.S. Pat. No. 8,671,243, Case IPR2017-00587-1023 (PTAB), filed Dec. 18, 2017, 1 page.
Elmhurst, D. et al., A 1.8-V 128-Mb 125-MHz Multilevel Cell Flash Memory With Flexible Read While Write, IEEE Journal of Solid-State Circuits, vol. 38, No. 11, Nov. 2003, 5 pages.
Ex. 1004—IPR2017-00587 Ron Maltiel CV, 7 pages.
Exhibit 1: Claim Chart Comparing Netlist's U.S. Pat. No. 8,001,434 to Smart Storage Ulltradimm, Case4:13-cv-05889-YGR Document 193-1, filed Apr. 10, 2014, 21 pages.
Exhibit, Decision Denying Institution of Inter Partes Review, *Smart Modular Technologies, Inc.* v. *Netlist, Inc.*, Case No. 4:13-cv-05889-YGR Document 309-3, filed Mar. 17, 2015, 23 pages.
Exhibit, Decision Denying Institution of Inter Partes Review, *Smart Modular Technologies, Inc.* v. *Netlist, Inc.*, Case No. 4:13-cv-05889-YGR Document 309-4, filed Mar. 17, 2015, 20 pages.
Exhibit, Decision Denying Institution of Inter Partes Review, *Smart Modular Technologies, Inc.* v. *Netlist, Inc.*, Case No. 4:13-cv-05889-YGR Document 309-5, filed Mar. 17, 2015, 28 pages.

(56) **References Cited**

OTHER PUBLICATIONS

Exhibit 12, Decision Denying Institution of Inter Partes Review, *Smart Modular Technologies, Inc.* v. *Netlist, Inc.*, Case No. 4:13-cv-05889-YGR Document 309-7, filed Mar. 17, 2015, 17 pages.
Exhibit 9, Email from Defendant's Counsel, *Netlist* v. *Smart Storage Systems, Inc. et al.*, Case 4:13-cv-05889-YGR Document 305-17, filed Mar. 10, 2015, 6 pages.
Exhibit, Institution of Inter Partes Review, *Sandisk Corporation* v. *Netlist, Inc.*, Case 4:13-cv-05889-YGR Document 316-9, filed Mar. 24, 2015, 29 pages.
Exhibit, Institution of Inter Partes Review, *Smart Modular Technologies, Inc.* v. *Netlist, Inc.*, Case 4:13-cv-05889-YGR Document 309-1, filed Mar. 17, 2015, 22 pages.
Exhibit, Letter from Defendant's Counsel, *Netlist* v. *Smart Storage Systems, Inc. et. al.*, Case 4:13-cv-05889-YGR Document 305-16, filed Mar. 10, 2015, 3 pages.
Exhibit, Letter sent via email on Dec. 6, 2013, Case 4:13-cv-03901-YGR Document 53-2, filed Jan. 6, 2014.
Exhibit, Order Denying Defendant's Motion to Stay Pending Inter Partes Review (Doc.59), *The Procter and Gamble Company* v. *Team Technologies, Inc, et al.*, Case 4:13-cv-05889-YGR Document 316-6, filed Mar. 24, 2015, 10 pages.
Exhibit, Patent Public Advisory Committee Quarterly Meeting, Appeals Statistics USPTO, Case No. 4:13-cv-05889-YGR Document 309-8, filed Mar. 17, 2015, 23 pages.
Exhibit, Reporter's Transcript of Proceedings, *Netlist, Inc* v. *Smart Modular Technologies, Inc., et al.*, Case 4:13-cv-05889-YGR Document 316-3, filed Mar. 24, 2015, 15 pages.
Exhibit, Transcript of Official Electronic Sound Recording Proceeding, *Netlist* v. *Smart Modular Technologies, Inc, et al.*, Case 4:13-cv-05889-YGR Document 305-7, filed Mar. 10, 2015, 10 pages.
Hasan, J. et al. Efficient Use of Memory Bandwidth to Improve Network Processor Throughput, Proceedings of the 30th Annual International Symposium on Computer Architecture (ISCA'03), IEEE, 2003, 12 pages.
Horowitz, P. et al., "The Art of Electronics", Cambridge University Press 2nd Ed. 1989, pp. 471, 495-496.
Innis, J., "MPC8560 PowerQUICC III Compact Flash Interface Design", Freescale Semiconductor, Inc., 2004-2006, pp. 1-23.
Intel 1.8 Volt Intel StrataFlash Wireless Memory (L18), 2003, 100 pages.
Inter Partes Review No. IPR2017-00692 (PTAB), U.S. Pat. No. 8,874,831, filed Jul. 26, 2012, 78 pages.
JEDEC Definition of DIMM, Exhibit 1029, IPR No. 2017-00587, Dec. 18, 2017, 2 pages.
JEDEC Global Standard for the Microelectronics Industry, Why JEDEC Standards Matter, 2014, 1 page.
JEDEC Standard No. 21-C (Release 17), Annex J: Serial Presence Detects for DDR2 SDRAM (Revision 1.3), 60 pages.
JEDEC Standard, Double Data Rate (DDR): SDRAM Specification: JESD79C (Revision JESD79B), Mar. 2003, pp. 1-75.
JEDEC Standard, FBDIMM Specification: DDR2 SDRAM Fully Buffered DIMM (FBDIMM) Design Specification: JESD205, JEDEC Solid State Tech. Assoc., Mar. 2007, 129 pages.
Patent Owner's Preliminary Response, *SanDisk Corporation* v. *Netlist*, U.S. Pat. No. 8,516,187, IPR2014-00982 (PTAB), dated Sep. 26, 2014, 57 pages.
Patent Owner's Preliminary Response, *SK hynix Inc., et al.,* v. *Netlist*, U.S. Pat. No. 8,301,833, Case No. IPR2017-00649 (PTAB), filed May 1, 2017, 67 pages.
Patent Owner's Preliminary Response, *SK hynix Inc., et al.,* v. *Netlist*, U.S. Pat. No. 8,874,831, Case No. IPR2017-00692 (PTAB), filed May 1, 2017, 48 pages.
Patent Owner's Preliminary Response, *Smart Modular Tech* v *Netlist Inc.*, U.S. Pat. No. 8,301,833, Case IPR2014-01370 (PTAB), filed Dec. 16, 2014, 66 pages.
Patent Owner's Response, U.S. Pat. No. 8,671,243, Case No. IPR2017-00587-12 (PTAB), filed Oct. 13, 2017, 80 pages.
Patent Owner's Response, *SK hynix Inc., et al.,* v. *Netlist*, U.S. Pat. No. 8,874,831, Case No. IPR2017-00692-12 (PTAB), filed Nov. 10, 2017, 77 pages.
Patterson et al., "Computer Organization & Design: The Hardware/Software Interface" Morgan Kaufmann Publishers, Inc. (1998), 71 pages.
Petition for Inter Partes Review of Claims 1-30 of U.S. Pat. No. 8,301,833, IPR2014-01370 (PTAB), filed Aug. 22, 2014, 68 pages.
Petition for Inter Partes Review of U.S. Pat. No. 8,301,833, filed Sep. 29, 2008.
Petition for Inter Partes Review of U.S. Pat. No. 8,671,243, Case No. IPR2017-00587-1(PTAB), May 29, 2013, 82 pages.
Petitioners Demonstratives, *SK hynix Inc., et al.,* v. *Netlist, Inc.*, U.S. Pat. No. 8,671,243, IPR2017-00587-1037, 89 pages.
Petitioners' Reply to Patent Owner's Response, Inter Partes Review No. IPR2017-00587-16, U.S. Pat. No. 8,671,243, Filed Jan. 12, 2018, 38 pages.
Petitioners' Reply, U.S. Pat. No. 8,874,831, Case No. IPR2017-00692 (PTAB), Paper No. 15, filed Mar. 2, 2018, 35 pages.
Provisional Application for Advance Dynamic Disk Memory Module, Specification, *SK hynix Inc.,* V. *Netlist, Inc.*, Case No. IPR2017-00649-1006 (PTAB), dated Dec. 8, 2005, 53 pages.
Supplemental Declaration of Daniel E. Alberti in Response to Court Order in Support of Motion to Withdraw as Counsel of Record, *Netlist* v. *Smart Storage Systems, Inc. et. al.*, Case 4:13-cv-05889-YGR Document 362, filed Sep. 5, 2018, 3 pages.
Supplemental Declaration of Ronald H. Spuhler, *Netlist* v. *Smart Storage Systems, Inc. et. al.*, Case 4:13-cv-05889-YGR Document 305-1, filed Mar. 10, 2015, 2 pages.
Supplemental Declaration of Steven J. Corr, *Netlist* v. *Smart Storage Systems, Inc. et. al.*, Case 4:13-cv-05889-YGR Document 309, filed Mar. 17, 2015, 3 pages.
Switches—D&T Online, Switches, http://wiki.dtonline.org, Exhibit 1022, Case No. IPR2017-00587, accessed Oct. 13, 2017, 3 pages.
U.S. Office Action in U.S. Appl. No. 13/536,173, dated Apr. 15, 2013, pp. 1-10.
U.S. Appl. No. 12/240,916, Case No. IPR2017-00692-2018 (PTAB), dated Sep. 29, 2008, 52 pages.
Webster's II New College Dictionary, Houghton Mifflin Company, Boston, MA, 2001, pp. 257, 259; 1114,1115, Exhibit 1017, IPR2014-01370 (PTAB), *Smart Modular Tech* v *Netlist Inc*, U.S. Pat. No. 8,301,833, Sep. 22, 2014, 5 pages.
Ian R. Sinclair and John Dunton, Practical Electronics Handbook, Sixth Edition 2007, ELSEVIER, <http://al-marefah.blogspot.com>, retrieved from the Internet Nov. 18, 2021.

* cited by examiner

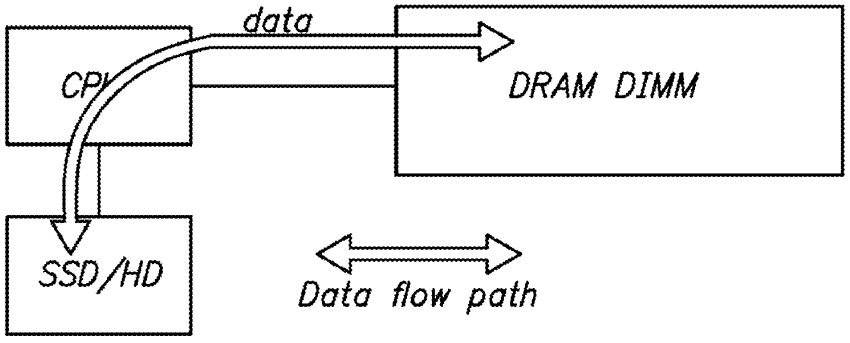


FIG. 1
(PRIOR ART)

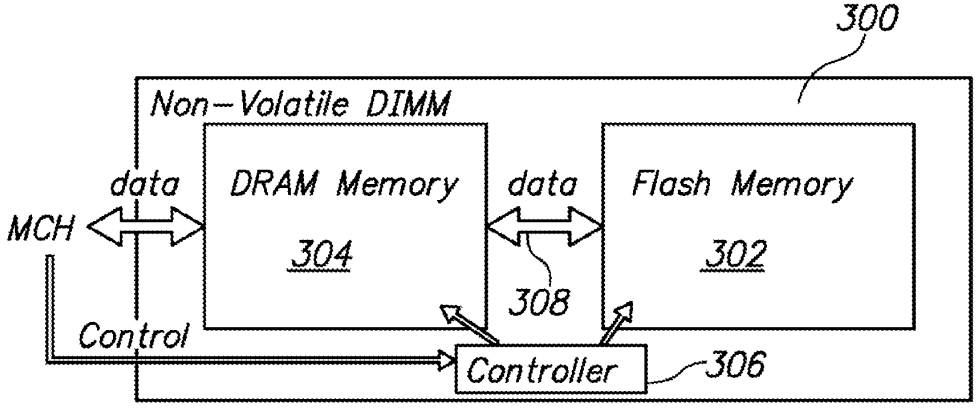


FIG. 3A

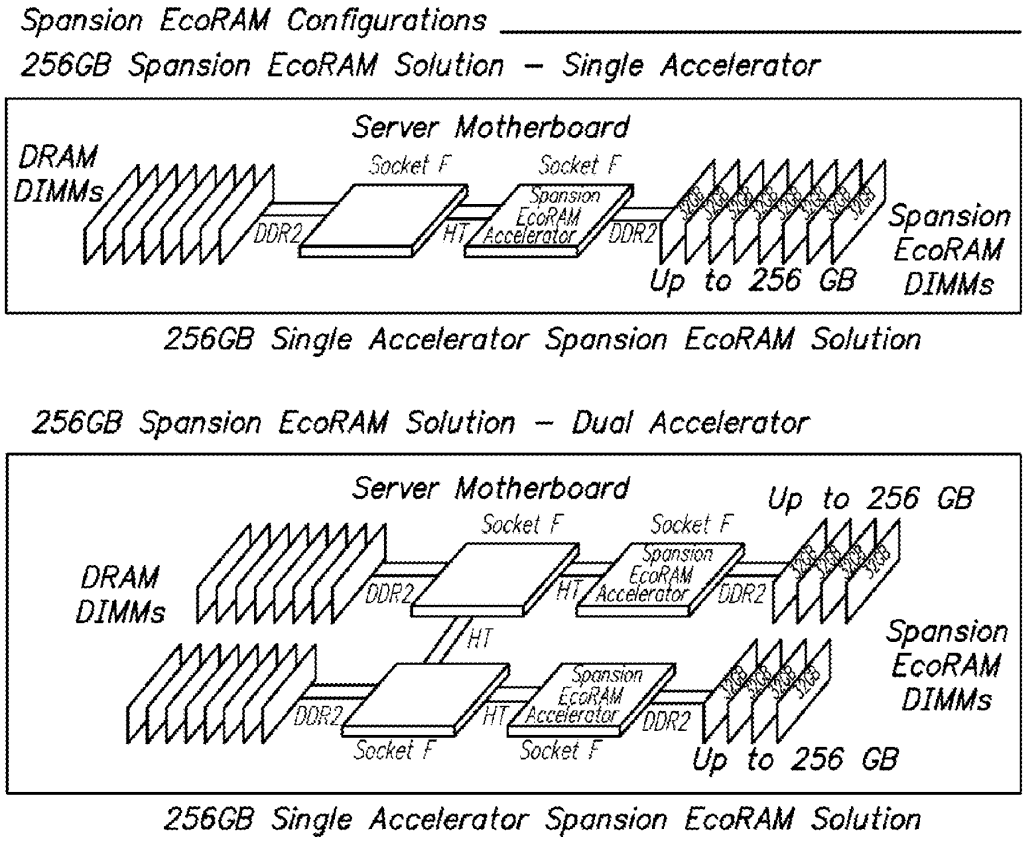


FIG. 2
(PRIOR ART)

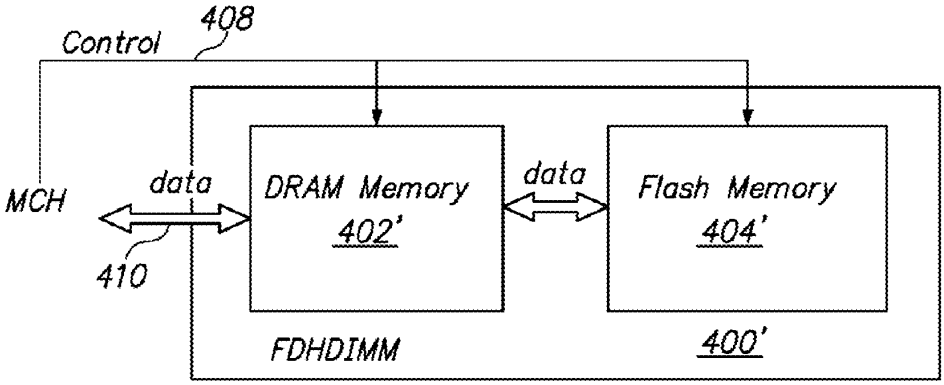


FIG. 4B

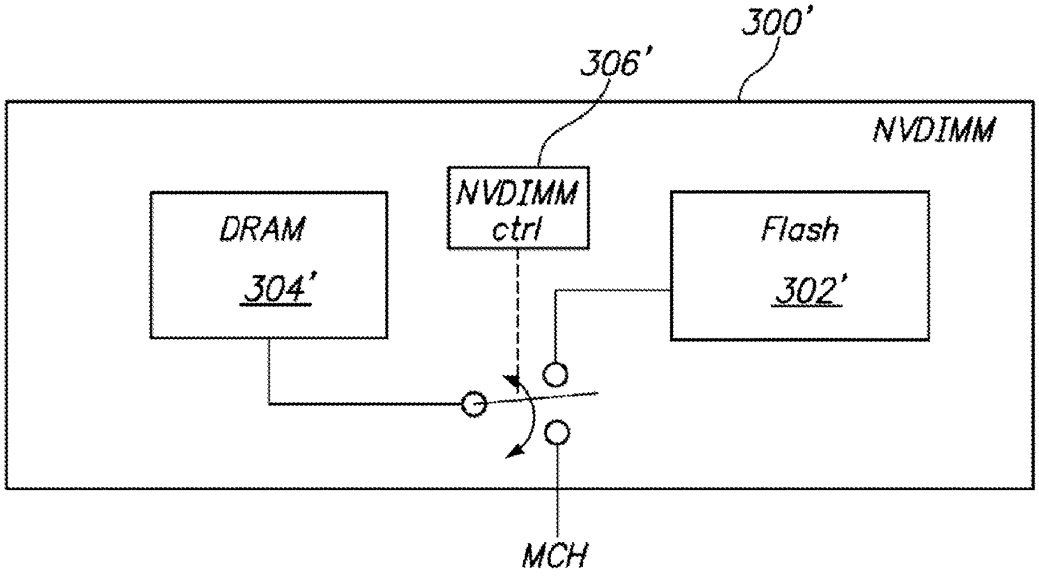

300'
306'
NVDIMM
NVDIMM ctrl
DRAM
304'
Flash
302'
MCH

FIG. 3B

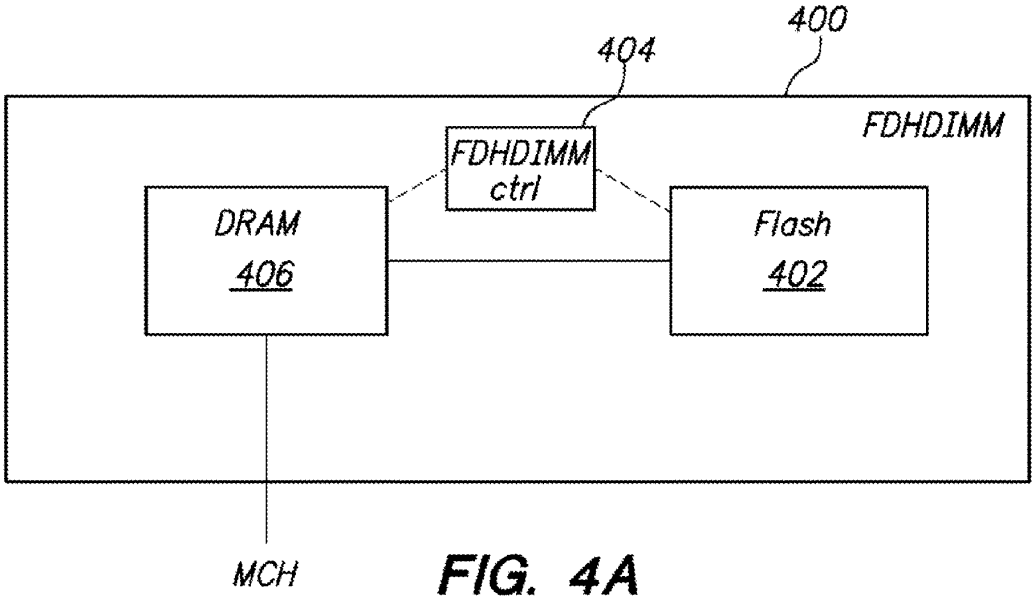

400
404
FDHDIMM
FDHDIMM ctrl
DRAM
406
Flash
402
MCH

FIG. 4A

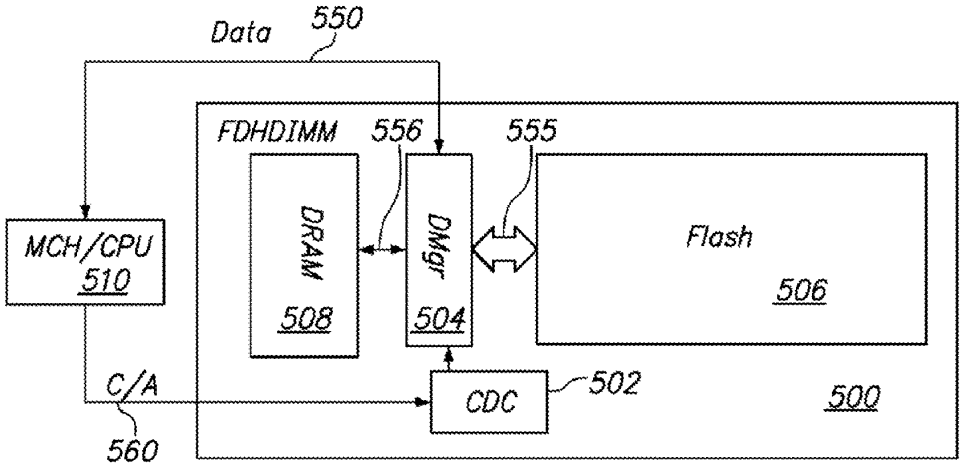

Data
550
FDHDIMM
556
555
DRAM
508
DMgr
504
Flash
506
MCH/CPU
510
C/A
560
CDC
502
500

FIG. 5A

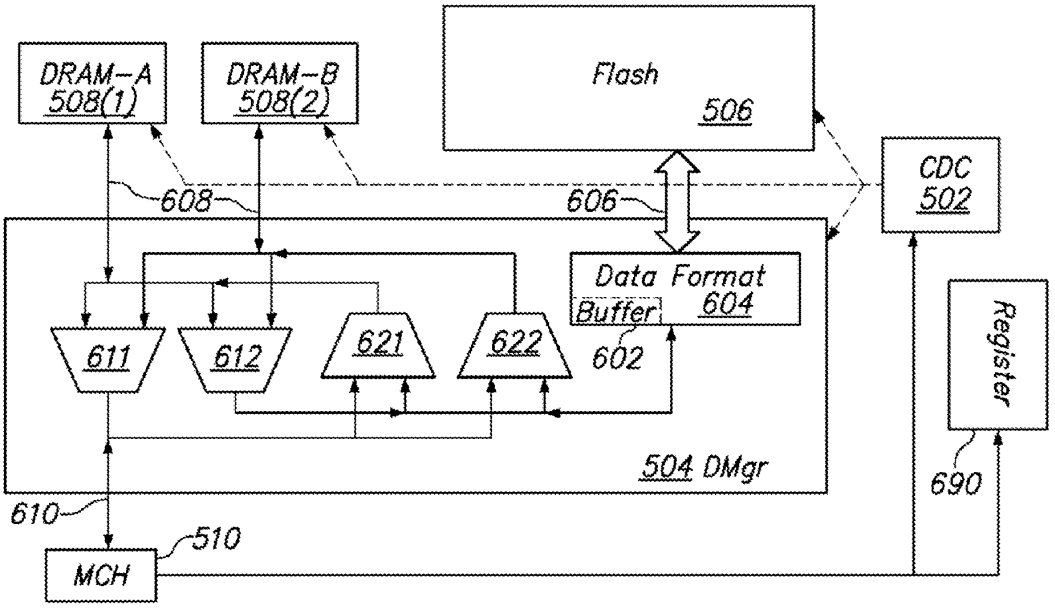

DRAM-A
508(1)
DRAM-B
508(2)
Flash
506
CDC
502
608
606
Data Format
Buffer
604
611
612
621
622
602
Register
504 DMgr
690
610
510
MCH

FIG. 6

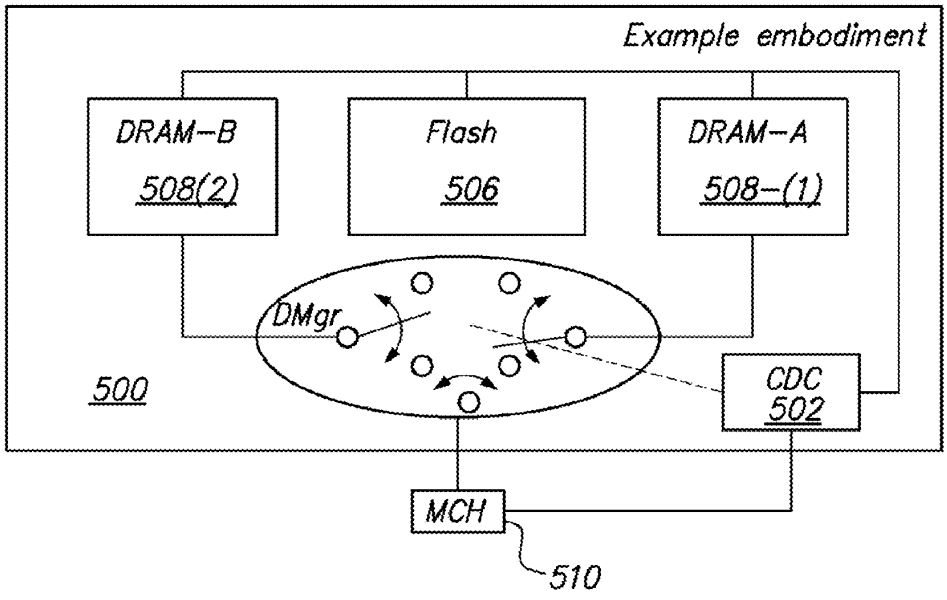

Example embodiment
DRAM–B
508(2)
Flash
506
DRAM–A
508–(1)
DMgr
500
CDC
502
MCH
510

FIG. 5B

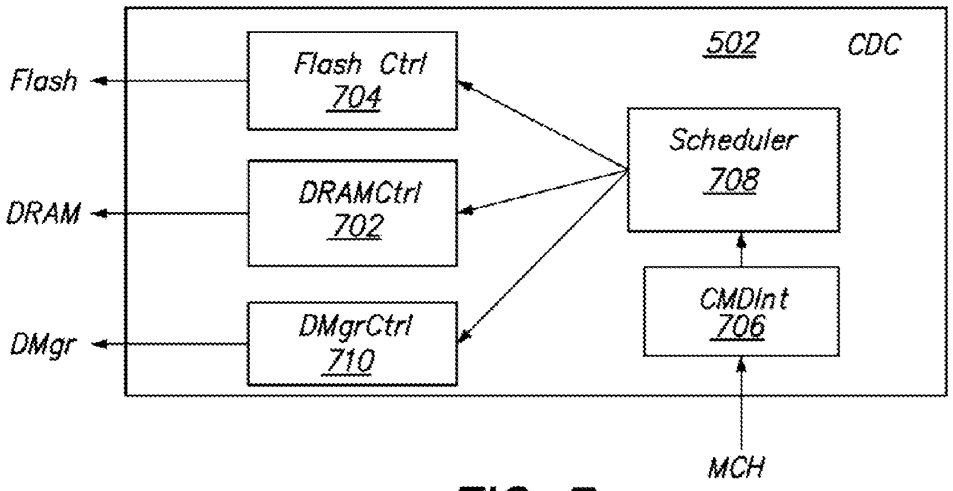

Flash
DRAM
DMgr
Flash Ctrl
704
DRAMCtrl
702
DMgrCtrl
710
502
CDC
Scheduler
708
CMDInt
706
MCH

FIG. 7

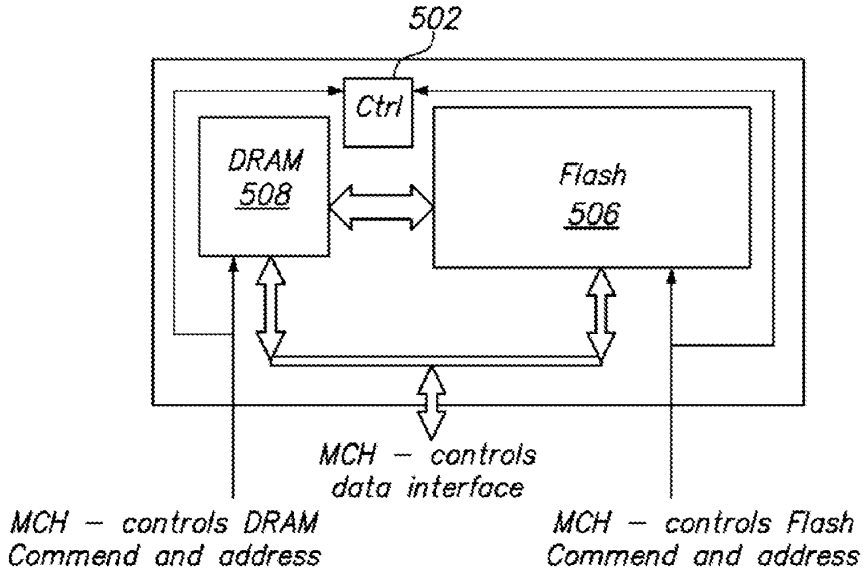

502
Ctrl
DRAM
508
Flash
506
MCH – controls
data interface
MCH – controls DRAM
Commend and address
MCH – controls Flash
Commend and address

FIG. 8A

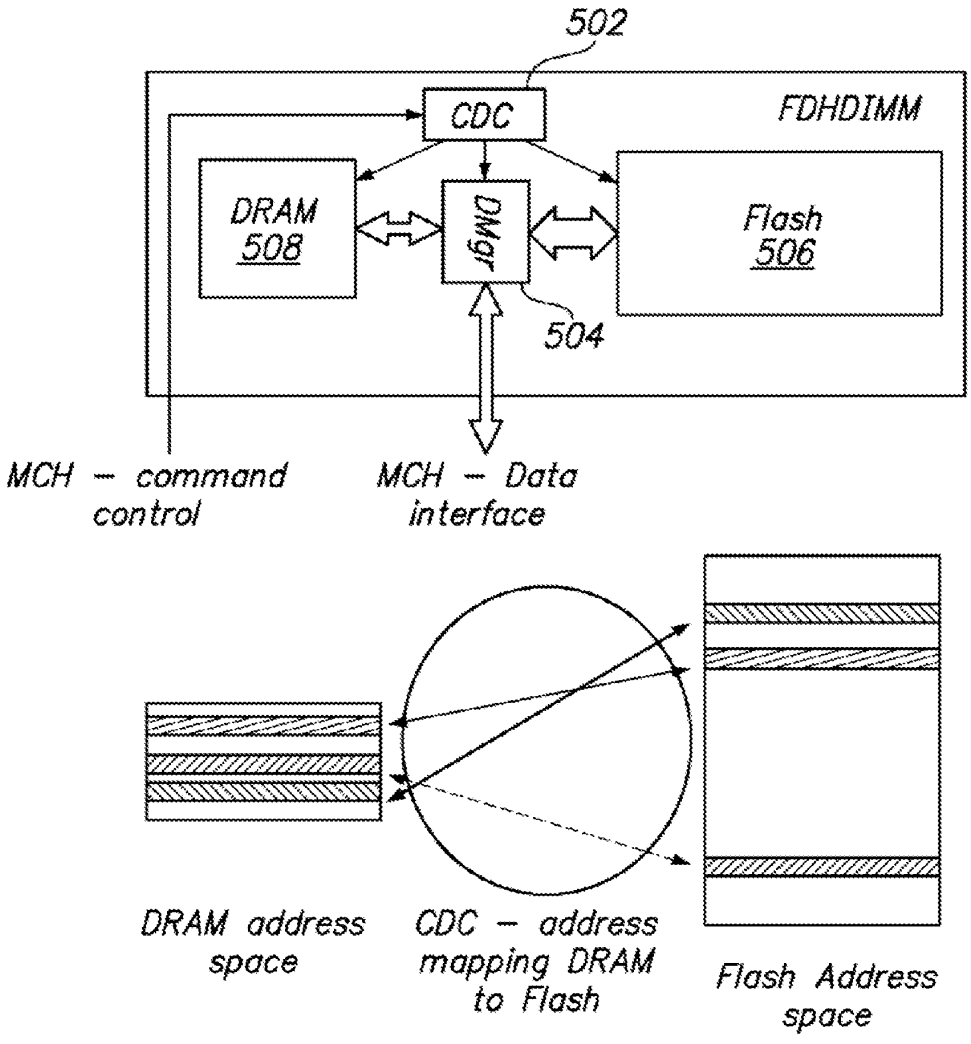

502
CDC
FDHDIMM
DRAM
508
DMgr
Flash
506
504
MCH – command control
MCH – Data interface
DRAM address space
CDC – address mapping DRAM to Flash
Flash Address space

FIG. 8B

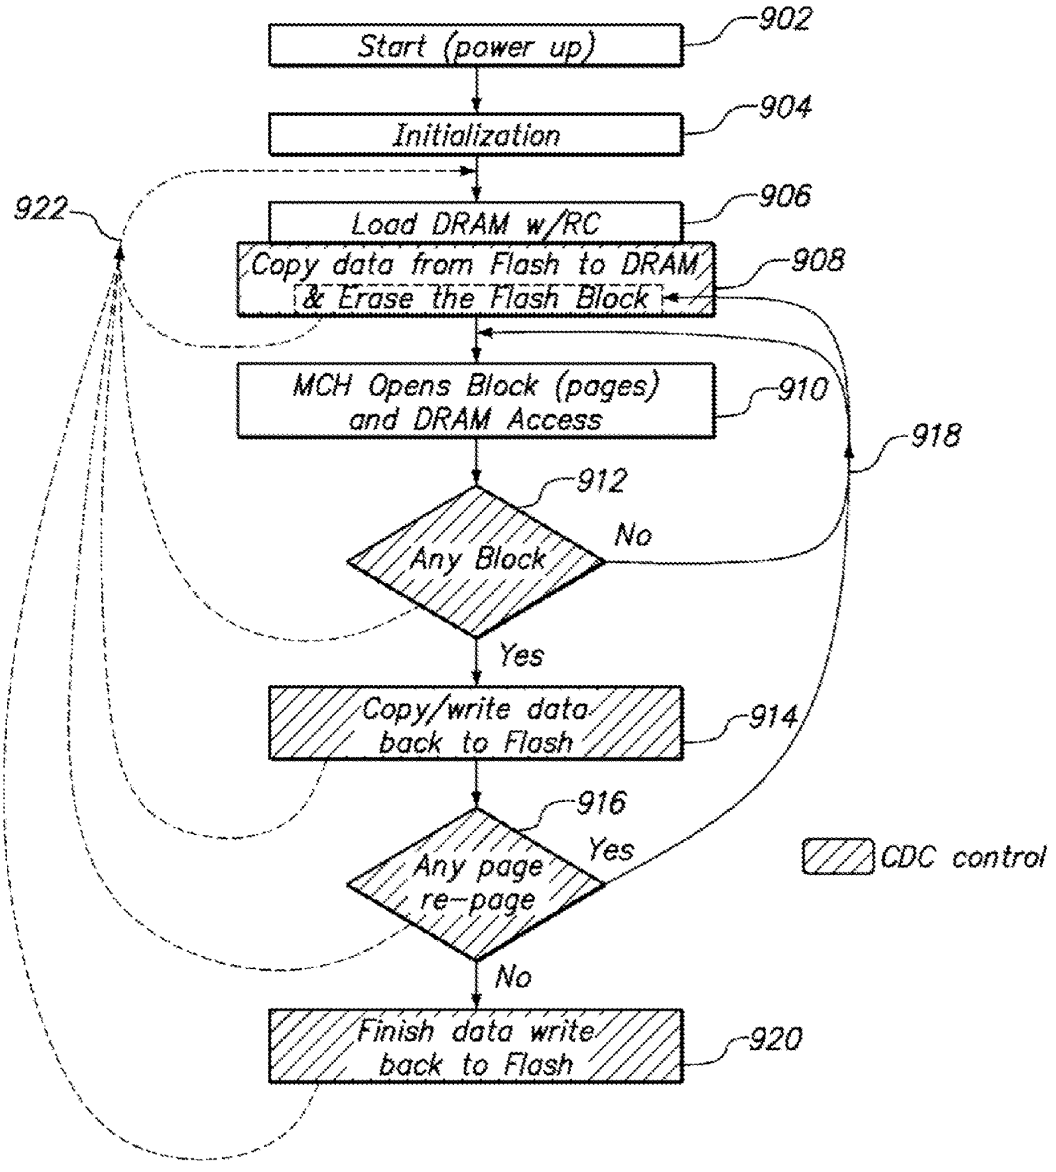

Start (power up)
902
Initialization
904
Load DRAM w/RC
906
Copy data from Flash to DRAM & Erase the Flash Block
908
MCH Opens Block (pages) and DRAM Access
910
Any Block
912
No
Yes
Copy/write data back to Flash
914
Any page re-page
916
Yes
No
Finish data write back to Flash
920
922
918
CDC control

FIG. 9

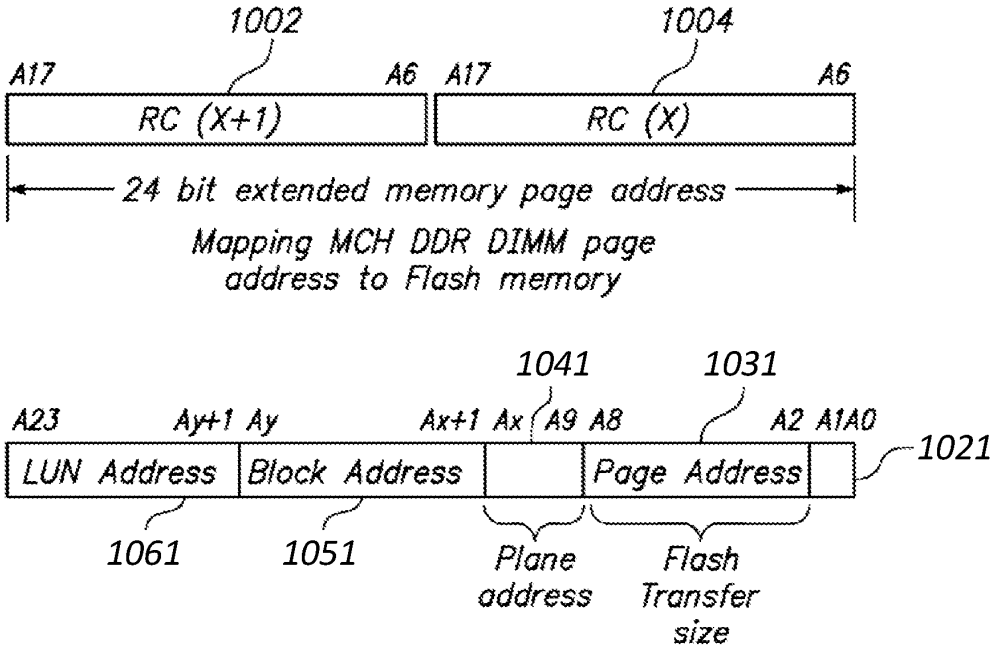

1002
1004
A17
A6
A17
A6
RC (X+1)
RC (X)
24 bit extended memory page address
Mapping MCH DDR DIMM page address to Flash memory
1041
1031
A23
Ay+1
Ay
Ax+1
Ax
A9
A8
A2
A1A0
1021
LUN Address
Block Address
Page Address
1061
1051
Plane address
Flash Transfer size

FIG. 10

DRAM density (GB)	# of blocks per bank	Flash wr-time to rd-time ratio	Avg block use time (sec)	Flash write time (sec)	Max allowed Closed Blk in queue to be written back to Flash
1	250	55	1.00E-03	2.00E-02	0
1	250	55	1.00E-02	2.00E-02	2
1	250	55	2.00E-02	2.00E-02	5
1	250	55	5.00E-02	2.00E-02	11
2	500	55	1.00E-03	2.00E-02	0
2	500	55	1.00E-02	2.00E-02	5
2	500	55	2.00E-02	2.00E-02	9
2	500	55	5.00E-02	2.00E-02	23
4	1000	55	1.00E-03	2.00E-02	1
4	1000	55	1.00E-02	2.00E-02	9
4	1000	55	2.00E-02	2.00E-02	18
4	1000	55	5.00E-02	2.00E-02	45

FIG. 11

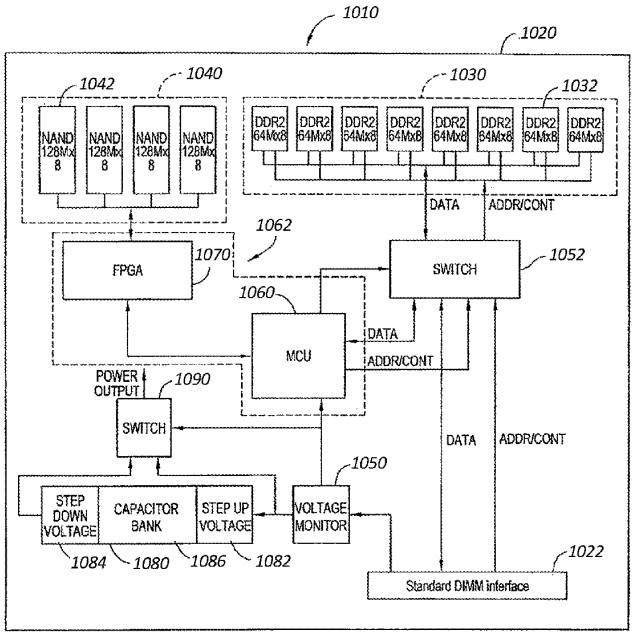


FIG. 12

Samsung Electronics Co., Ltd.
Ex. 1001, p. 16

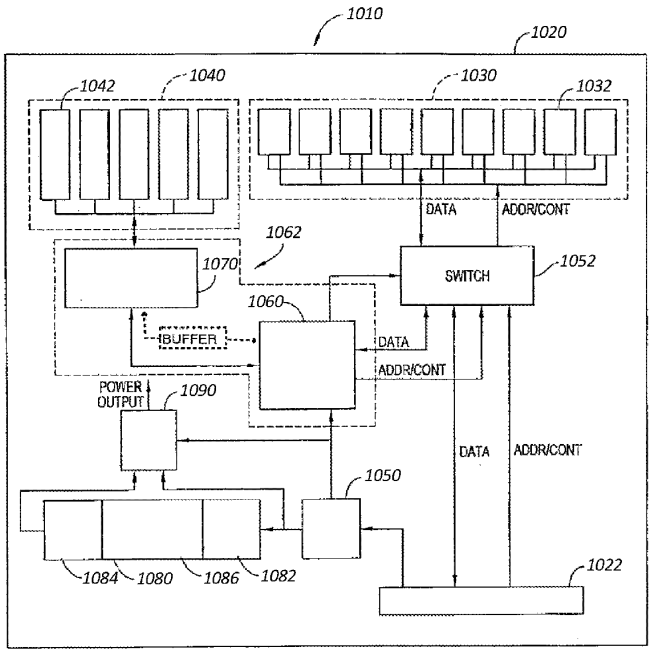


FIG. 13

Samsung Electronics Co., Ltd.
Ex. 1001, p. 17

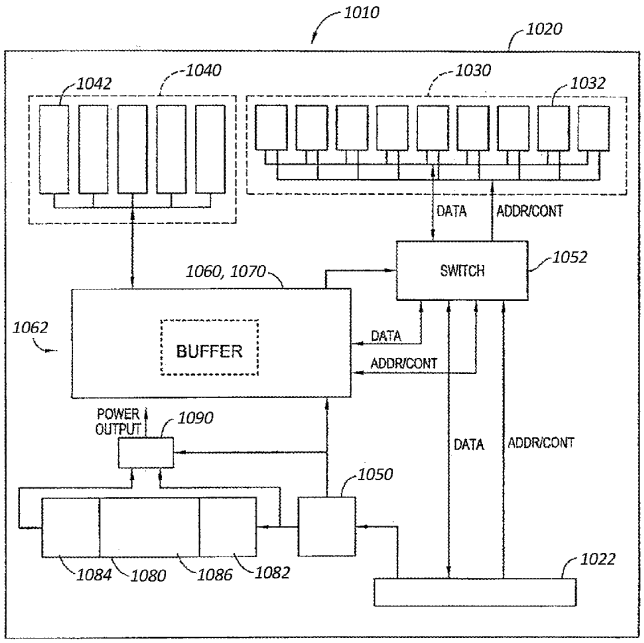


FIG. 14

Samsung Electronics Co., Ltd.
Ex. 1001, p. 18

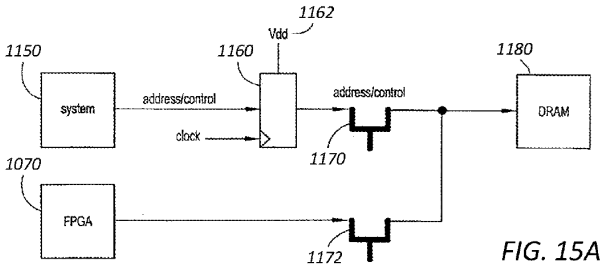


FIG. 15A

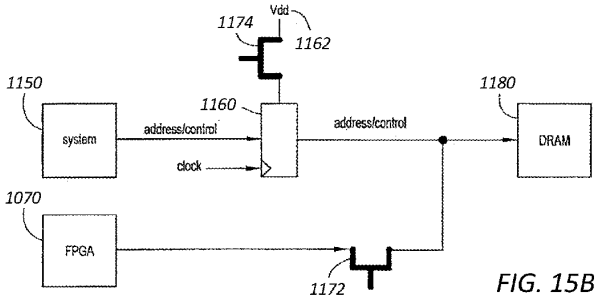


FIG. 15B

Samsung Electronics Co., Ltd.
Ex. 1001, p. 19

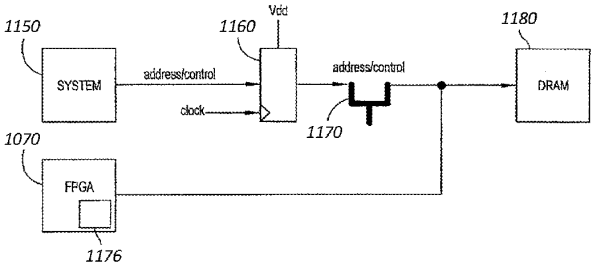


FIG. 15C

Samsung Electronics Co., Ltd.
Ex. 1001, p. 20

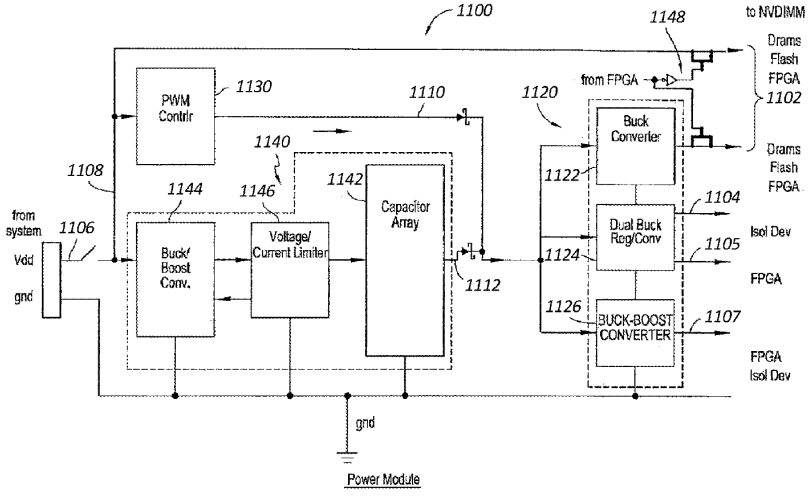


FIG. 16

Samsung Electronics Co., Ltd.
Ex. 1001, p. 21

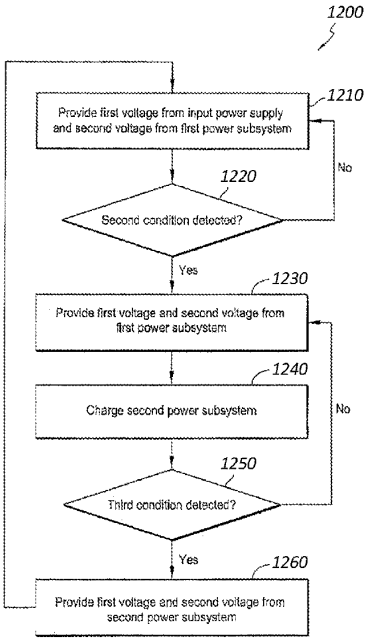


FIG. 17

Samsung Electronics Co., Ltd.
Ex. 1001, p. 22

Samsung Electronics Co., Ltd.
Ex. 1077, p. 757

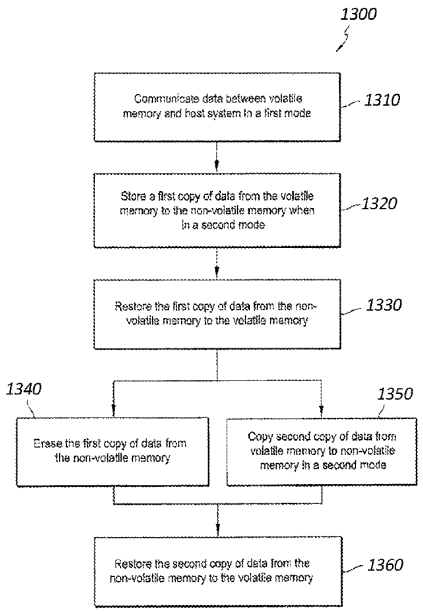


FIG. 18

Samsung Electronics Co., Ltd.
Ex. 1001, p. 23

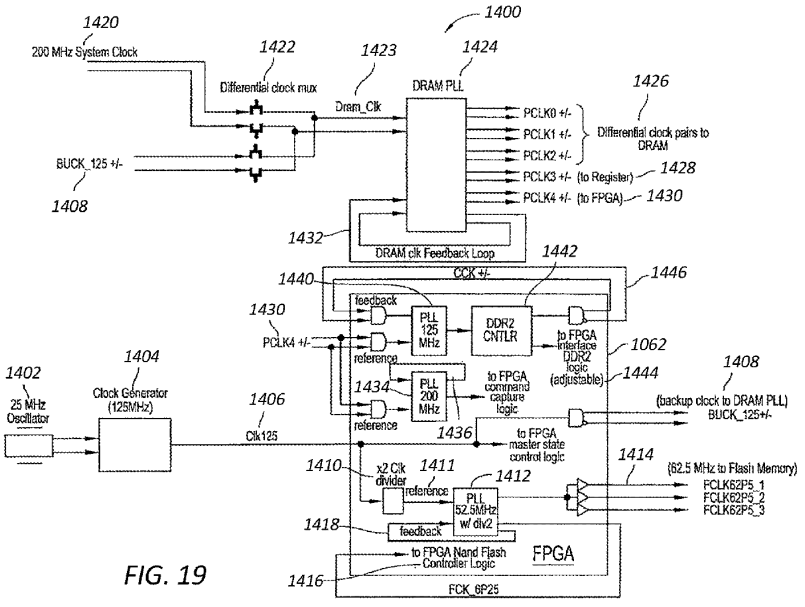


FIG. 19

Samsung Electronics Co., Ltd.
Ex. 1001, p. 24

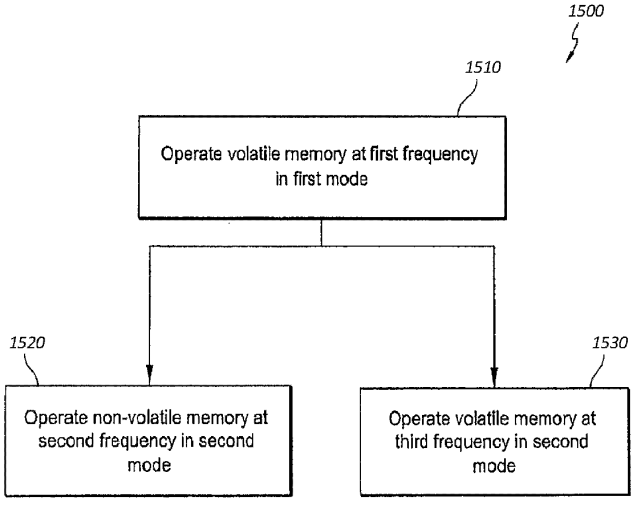


FIG. 20

Samsung Electronics Co., Ltd.
Ex. 1001, p. 25

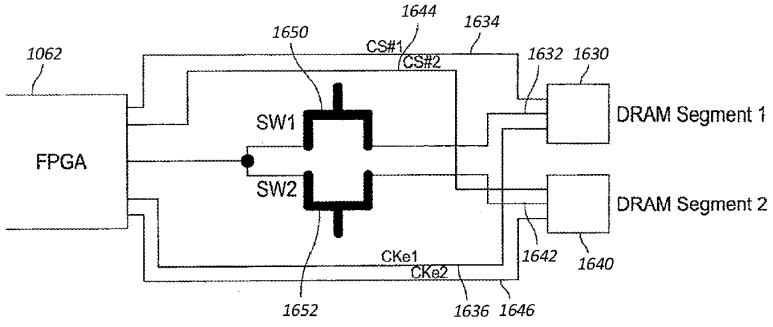


FIG. 21

Samsung Electronics Co., Ltd.
Ex. 1001, p. 26

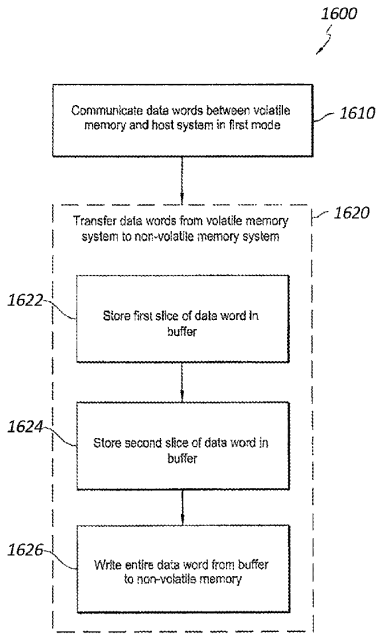


FIG. 22

Samsung Electronics Co., Ltd.
Ex. 1001, p. 27

FLASH-DRAM HYBRID MEMORY MODULE

PRIORITY CLAIM

This application is a continuation of U.S. patent application Ser. No. 17/138,766, filed Dec. 30, 2020, titled "Flash-Dram Hybrid Memory", now U.S. Pat. No. 11,016,918, which is a continuation of U.S. patent application Ser. No. 15/934,416, filed Mar. 23, 2018, titled "Flash-Dram Hybrid Memory Module," which is a continuation of U.S. patent application Ser. No. 14/840,865, filed Aug. 31, 2015, titled "Flash-Dram Hybrid Memory Module," now U.S. Pat. No. 9,928,186, which is a continuation of U.S. patent application Ser. No. 14/489,269, filed Sep. 17, 2014, titled "Flash-Dram Hybrid Memory Module," now U.S. Pat. No. 9,158,684, which is a continuation of U.S. patent application Ser. No. 13/559,476, filed Jul. 26, 2012, titled "Flash-Dram Hybrid Memory Module," now U.S. Pat. No. 8,874,831, which claims the benefit of U. S. Provisional Patent Application No. 61/512,871, filed Jul. 28, 2011, and is a continuation-in-part of U.S. patent application Ser. No. 12/240,916, filed Sep. 29, 2008, titled "Non-Volatile Memory Module," now U.S. Pat. No. 8,301,833, which is a continuation of U.S. patent application Ser. No. 12/131,873, filed Jun. 2, 2008, which claims the benefit of U. S. Provisional Patent Application No. 60/941,586, filed Jun. 1, 2007, the contents of all of which are incorporated herein by reference in their entirety.

This application may be considered related to U.S. patent application Ser. No. 14/173,242, titled "Isolation Switching For Backup Of Registered Memory," filed Feb. 5, 2014, which is a continuation of U.S. patent application Ser. No. 13/905,053, titled "Isolation Switching For Backup Of Registered Memory," filed May 29, 2013, now U.S. Pat. No. 8,677,060, issued Mar. 18, 2014, which is a continuation of U.S. patent application Ser. No. 13/536,173, titled "Data Transfer Scheme For Non-Volatile Memory Module," filed Jun. 28, 2012, now U.S. Pat. No. 8,516,187, issued Aug. 20, 2013, which is a divisional of U.S. patent application Ser. No. 12/240,916, titled "Non-Volatile Memory Module," filed Sep. 29, 2008, now U.S. Pat. No. 8,301,833, issued Oct. 30, 2012, which is a continuation of U.S. patent application Ser. No. 12/131,873, filed Jun. 2, 2008, now abandoned, which claims the benefit of U.S. Provisional Application No. 60/941,586, filed Jun. 1, 2007, the contents of which are incorporated by reference herein in their entirety.

This application may also be considered related to U.S. patent application Ser. No. 15/000,834, filed Jan. 19, 2016 (abandoned), which is a continuation of U.S. patent application Ser. No. 14/489,332, filed Sep. 17, 2014, now U.S. Pat. No. 9,269,437, which is a continuation of U.S. patent application Ser. No. 14/173,219, filed Feb. 5, 2014, now U.S. Pat. No. 8,904,099, which is a continuation of U.S. patent application Ser. No. 13/905,048, filed May 29, 2013, now U.S. Pat. No. 6,671,243, which is a continuation U.S. patent application Ser. No. 13/536,173 above.

This application may also be considered related to U.S. patent application Ser. No. 15/924,866, (abandoned), which is a continuation of U.S. patent application Ser. No. 14/489,281, filed Sep. 17, 2014, now U.S. Pat. No. 9,921,762, which is a continuation of U.S. patent application Ser. No. 13/625,563, filed Sep. 24, 2012, now U.S. Pat. No. 8,904,098, which claims the benefit of U.S. Provisional Application No. 61/583,775, filed Sep. 23, 2011.

TECHNICAL FIELD

The present disclosure relates generally to computer memory devices, and more particularly, to devices that employ different types of memory devices such as combinations of Flash and random access memories.

BACKGROUND

As technology advances and the usage of portable computing devices, such as tablet notebook computers, increases, more data needs to be transferred among data centers and to/from end users. In many cases, data centers are built by clustering multiple servers that are networked to increase performance.

Although there are many types of networked servers that are specific to the types applications envisioned, the basic concept is generally to increase server performance by dynamically allocating computing and storage resources. In recent years, server technology has evolved to be specific to particular applications such as 'finance transactions' (for example, point-of-service, inter-bank transaction, stock market transaction), 'scientific computation' (for example, fluid dynamic for automobile and ship design, weather prediction, oil and gas expeditions), 'medical diagnostics' (for example, diagnostics based on the fuzzy logic, medical data processing), 'simple information sharing and searching' (for example, web search, retail store website, company home page), 'email' (information distribution and archive), 'security service', 'entertainment' (for example, video-on-demand), and so on. However, all of these applications suffer from the same information transfer bottleneck due to the inability of a high speed CPU (central processing unit) to efficiently transfer data in and out of relatively slower speed storage or memory subsystems, particularly since data transfers typically pass through the CPU input/output (I/O) channels.

The data transfer limitations by the CPU are exemplified by the arrangement shown in FIG. **1**, and apply to data transfers between main storage (for example the hard disk (HD) or solid state drive (SSD) and the memory subsystems (for example DRAM DIMM (Dynamic Random Access Memory Dual In-line Memory Module) connected to the front side bus (FSB)). In arrangements such as that of FIG. **1**, the SSD/HD and DRAM DIMM of a conventional memory arrangement are connected to the CPU via separate memory control ports (not shown). FIG. **1** specifically shows, through the double-headed arrow, the data flow path between the computer or server main storage (SSD/HD) to the DRAM DIMMs. Since the SSD/HD data I/O and the DRAM DIMM data I/O are controlled by the CPU, the CPU needs to allocate its process cycles to control these I/Os, which may include the IRQ (Interrupt Request) service which the CPU performs periodically. As will be appreciated, the more time a CPU allocates to controlling the data transfer traffic, the less time the CPU has to perform other tasks. Therefore, the overall performance of a server will deteriorate with the increased amount of time the CPU has to expend in performing data transfer.

There have been various approaches to increase the data transfer throughput rates from/to the main storage, such as SSD/HD, to local storage, such as DRAM DIMM. In one example as illustrated in FIG. **2**, EcoRAM™ developed by Spansion provides a storage SSD based system that assumes a physical form factor of a DIMM. The EcoRAM™ is populated with Flash memories and a relatively small memory capacity using DRAMs which serve as a data buffer. This arrangement is capable of delivering higher throughput rate than a standard SSD based system since the EcoRAM™ is connected to the CPU (central processing unit) via a high speed interface, such as the HT (Hyper

Transport) interface, while an SSD/HD is typically connected via SATA (serial AT attachment), USB (universal serial bus), or PCI Express (peripheral component interface express). For example, the read random access throughput rate of EcoRAM™ is near 3 GB/s compared with 400 MB/s for a NAND SSD memory subsystem using the standard PCI Express-based. This is a 7.5× performance improvement. However, the performance improvement for write random access throughput rate is less than 2× (197 MBs for the EcoRAM vs. 104 MBs for NAND SSD). This is mainly due to the fact that the write speed is cannot be faster than the NAND Flash write access time. FIG. **2** is an example of EcoRAM™ using SSD with the form factor of a standard DIMM such that it can be connected to the FSB (front side bus). However, due to the interface protocol difference between DRAM and Flash, an interface device, EcoRAM Accelerator™), which occupies one of the server's CPU sockets is used, and hence further reducing server's performance by reducing the number of available CPU sockets available, and in turn reducing the overall computation efficiency. The server's performance will further suffer due to the limited utilization of the CPU bus due to the large difference in the data transfer throughput rate between read and write operations.

The EcoRAM™ architecture enables the CPU to view the Flash DIMM controller chip as another processor with a large size of memory available for CPU access.

In general, the access speed of a Flash based system is limited by four items: the read/write speed of the Flash memory, the CPU's FSB bus speed and efficiency, the Flash DIMM controller's inherent latency, and the HT interconnect speed and efficiency which is dependent on the HT interface controller in the CPU and Flash DIMM controller chip.

The published results indicate that these shortcomings are evident in that the maximum throughput rate is 1.56 GB s for the read operation and 104 MBs for the write operation. These access rates are 25% of the DRAM read access speed, and 1.7% of the DRAM access speed at 400 MHz operation. The disparity in the access speed (15 to 1) between the read operation and write operation highlight a major disadvantage of this architecture. The discrepancy of the access speed between this type of architecture and JEDEC standard DRAM DIMM is expected to grow wider as the DRAM memory technology advances much faster than the Flash memory.

Certain types of memory modules comprise a plurality of dynamic random-access memory (DRAM) devices mounted on a printed circuit board (PCB). These memory modules are typically mounted in a memory slot or socket of a computer system (e.g., a server system or a personal computer) and are accessed by the computer system to provide volatile memory to the computer system.

Volatile memory generally maintains stored information only when it is powered. Batteries have been used to provide power to volatile memory during power failures or interruptions. However, batteries may require maintenance, may need to be replaced, are not environmentally friendly, and the status of batteries can be difficult to monitor.

Non-volatile memory can generally maintain stored information while power is not applied to the non-volatile memory. In certain circumstances, it can therefore be useful to backup volatile memory using non-volatile memory.

OVERVIEW

Described herein is a memory module couplable to a memory controller of a host system. The memory module

includes a non-volatile memory subsystem, a data manager coupled to the non-volatile memory subsystem, a volatile memory subsystem coupled to the data manager and operable to exchange data with the non-volatile memory subsystem by way of the data manager, and a controller operable to receive commands from the memory controller and to direct (i) operation of the non-volatile memory subsystem, (ii) operation of the volatile memory subsystem, and (iii) transfer of data between any two or more of the memory controller, the volatile memory subsystem, and the non-volatile memory subsystem based on at least one received command from the memory controller.

Also described herein is a method for managing a memory module by a memory controller, the memory module including volatile and non-volatile memory subsystems. The method includes receiving control information from the memory controller, wherein the control information is received using a protocol of the volatile memory subsystem. The method further includes identifying a data path to be used for transferring data to or from the memory module using the received control information, and using a data manager and a controller of the memory module to transfer data between any two or more of the memory controller, the volatile memory subsystem, and the non-volatile memory subsystem based on at least one of the received control information and the identified data path.

Also described herein is a memory module wherein the data manager is operable to control one or more of data flow rate, data transfer size, data buffer size, data error monitoring, and data error correction in response to receiving at least one of a control signal and control information from the controller.

Also described herein is a memory module wherein the data manager controls data traffic between any two or more of the memory controller, the volatile memory subsystem, and the non-volatile memory subsystem based on instructions received from the controller.

Also described herein is a memory module wherein data traffic control relates to any one or more of data flow rate, data transfer size, data buffer size, data transfer bit width, formatting information, direction of data flow, and the starting time of data transfer.

Also described herein is a memory module wherein the controller configures at least one of a first memory address space of the volatile memory subsystem and a second memory address space of the non-volatile memory subsystem in response to at least one of a received command from the memory controller and memory address space initialization information of the memory module.

Also described herein is a memory module wherein the data manager is configured as a bi-directional data transfer fabric having two or more sets of data ports coupled to any one of the volatile and non-volatile memory subsystems.

Also described herein is a memory module wherein at least one of the volatile and non-volatile memory subsystems comprises one or more memory segments.

Also described herein is a memory module wherein each memory segment comprises at least one memory circuit, memory device, or memory die.

Also described herein is a memory module wherein the volatile memory subsystem comprises DRAM memory.

Also described herein is a memory module wherein the non-volatile memory subsystem comprises flash memory.

Also described herein is a memory module wherein at least one set of data ports is operated by the data manager to

independently and/or concurrently transfer data to or from one or more memory segments of the volatile or non-volatile memory subsystems.

Also described herein is a memory module wherein the data manager and controller are configured to effect data transfer between the memory controller and the non-volatile memory subsystem in response to memory access commands received by the controller from the memory controller.

Also described herein is a memory module wherein the volatile memory subsystem is operable as a buffer for the data transfer between the memory controller and non-volatile memory.

Also described herein is a memory module wherein the data manager further includes a data format module configured to format data to be transferred between any two or more of the memory controller, the volatile memory subsystem, and the non-volatile memory subsystem based on control information received from the controller.

Also described herein is a memory module wherein the data manager further includes a data buffer for buffering data delivered to or from the non-volatile memory subsystem.

Also described herein is a memory module wherein the controller is operable to perform one or more of memory address translation, memory address mapping, address domain conversion, memory access control, data error correction, and data width modulation between the volatile and non-volatile memory subsystems.

Also described herein is a memory module wherein the controller is configured to effect operation with the host system in accordance with a prescribed protocol.

Also described herein is a memory module wherein the prescribed protocol is selected from one or more of DDR, DDR2, DDR3, and DDR4 protocols.

Also described herein is a memory module wherein the controller is operable to configure memory space in the memory module based on at least one of a command received from the memory controller, a programmable value written into a register, a value corresponding to a first portion of the volatile memory subsystem, a value corresponding to a first portion of the non-volatile memory subsystem, and a timing value.

Also described herein is a memory module wherein the controller configures the memory space of the memory module using at least a first portion of the volatile memory subsystem and a first portion of the non-volatile memory subsystem, and the controller presents a unified memory space to the memory controller.

Also described herein is a memory module wherein the controller configures the memory space in the memory module using partitioning instructions that are application-specific.

Also described herein is a memory module wherein the controller is operable to copy booting information from the non-volatile to the volatile memory subsystem during power up.

Also described herein is a memory module wherein the controller includes a volatile memory control module, a non-volatile memory control module, data manager control module, a command interpreter module, and a scheduler module.

Also described herein is a memory module wherein commands from the volatile memory control module to the volatile memory subsystem are subordinated to commands from the memory controller to the controller.

Also described herein is a memory module wherein the controller effects pre-fetching of data from the non-volatile to the volatile memory.

Also described herein is a memory module wherein the pre-fetching is initiated by the memory controller writing an address of requested data into a register of the controller.

Also described herein is a memory module wherein the controller is operable to initiate a copy operation of data of a closed block in the volatile memory subsystem to a target block in the non-volatile memory subsystem.

Also described herein is a memory module wherein, if the closed block is re-opened, the controller is operable to abort the copy operation and to erase the target block from the non-volatile memory subsystem.

Also described herein is a method for managing a memory module wherein the transfer of data includes a bidirectional transfer of data between the non-volatile and the volatile memory subsystems.

Also described herein is a method for managing a memory module further comprising operating the data manager to control one or more of data flow rate, data transfer size, data width size, data buffer size, data error monitoring, data error correction, and the starting time of the transfer of data.

Also described herein is a method for managing a memory module further comprising operating the data manager to control data traffic between the memory controller and at least one of the volatile and non-volatile memory subsystems.

Also described herein is a method for managing a memory module wherein data traffic control relates to any one or more of data transfer size, formatting information, direction of data flow, and the starting time of the transfer of data.

Also described herein is a method for managing a memory module wherein data traffic control by the data manager is based on instructions received from the controller.

Also described herein is a method for managing a memory module further comprising operating the data manager as a bi-directional data transfer fabric with two or more sets of data ports coupled to any one of the volatile and non-volatile memory subsystems.

Also described herein is a method for managing a memory module wherein at least one of the volatile and non-volatile memory subsystems comprises one or more memory segments.

Also described herein is a method for managing a memory module wherein each memory segment comprises at least one memory circuit, memory device, or memory die.

Also described herein is a method for managing a memory module wherein the volatile memory subsystem comprises DRAM memory.

Also described herein is a method for managing a memory module wherein the non-volatile memory subsystem comprises Flash memory.

Also described herein is a method for managing a memory module further comprising operating the data ports to independently and/or concurrently transfer data to or from one or more memory segments of the volatile or non-volatile memory subsystems.

Also described herein is a method for managing a memory module further comprising directing transfer of data bi-directionally between the volatile and non-volatile memory subsystems using the data manager and in response to memory access commands received by the controller from the memory controller.

Also described herein is a method for managing a memory module further comprising buffering the data transferred between the memory controller and non-volatile memory subsystem using the volatile memory subsystem.

Also described herein is a method for managing a memory module further comprising using the controller to perform one or more of memory address translation, memory address mapping, address domain conversion, memory access control, data error correction, and data width modulation between the volatile and non-volatile memory subsystems.

Also described herein is a method for managing a memory module further comprising using the controller to effect communication with a host system by the volatile memory subsystem in accordance with a prescribed protocol.

Also described herein is a method for managing a memory module wherein the prescribed protocol is selected from one or more of DDR, DDR2, DDR3, and DDR4 protocols.

Also described herein is a method for managing a memory module further comprising using the controller to configure memory space in the memory module based on at least one of a command received from the memory controller, a programmable value written into a register, a value corresponding to a first portion of the volatile memory subsystem, a value corresponding to a first portion of the non-volatile memory subsystem, and a timing value.

Also described herein is a method for managing a memory module wherein the controller configures the memory space of the memory module using at least a first portion of the volatile memory subsystem and a first portion of the non-volatile memory subsystem, and the controller presents a unified memory space to the memory controller.

Also described herein is a method for managing a memory module wherein the controller configures the memory space in the memory module using partitioning instructions that are application-specific.

Also described herein is a method for managing a memory module further comprising using the controller to copy booting information from the non-volatile to the volatile memory subsystem during power up.

Also described herein is a method for managing a memory module wherein the controller includes a volatile memory control module, the method further comprising generating commands by the volatile memory control module in response to commands from the memory controller, and transmitting the generated commands to the volatile memory subsystem.

Also described herein is a method for managing a memory module further comprising pre-fetching of data from the non-volatile memory subsystem to the volatile memory subsystem.

Also described herein is a method for managing a memory module wherein the pre-fetching is initiated by the memory controller writing an address of requested data into a register of the controller.

Also described herein is a method for managing a memory module further comprising initiating a copy operation of data of a closed block in the volatile memory subsystem to a target block in the non-volatile memory subsystem.

Also described herein is a method for managing a memory module further comprising aborting the copy operation when the closed block of the volatile memory subsystem is re-opened, and erasing the target block in the non-volatile memory subsystem.

Also described herein is a memory system having a volatile memory subsystem, a non-volatile memory subsystem, a controller coupled to the non-volatile memory subsystem, and a circuit coupled to the volatile memory subsystem, to the controller, and to a host system. In a first mode of operation, the circuit is operable to selectively isolate the controller from the volatile memory subsystem, and to selectively couple the volatile memory subsystem to the host system to allow data to be communicated between the volatile memory subsystem and the host system. In a second mode of operation, the circuit is operable to selectively couple the controller to the volatile memory subsystem to allow data to be communicated between the volatile memory subsystem and the nonvolatile memory subsystem using the controller, and the circuit is operable to selectively isolate the volatile memory subsystem from the host system.

Also described herein is a method for operating a memory system. The method includes coupling a circuit to a host system, a volatile memory subsystem, and a controller, wherein the controller is coupled to a non-volatile memory subsystem. In a first mode of operation that allows data to be communicated between the volatile memory subsystem and the host system, the circuit is used to (i) selectively isolate the controller from the volatile memory subsystem, and (ii) selectively couple the volatile memory subsystem to the host system. In a second mode of operation that allows data to be communicated between the volatile memory subsystem and the nonvolatile memory subsystem via the controller, the circuit is used to (i) selectively couple the controller to the volatile memory subsystem, and (ii) selectively isolate the volatile memory subsystem from the host system.

Also described herein is a nontransitory computer readable storage medium storing one or more programs configured to be executed by one or more computing devices. The programs, when executing on the one or more computing devices, cause a circuit that is coupled to a host system, to a volatile memory subsystem, and to a controller that is coupled to a nonvolatile memory subsystem, to perform a method in which, in a first mode of operation that allows data to be communicated between the volatile memory subsystem and the host system, operating the circuit to (i) selectively isolate the controller from the volatile memory subsystem, and (ii) selectively couple the volatile memory subsystem to the host system. In a second mode of operation that allows data to be communicated between the volatile memory subsystem and the nonvolatile memory subsystem via the controller, operating the circuit to (i) selectively couple the controller to the volatile memory subsystem, and (ii) selectively isolate the volatile memory subsystem from the host system.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated into and constitute a part of this specification, illustrate one or more examples of embodiments and, together with the description of example embodiments, serve to explain the principles and implementations of the embodiments.

In the drawings:

FIG. **1** is a block diagram illustrating the path of data transfer, via a CPU, of a conventional memory arrangement;

FIG. **2** is a block diagram of a known EcoRAM™ architecture;

FIGS. **3**A and **3**B are block diagrams of a non-volatile memory DIMM or NVDIMM;

FIGS. **4**A and **4**B are block diagrams of a Flash-DRAM hybrid DIMM or FDHDIMM;

FIG. **5**A is a block diagram of a memory module **500** in accordance with certain embodiments described herein;

FIG. **5**B is a block diagram showing some functionality of a memory module such as that shown in FIG. **5**A;

FIG. **6** is a block diagram showing some details of the data manager (DMgr);

FIG. **7** is a functional block diagram of the on-module controller (CDC);

FIG. **8**A is a block diagram showing more details of the prior art Flash-DRAM hybrid DIMM (FDHDIMM) of FIGS. **4**A and **4**B;

FIG. **8**B is a block diagram of a Flash-DRAM hybrid DIMM (FDHDIMM) in accordance with certain embodiments disclosed herein;

FIG. **9** is a flow diagram directed to the transfer of data from Flash memory to DRAM memory and vice versa in an exemplary FDHDIMM;

FIG. **10** is a block diagram showing an example of mapping of DRAM address space to Flash memory address space; and

FIG. **11** is a table showing estimates of the maximum allowed closed blocks in a queue to be written back to Flash memory for different DRAM densities using various average block use time.

FIG. **12** is a block diagram of an example memory system compatible with certain embodiments described herein.

FIG. **13** is a block diagram of an example memory module with ECC (error-correcting code) having a volatile memory subsystem with nine volatile memory elements and a non-volatile memory subsystem with five non-volatile memory elements in accordance with certain embodiments described herein.

FIG. **14** is a block diagram of an example memory module having a microcontroller unit and logic element integrated into a single device in accordance with certain embodiments described herein.

FIGS. **15**A-**15**C schematically illustrate example embodiments of memory systems having volatile memory subsystems comprising registered dual in-line memory modules in accordance with certain embodiments described herein.

FIG. **16** schematically illustrates an example power module of a memory system in accordance with certain embodiments described herein.

FIG. **17** is a flowchart of an example method of providing a first voltage and a second voltage to a memory system including volatile and non-volatile memory subsystems.

FIG. **18** is a flowchart of an example method of controlling a memory system operatively coupled to a host system and which includes at least 100 percent more storage capacity in non-volatile memory than in volatile memory.

FIG. **19** schematically illustrates an example clock distribution topology of a memory system in accordance with certain embodiments described herein.

FIG. **20** is a flowchart of an example method of controlling a memory system operatively coupled to a host system, the method including operating a volatile memory subsystem at a reduced rate in a back-up mode.

FIG. **21** schematically illustrates an example topology of a connection to transfer data slices from two DRAM segments of a volatile memory subsystem of a memory system to a controller of the memory system.

FIG. **22** is a flowchart of an example method of controlling a memory system operatively coupled to a host system, the method including backing up and/or restoring a volatile memory subsystem in slices.

DESCRIPTION OF EXAMPLE EMBODIMENTS

Example embodiments are described herein in the context of a system of computers, servers, controllers, memory modules, hard disk drives and software. Those of ordinary skill in the art will realize that the following description is illustrative only and is not intended to be in any way limiting. Other embodiments will readily suggest themselves to such skilled persons having the benefit of this disclosure. Reference will now be made in detail to implementations of the example embodiments as illustrated in the accompanying drawings. The same reference indicators will be used to the extent possible throughout the drawings and the following description to refer to the same or like items.

In the interest of clarity, not all of the routine features of the implementations described herein are shown and described. It will, of course, be appreciated that in the development of any such actual implementation, numerous implementation-specific decisions must be made in order to achieve the developer's specific goals, such as compliance with application- and business-related constraints, and that these specific goals will vary from one implementation to another and from one developer to another. Moreover, it will be appreciated that such a development effort might be complex and time-consuming, but would nevertheless be a routine undertaking of engineering for those of ordinary skill in the art having the benefit of this disclosure.

In accordance with this disclosure, the components, process steps, and/or data structures described herein may be implemented using various types of operating systems, computing platforms, computer programs, and/or general purpose machines. In addition, those of ordinary skill in the art will recognize that devices of a less general purpose nature, such as hardwired devices, field programmable gate arrays (FPGAs), application specific integrated circuits (ASICs), or the like, may also be used without departing from the scope and spirit of the inventive concepts disclosed herein. Where a method comprising a series of process steps is implemented by a computer or a machine and those process steps can be stored as a series of instructions readable by the machine, they may be stored on a tangible medium such as a computer memory device (e.g., ROM (Read Only Memory), PROM (Programmable Read Only Memory), EEPROM (Electrically Eraseable Programmable Read Only Memory), Flash memory, Jump Drive, and the like), magnetic storage medium (e.g., tape, magnetic disk drive, and the like), optical storage medium (e.g., CD-ROM, DVD-ROM, paper card, paper tape and the like) and other types of program memory.

The term "exemplary" where used herein is intended to mean "serving as an example, instance or illustration." Any embodiment described herein as "exemplary" is not necessarily to be construed as preferred or advantageous over other embodiments.

Disclosed herein are arrangements for improving memory access rates and addressing the high disparity (15 to 1 ratio) between the read and write data throughput rates. In one arrangement, a Flash-DRAM-hybrid DIMM (FDHDIMM) with integrated Flash and DRAM is used. Methods for controlling such an arrangement are described.

In certain embodiments, the actual memory density (size or capacity) of the DIMM and/or the ratio of DRAM memory to Flash memory are configurable for optimal use with a particular application (for example, POS, inter-bank transaction, stock market transaction, scientific computation such as fluid dynamics for automobile and ship design, weather prediction, oil and gas expeditions, medical diagnostics such as diagnostics based on the fuzzy logic, medical data processing, simple information sharing and searching such as web search, retail store website, company home

page, email or information distribution and archive, security service, and entertainment such as video-on-demand).

In certain embodiments, the device contains a high density Flash memory with a low density DRAM, wherein the DRAM is used as a data buffer for read/write operation. The Flash serves as the main memory. Certain embodiments described herein overcome the needs of having a long separation period between an Activate command (may be referred to as RAS) and a corresponding read or write command (may be referred to as first CAS command).

In accordance with one embodiment, described with reference to FIGS. **3**A and **3**B, a memory system **300** includes a non-volatile (for example Flash) memory subsystem **302** and a volatile (for example DRAM) memory subsystem **304**. The examples of FIGS. **3**A and **3**B are directed to architectures of a non-volatile DIMM (NVDIMM) NVDIMM system that may use a power subsystem (not shown) that can include a battery or a capacitor as a means for energy storage to copy DRAM memory data into Flash memory when power loss occurs, is detected, or is anticipated to occur during operation. When normal power is restored, a restore NVDIMM operation is initiated and the data stored in the Flash memory is properly restored to the DRAM memory. In this architecture, the density of the Flash is about the same as the DRAM memory size or within a few multiples, although in some applications it may be higher. This type of architecture may also be used to provide non-volatile storage that is connected to the FSB (front side bus) to support RAID (Redundant Array of Independent Disks) based systems or other type of operations. An NVDIMM controller **306** receives and interprets commands from the system memory controller hub (MCH). The NVDIMM controller **306** control the NVDIMM DRAM and Flash memory operations. In FIG. **3**A, the DRAM **304** communicates data with the MCH, while an internal bus **308** is used for data transfer between the DRAM and Flash memory subsystems. In FIG. **3**B, the NVDIMM controller **306'** of NVDIMM **300'** monitors events or commands and enables data transfer to occur in a first mode between the DRAM **304'** and Flash **302'** or in a second mode between the DRAM and the MCH.

In accordance with one embodiment, a general architecture for a Flash and DRAM hybrid DIMM (FDHDIMM) system **400** is shown in FIG. **4**A. The FDHDIMM interfaces with an MCH (memory controller hub) to operate and behave as a high density DIMM, wherein the MCH interfaces with the non-volatile memory subsystem (for example Flash) **402** is controlled by an FDHDIMM controller **404**. Although the MCH interfaces with the Flash via the FDHDIMM controller, the FDHDIMM overall performance is governed by the Flash access time. The volatile memory subsystem (for example DRAM) **406** is primarily used as a data buffer or a temporary storage location such that data from the Flash memory **402** is transferred to the DRAM **406** at the Flash access speed, and buffered or collected into the DRAM **406**, which then transfers the buffered data to the MCH based on the access time of DRAM. Similarly, when the MCH transfers data to the DRAM **406**, the FDHDIMM controller **404** manages the data transfer from the DRAM **406** to the Flash **402**. Since the Flash memory access speed (both read and write) is relatively slower than DRAM, (e.g. for example a few hundred microseconds for read access), the average data throughput rate of FDHDIMM **400** is limited by the Flash access speed. The DRAM **406** serves as a data buffer stage that buffers the MCH read or write data. Thus, the DRAM **406** serves as a temporary storage for the data to be transferred from/to the Flash **402**. Furthermore, in accordance with one embodiment, the MCH recognizes the physical density of an FDHDIMM operating as a high density DIMM as the density of Flash alone.

In accordance with one embodiment, a read operation can be performed by the MCH by sending an activate command (may be simply referred to as RAS, or row address strobe) to the FDHDIMM **400** to conduct a pre-fetch read data operation from the Flash **402** to the DRAM **406**, with the pre-fetch data size being for example a page (1 KB or 2 KB, or may be programmable to any size). The MCH then sends a read command (may be simply referred to as CAS, or column address strobe) to read the data out input of the DRAM. In this embodiment, the data transfer from Flash to DRAM occurs at Flash access speed rates, while data transfer from DRAM to MCH occurs at DRAM access speed rates. In this example, data latency and throughput rates are the same as any DRAM operation as long as the read operations are executed onto the pages that were opened with the activate command previously sent to pre-fetch data from the Flash to DRAM. Thus, a longer separation time period between the RAS (e.g. Activate command) and the first CAS (column address strobe e.g. read or write command) is required to account for the time it takes to pre-fetch data from the Flash to DRAM.

An example of FDHDIMM operating as a DDR DIMM with SSD is shown in FIG. **4**B, wherein the FDHDIMM **400'** supports two different interface interpretations to the MCH. In the first interface interpretation, the MCH views the FDHDIMM **400'** as a combination of DRAM DIMM and SSD (not illustrated). In this mode the MCH needs to manage two address spaces, one for the DRAMs **402'** and one for the Flash **404'**. The MCH is coupled to, and controls, both of the DRAM and Flash memory subsystems. One advantage of this mode is that the CPU does not need to be in the data path when data is moved from DRAM to Flash or from Flash to DRAM. In the second interface interpretation, the MCH views the FDHDIMM **400'** as an on-DIMM Flash with the SSD in an extended memory space that is behind the DRAM space. Thus, in this mode, the MCH physically fetches data from the SSD to the DDR DRAM and then the DRAM sends the data to the MCH. Since all data movement occurs on the FDHDIMM, this mode will provide better performance than if the data were to be moved through or via the CPU.

In accordance with one embodiment and as shown in FIG. **4**B, the FDHDIMM **400'** receives control signals **408** from the MCH, where the control signals may include one or more control signals specifically for the DRAM **402'** operation and one or more control signals specifically for the Flash **404'** operation. In this embodiment, the MCH or CPU is coupled to the FDHDIMM via a single data bus interface **410** which couples the MCH to the DRAM.

FIGS. **5**A and **5**B are block diagrams of a memory module **500** that is couplable to a host system (not shown). The host system may be a server or any other system comprising a memory system controller or an MCH for providing and controlling the read/write access to one or more memory systems, wherein each memory system may include a plurality of memory subsystems, a plurality of memory devices, or at least one memory module. The term "read/write access" means the ability of the MCH to interface with a memory system or subsystem in order to write data into it or read data from it, depending on the particular requirement at a particular time.

In certain embodiments, memory module **500** is a Flash-DRAM hybrid memory subsystem which may be integrated with other components of a host system. In certain embodiments, memory module **500** is a Flash-DRAM hybrid

memory module that has the DIMM (dual-inline memory module) form factor, and may be referred to as a FDHDIMM, although it is to be understood that in both structure and operation it may be different from the FDHDIMM discussed above and described with reference to FIGS. **4**A and **4**B. Memory module **500** includes two on-module intermediary components: a controller and a data manager. These on-module intermediary components may be physically separate components, circuits, or modules, or they may be integrated onto a single integrated circuit or device, or integrated with other memory devices, for example in a three dimensional stack, or in any one of several other possible expedients for integration known to those skilled in the art to achieve a specific design, application, or economic goal. In the case of a DIMM, these on-module intermediary components are an on-DIMM Controller (CDC) **502** and an on-DIMM data manager (DMgr) **504**. While the DIMM form factor will predominate the discussion herein, it should be understood that this is for illustrative purposes only and memory systems using other form factors are contemplated as well. CDC **502** and data manager DMgr **504** are operative to manage the interface between a non-volatile memory subsystem such as a Flash **506**, a volatile memory subsystem such as a DRAM **508**, and a host system represented by MCH **510**.

In certain embodiments, CDC **502** controls the read/write access to/from Flash memory **506** from/to DRAM memory **508**, and to/from DRAM memory from/to MCH **510**. Read/write access between DRAM **508**, Flash **506** and MCH **510** may be referred to herein generally as communication, wherein control and address information C/A **560** is sent from MCH **510** to CDC **502**, and possible data transfers follow as indicated by Data **550**, Data **555**, and/or Data **556**. In certain embodiments, the CDC **502** performs specific functions for memory address transformation, such as address translation, mapping, or address domain conversion, Flash access control, data error correction, manipulation of data width or data formatting or data modulation between the Flash memory and DRAM, and so on. In certain embodiments, the CDC **502** ensures that memory module **500** provides transparent operation to the MCH in accordance with certain industry standards, such as DDR, DDR2, DDR3, DDR4 protocols. In the arrangement shown in FIGS. **5**A and **5**B, there is no direct access from the MCH **510** to the Flash **506** memory subsystem. Thus in accordance with certain embodiments, the Flash access speed has minimal impact on the overall FDHDIMM access speed. In the schematic illustration of FIG. **5**B and in accordance with one embodiment, the CDC controller **502** receives standard DDR commands from the MCH, interprets, and produces commands and/or control signals to control the operation of the Data manager (DMgr), the Flash memory and the DRAM memory. The DMgr controls the data path routing amongst DRAMs, Flash and MCH, as detailed below. The data path routing control signals are independently operated without any exclusivity.

An exemplary role of DMgr **504** is described with reference to FIG. **6**. In certain embodiments and in response to communication from CDC **502**, DMgr **504** provides a variety of functions to control data flow rate, data transfer size, data buffer size, data error monitoring or data error correction. For example, these functions or operations can be performed on-the-fly (while data is being transferred via the DMgr **504**) or performed on buffered or stored data in DRAM or a buffer. In addition, one role of DMgr **504** is to provide interoperability among various memory subsystems or components and/or MCH **510**.

In one embodiment, an exemplary host system operation begins with initialization. The CDC **502** receives a first command from the MCH **510** to initialize FDHDIMM **500** using a certain memory space. The memory space as would be controlled by MCH **510** can be configured or programmed during initialization or after initialization has completed. The MCH **510** can partition or parse the memory space in various ways that are optimized for a particular application that the host system needs to run or execute. In one embodiment, the CDC **502** maps the actual physical Flash **506** and DRAM **508** memory space using the information sent by MCH **510** via the first command. In one embodiment, the CDC **502** maps the memory address space of any one of the Flash **506** and DRAM **508** memory subsystems using memory address space information that is received from the host system, stored in a register within FDHDIMM **500**, or stored in a memory location of a non-volatile memory subsystem, for example a portion of Flash **506** or a separate non-volatile memory subsystem. In one embodiment, the memory address space information corresponds to a portion of initialization information of the FDHDIMM **500**.

In one embodiment, MCH **510** may send a command to restore a certain amount of data information from Flash **506** to DRAM **508**. The CDC **502** provides control information to DMgr **504** to appropriately copy the necessary information from Flash **506** to the DRAM **508**. This operation can provide support for various host system booting operations and/or a special host system power up operation.

In one embodiment, MCH **510** sends a command which may include various fields comprising control information regarding data transfer size, data format options, and/or startup time. CDC **502** receives and interprets the command and provides control signals to DMgr **504** to control the data traffic between the Flash **506**, the DRAM **508**, and the MCH **510**. For example, DMgr **504** receives the data transfer size, formatting information, direction of data flow (via one or more multiplexers such as **611**, **612**, **621**, **622** as detailed below), and the starting time of the actual data transfer from CDC **502**. DMgr **504** may also receive additional control information from the CDC **502** to establish a data flow path and/or to correctly establish the data transfer fabric. In certain embodiments, DMgr **504** also functions as a bi-directional data transfer fabric. For example, DMgr **504** may have more than 2 sets of data ports facing the Flash **506** and the DRAM **508**. Multiplexers **611** and **612** provide controllable data paths from any one of the DRAMs **508**(**1**) and **508**(**2**) (DRAM-A and DRAM-B) to any one of the MCH **510** and the Flash **506**. Similarly multiplexers **621** and **622** provide controllable data paths from any one of the MCH and the Flash memory to any one of the DRAMs **508**(**1**) and **508**(**2**) (DRAM-A and DRAM-B). In one embodiment, DRAM **508**(**1**) is a segment of DRAM **508**, while in other embodiments, DRAM **508**(**1**) is a separate DRAM memory subsystem. It will be understood that each memory segment can comprise one or more memory circuits, a memory devices, and/or memory integrated circuits. Of course other configurations for DRAM **508** are possible, and other data transfer fabrics using complex data paths and suitable types of multiplexing logic are contemplated.

In accordance with one embodiment, the two sets of multiplexors **611**, **612** and **621**, **622** allow independent data transfer to Flash **506** from DRAM-A **508**(**1**) and DRAM-B **508**(**2**). For example, in response to one or more control signals or a command from CDC **502**, DMgr **504** can transfer data from DRAM-A **508**(**1**) to MCH **510**, via multiplexer **611**, at the same time as from DRAM-B **508**(**2**)

to the Flash **506**, via multiplexer **612**; or data is transferred from DRAM-B **508**(**2**) to MCH **510**, via multiplexer **611**, and simultaneously data is transferred from the Flash **506** to DRAM-A **508**(**1**), via multiplexer **621**. Further, in the same way that data can be transferred to or from the DRAM in both device-wide or segment-by-segment fashion, data can be transferred to or from the flash memory in device-wide or segment-by-segment fashion, and the flash memory can be addressed and accessed accordingly.

In accordance with one embodiment the illustrated arrangement of data transfer fabric of DMgr **504** also allows the CDC **502** to control data transfer from the Flash memory to the MCH by buffering the data from the Flash **506** using a buffer **602**, and matching the data rate and/or data format of MCH **510**. The buffer **602** is shown in FIG. **6** as a portion of a data format module **604**; however, buffer **602** may also be a distributed buffer such that one buffer is used for each one of the set of multiplexer logic elements shown as multiplexers **611**, **612**, **621**, and **622**. Various buffer arrangements may be used, such as a programmable size buffer to meet the requirement of a given system design requirement, for example the disparity between read/write access time; or overall system performance, for example latency. In certain embodiments, the buffer **604** may introduce one or more clock cycle delays into a data communication path between MCH **510**, DRAM **508**, and Flash **506**.

In certain embodiments, data format module **604** contains a data formatting subsystem (not shown) to enable DMgr **504** to format and perform data transfer in accordance with control information received from CDC **502**. Data buffer **604** of data format module **602**, discussed above, also supports a wide data bus **606** coupled to the Flash memory **506** operating at a first frequency, while receiving data from DRAM **508** using a relatively smaller width data bus **608** operating at a second frequency, the second frequency being larger than the first frequency in certain embodiments. The buffer **602** is designed to match the data flow rate between the DRAM **508** and the Flash **506**.

A register **690** provides the ability to register commands received from MCH **510** via C/A **560** (FIG. **5**A). The register **690** may communicate these commands to CDC **502** and/or to the DRAM **508** and/or Flash **506**. The register **690** communicates these registered commands to CDC **502** for processing. The register **690** may also include multiple registers (not shown), such that it can provide the ability to register multiple commands, a sequence of commands, or provide a pipeline delay stage for buffering and providing a controlled execution of certain commands received form MCH **510**.

In certain embodiments, the register **690** may register commands from MCH **510** and transmit the registered commands to DRAM **508** and/or Flash **506** memory subsystems. In certain embodiments, the CDC **502** monitors commands received from MCH **510**, via control and address bus C/A **560**, and provides appropriate control information to DMgr **504**, DRAM **508**, or Flash **506** to execute these commands and perform data transfer operations between MCH **510** and FDHDIMM **500** via MCH data bus **610**.

FIG. **7** illustrates a functional block diagram of the CDC **502**. In certain embodiments, the major functional blocks of the CDC **502** are a DRAM control block DRAMCtrl **702**, Flash control block FlashCtrl **704**, MCH command interpreter CmdInt **706**, DRAM-Flash interface scheduler Scheduler **708**, and DMgr control block (DMgrCtrl) **710**.

In accordance with one embodiment, DRAMCtrl **702** generates DRAM commands that are independent from the commands issued by the MCH **510**. In accordance with one embodiment, when the MCH **510** initiates a read/write operation from/to the same DRAM **508** that is currently executing a command from the DRAMCtrl **702**, then the CDC **502** may choose to instruct DRAMCtrl **702** to abort its operation in order to execute the operation initiated by the MCH. However, the CDC **502** may also pipeline the operation so that it causes DRAMCtrl **702** to either halt or complete its current operation prior to executing that of the MCH. The CDC **502** may also instruct DRAMCtrl **702** to resume its operation once the command from MCH **510** is completed.

In accordance with one embodiment, the FlashCtrl **704** generates appropriate Flash commands for the proper read/write operations. The CmdInt **706** intercepts commands received from MCH **510** and generates the appropriate control information and control signals and transmit them to the appropriate FDHDIMM functional block. For example, CmdInt **706** issues an interrupt signal to the DRAMCtrl **702** when the MCH issues a command that collides (conflicts) with the currently executing or pending commands that DRAMCtrl **702** has initiated independently from MCH **510**, thus subordinating these commands to those from the MCH. The Scheduler **708** schedules the Flash-DRAM interface operation such that there is no resource conflict in the DMgr **504**. In accordance with one embodiment, the Scheduler **708** assigns time slots for the DRAMCtrl **702** and FlashCtrl **704** operation based on the current status and the pending command received or to be received from the MCH. The DMgrCtrl **710** generates and sends appropriate control information and control signals for the proper operation and control of the data transfer fabric to enable or disable data paths between Flash **506**, DRAM **508**, and the MCH **510**.

FIG. **8**A is a block diagram showing a Flash-DRAM hybrid DIMM (FDHDIMM). As seen from FIG. **8**A, this Flash-DRAM hybrid DIMM requires two separate and independent address buses to separately control the address spaces: one for the Flash memory Flash **506** and the other for the DRAM memory DRAM **508**. The MCH treats the DRAM **508** and Flash **506** as separate memory subsystems, for example DRAM and SSD/HD memory subsystems. The memory in each address space is controlled directly by the MCH. However, the on-DIMM data path between Flash **506** and DRAM **508** allows for direct data transfer to occur between the Flash **506** and the DRAM **508** in response to control information from Ctrl **502**. In this embodiment, this data transfer mechanism provides direct support for executing commands from the MCH without having the MCH directly controlling the data transfer, and thus improving data transfer performance from Flash **506** to the DRAM **508**. However, the MCH needs to manage two address spaces and two different memory protocols simultaneously. Moreover, the MCH needs to map the DRAM memory space into the Flash memory space, and the data interface time suffers due to the difference in the data access time between the Flash memory and the DRAM memory.

In accordance with one embodiment, a memory space mapping of a Flash-DRAM hybrid DIMM is shown in FIG. **8**B. A memory controller of a host system (not shown) controls both of the DRAM **508** address space and the Flash **506** address space using a single unified address space. The CDC **502** receives memory access commands from the MCH and generates control information for appropriate mapping and data transfer between Flash and DRAM memory subsystem to properly carry out the memory access commands. In one embodiment, the memory controller of the host system views the large Flash memory space as a DRAM memory space, and accesses this unified memory

space with a standard DDR (double data rate) protocol used for accessing DRAM. The unified memory space in this case can exhibit overlapping memory address space between the Flash **506** and the DRAM **508**. The overlapping memory address space may be used as a temporary storage or buffer for data transfer between the Flash **506** and the DRAM **508**. For example, the DRAM memory space may hold a copy of data from the selected Flash memory space such that the MCH can access this data normally via DDR memory access commands. The CDC **502** controls the operation of the Flash **506** and DRAM **508** memory subsystems in response to commands received from a memory controller of a host system.

In one embodiment, the unified memory space corresponds to a contiguous address space comprising a first portion of the address space of the Flash **506** and a first portion of the address space of the DRAM **508**. The first portion of the address space of the Flash **506** can be determined via a first programmable register holding a first value corresponding to the desired Flash memory size to be used. Similarly, the first portion of the address space of the DRAM **508** can be determined via a second programmable register holding a second value corresponding to the desired DRAM memory size to be used. In one embodiment, any one of the first portion of the address space of the Flash **506** and the first portion of the address space of the DRAM **508** is determined via a first value corresponding to a desired performance or memory size, the first value being received by the CDC **502** via a command sent by memory controller of the host system.

In accordance with one embodiment, a flow diagram directed to the transfer of data from Flash memory to DRAM memory and vice versa in an exemplary FDHDIMM is shown in FIG. **9**. In certain embodiments, data transfer from the Flash **506** to the DRAM **508** occurs in accordance with memory access commands which the CDC **502** receives from the memory controller of the host system. In certain embodiments, the CDC **502** controls the data transfer from the DRAM **508** to the Flash **506** so as to avoid conflict with any memory operation that is currently being executed. For example, when all the pages in a particular DRAM memory block are closed. The CDC **502** partitions the DRAM memory space into a number of blocks for the purpose of optimally supporting the desired application. The controller can configure memory space in the memory module based on at least one of one or more commands received from the MCH, instructions received from the MCH, a programmable value written into a register, a value corresponding to a first portion of the volatile memory subsystem, a value corresponding to a first portion of the non-volatile memory subsystem, and a timing value. Furthermore, the block size can be configurable by the memory controller of the host system, such that the number pages in a block can be optimized to support a particular application or a task. Furthermore, the block size may be configured on-the-fly, e.g. CDC **502** can receive instruction regarding a desired block size from the memory controller via a memory command, or via a programmable value.

In certain embodiments, a memory controller can access the memory module using a standard access protocol, such as JEDEC's DDR DRAM, by sending a memory access command to the CDC **502** which in turn determines what type of a data transfer operation it is and the corresponding target address where the data information is stored, e.g. data information is stored in the DRAM **508** or Flash **506** memory subsystems. In response to a read operation, if the CDC **502** determines that data information, e.g. a page (or block), does not reside in the DRAM **508** but resides in Flash **506**, then the CDC **502** initiates and controls all necessary data transfer operations from Flash **506** to DRAM **508** and subsequently to the memory controller. In one embodiment, once the CDC **502** completes the data transfer operation of the requested data information from the Flash **506** to the DRAM **508**, the CDC **502** alerts the memory controller to retrieve the data information from the DRAM **508**. In on embodiment, the memory controller initiates the copying of data information from Flash **506** to DRAM **508** by writing, into a register in the CDC **502**, the target Flash address along with a valid block size. The CDC **502** in turn, executes appropriate operations and generates control information to copy the data information to the DRAM **508**. Consequently, the memory controller can access or retrieve the data information using standard memory access commands or protocol.

An exemplary flow chart is shown in FIG. **9**, a starting step or power up **902**, is followed by an initialization step **904**, the memory controller initiates, at step **906**, a data move from the Flash **506** to the DRAM **508** by writing target address and size, to a control register in the CDC **502**, which then copies, at **908**, data information from the Flash **506** to the DRAM **508** and erases the block in the Flash. Erasing the data information from Flash may be accomplished independently from (or concurrently with) other steps that CDC **502** performs in this flow chart, i.e. other steps can be executed concurrently with the Erase the Flash block step. Once the data information or a block of data information is thus moved to the DRAM **508**, the memory controller can operate on this data block using standard memory access protocol or commands at **910**. The CDC **502** checks, at **912**, if any of the DRAM **508** blocks, or copied blocks, are closed. If the memory controller closed any open blocks in DRAM **508**, then the CDC **502** initiate a Flash write to write the closed block from the DRAM **508** to the Flash **506**, at **914**. In addition, the memory controller, at **916**, reopens the closed block that is currently being written into the Flash **506**, then the CDC **502** stops the Flash write operation and erases the Flash block which was being written to, as shown at **918**. Otherwise, the CDC **502** continues and completes the writing operation to the Flash at **920**.

The dashed lines in FIG. **9** indicate independent or parallel activities that can be performed by the CDC **502**. At any time the CDC **502** receives a DRAM load command from a memory controller which writes a Flash target address and/or block size information into the RC register(s) at **922**, as described above, then the CDC **502** executes a load DRAM w/RC step **906** and initiates another branch (or a thread) of activities that includes steps **908-922**. In one embodiment, the CDC **502** controls the data transfer operations between DRAM **508** and Flash **506** such that the Flash **506** is completely hidden from the memory controller. The CDC **502** monitors all memory access commands sent by the memory controller using standard DRAM protocol and appropriately configures and manipulate both Flash **506** and DRAM **508** memory subsystems to perform the requested memory access operation and thus achieve the desired results. The memory controller does not interface directly with the Flash memory subsystem. Instead, the memory controller interfaces with the CDC **502** and/or DMgr **504** as shown in FIG. **5** and FIG. **6**. Moreover, the memory controller may use one or more protocol, such as DDR, DDR2, DDR3, DDR4 protocols or the like.

In accordance with one embodiment, an example of mapping a DRAM address space to Flash memory address space is shown in FIG. **10**. Two sets (**1002**, **1004**) of address

bits AD**6** to AD**17**, forming a 24 bit extended memory page address, are allocated for the block address. For example, assuming a Block size of 256K Bytes, then a 24-bit block address space (using the two sets of AD**6** to AD**17** **1002** and **1004**) would enable access to 4 TB of Flash memory storage space. If a memory module has 1 GB of DRAM storage capacity, then it can hold approximately 4K Blocks of data in the DRAM memory, each Block comprise 256 K Bytes of data. The DRAM address space, corresponding to the 4K blocks, can be assigned to different virtual ranks and banks, where the number of virtual ranks and banks is configurable and can be manipulated to meet a specific design or performance needs. For example, if a 1G Bytes memory module is configured to comprise two ranks with eight banks per rank, then each bank would hold two hundred fifty (250) blocks or the equivalent of 62 M Bytes or 62K pages, where each page correspond to a 1K Bytes. Other configurations using different page, block, banks, or ranks numbers may also be used. Furthermore, an exemplary mapping of 24-bit DDR DIMM block address to Flash memory address, using Block addressing as described above, is shown in FIG. **10**. The 24-bit can be decomposed into fields, such as a logical unit number LUN address **1061** field, a Block address **1051** field, a Plane address **1041**, a Page address **1031**, and a group of least significant address bits A_0A_1 **1021**. The Plane address **1041** is a sub address of the block address, and it may be used to support multiple page IO so as to improve Flash memory subsystem operation. In this example, it is understood that different number of bits may be allocated to each field of the 24-bit

The CDC **502** manages the block write-back operation by queuing the blocks that are ready to be written back to the Flash memory. As described above, if any page in a queued block for a write operation is reopened, then the CDC **502** will stop the queued block write operation, and remove the block from the queue. Once all the pages in a block are closed, then the CDC **502** restarts the write-back operation and queue the block for a write operation.

In accordance with one embodiment, an exemplary read operation from Flash **506** to DRAM **508** can be performed in approximately 400 μs, while a write operation from DRAM **508** to Flash **506** can be performed in approximately 22 ms resulting in a read to write ratio of 55 to 1. Therefore, if the average time a host system's memory controller spends accessing data information in a Block of DRAM is about 22 ms (that is the duration that a Block comprises one or more pages that are open), then the block write-back operation from DRAM to Flash would not impact performance and hence the disparity between read and write access may be completely hidden from the memory controller. If the block usage time is 11 ms instead of 22 ms, then the CDC **502** control the data transfer operation between DRAM **508** and Flash **506** such that there are no more than 9 closed blocks in the queue to be written-back to the Flash memory, hence approximately an average of 100 ms can be maintained for a standard DDR DRAM operation. Moreover, the number of closed Blocks in the queue to be written-back to the Flash memory subsystem varies with the average block usage time and the desired performance for a specific host system or for a specific application running using the host system resources.

Consequently, the maximum number of closed Blocks to be written-back to Flash can be approximated to be

((#of blocks per bank)/(ratio of 'Flash_block_write_time' to 'Flash_read_time'))*((Block usage time)/('Flash_block_write_time'))

In order to maintain less than 100 ms time period for queued write-back Blocks, then using a Flash memory subsystem having 22 ms write access time per Block would results in a maximum number of four Blocks to be queued for write operation to Flash **506**. Therefore, on average approximately 88 ms (=22 ms*4) for blocks means that each bank should not have more than four Blocks that need to be written back to the Flash **506**.

The above equation also indicates that bigger DRAM memory space can support shorter block usage times. For example, 2 GB of DRAM memory allows the 8 closed blocks to be written-back to Flash. The table in FIG. **11** provides an estimation of the maximum allowed closed blocks in the queue to be written back to the Flash memory for different DRAM density using various average block use time.

Certain embodiments described herein include a memory system which can communicate with a host system such as a disk controller of a computer system. The memory system can include volatile and non-volatile memory, and a controller. The controller backs up the volatile memory using the non-volatile memory in the event of a trigger condition. Trigger conditions can include, for example, a power failure, power reduction, request by the host system, etc. In order to power the system in the event of a power failure or reduction, the memory system can include a secondary power source which does not comprise a battery and may include, for example, a capacitor or capacitor array.

In certain embodiments, the memory system can be configured such that the operation of the volatile memory is not adversely affected by the non-volatile memory or by the controller when the volatile memory is interacting with the host system. For example, one or more isolation devices may isolate the non-volatile memory and the controller from the volatile memory when the volatile memory is interacting with the host system and may allow communication between the volatile memory and the non-volatile memory when the data of the volatile memory is being restored or backed-up. This configuration generally protects the operation of the volatile memory when isolated while providing backup and restore capability in the event of a trigger condition, such as a power failure.

In certain embodiments described herein, the memory system includes a power module which provides power to the various components of the memory system from different sources based on a state of the memory system in relation to a trigger condition (e.g., a power failure). The power module may switch the source of the power to the various components in order to efficiently provide power in the event of the power failure. For example, when no power failure is detected, the power module may provide power to certain components, such as the volatile memory, from system power while charging a secondary power source (e.g., a capacitor array). In the event of a power failure or other trigger condition, the power module may power the volatile memory elements using the previously charged secondary power source.

In certain embodiments, the power module. transitions relatively smoothly from powering the volatile memory with system power to powering it with the secondary power source. For example, the power system may power volatile memory with a third power source from the time the memory system detects that power failure is likely to occur until the time the memory system detects that the power failure has actually occurred.

In certain embodiments, the volatile memory system can be operated at a reduced frequency during backup and/or

restore operations which can improve the efficiency of the system and save power. In some embodiments, during backup and/or restore operations, the volatile memory communicates with the non-volatile memory by writing and/or. reading data words in bit-wise slices instead of by writing entire words at once. In certain embodiments, when each slice is being written to or read from the volatile memory the unused slice(s) of volatile memory is not active, which can reduce the power consumption of the system.

In yet other embodiments, the non-volatile memory can include at least 100 percent more storage capacity than the volatile memory. This configuration can allow the memory system to efficiently handle subsequent trigger conditions.

FIG. **12** is a block diagram of an example memory system **1010** compatible with certain embodiments described herein. The memory system **1010** can be coupled to a host computer system and can include a volatile memory subsystem **1030**, a non-volatile memory subsystem **1040**, and a controller **1062** operatively coupled to the non-volatile memory subsystem **1040**. In certain embodiments, the memory system **1010** includes at least one circuit **1052** configured to selectively operatively decouple the controller **1062** from the volatile memory subsystem **1030**.

In certain embodiments, the memory system **1010** comprises a memory module. The memory system **1010** may comprise a printed-circuit board (PCB) **1020**. In certain embodiments, the memory system **1010** has a memory capacity of 512-MB, 1-GB, 2-GB, 4-GB, or 8-GB. Other volatile memory capacities are also compatible with certain embodiments described herein. In certain embodiments, the memory system **10** has a non-volatile memory capacity of 512-MB, 1-GB, 2-GB, 4-GB, 8-GB, 16-GB, or 32-GB. Other non-volatile memory capacities are also compatible with certain embodiments described herein. In addition, memory systems **1010** having widths of 4 bytes, 8 bytes, 16 bytes, 32 bytes, or 32 bits, 64 bits, 128 bits, 256 bits, as well as other widths (in bytes or in bits), are compatible with embodiments described herein. In certain embodiments, the PCB **1020** has an industry-standard form factor. For example, the PCB **1020** can have a low profile (LP) form factor with a height of 30 millimeters and a width of 133.35 millimeters. In certain other embodiments, the PCB **1020** has a very high profile (VHP) form factor with a height of 50 millimeters or more. In certain other embodiments, the PCB **1020** has a very low profile (VLP) form factor with a height of 18.3 millimeters. Other form factors including, but not limited to, small-outline (SO-DIMM), unbuffered (UDIMM), registered (RDIMM), fully-buffered (FB-DIMM), miniDIMM, mini-RDIMM, VLP mini-DIMM, micro-DIMM, and SRAM DIMM are also compatible with certain embodiments described herein. For example, in other embodiments, certain non-DIMM form factors are possible such as, for example, single in-line memory module (SIMM), multi-media card (MMC), and small computer system interface (SCSI).

In certain preferred embodiments, the memory system **1010** is in electrical communication with the host system. In other embodiments, the memory system **1010** may communicate with a host system using some other type of communication, such as, for example, optical communication. Examples of host systems include, but are not limited to, blade servers, 1 U servers, personal computers (PCs), and other applications in which space is constrained or limited. The memory system **1010** can be in communication with a disk controller of a computer system, for example. The PCB **1020** can comprise an interface **1022** that is configured to be in electrical communication with the host system (not shown). For example, the interface **1022** can comprise a plurality of edge connections which fit into a corresponding slot connector of the host system. The interface **1022** of certain embodiments provides a conduit for power voltage as well as data, address, and control signals between the memory system **1010** and the host system. For example, the interface **1022** can comprise a standard 240-pin DDR2 edge connector.

The volatile memory subsystem **1030** comprises a plurality of volatile memory elements **1032** and the non-volatile memory subsystem **1040** comprises a plurality of non-volatile memory elements **1042**. Certain embodiments described herein advantageously provide nonvolatile storage via the non-volatile memory subsystem **1040** in addition to high-performance (e.g., high speed) storage via the volatile memory subsystem **1030**. In certain embodiments, the first plurality of volatile memory elements **1032** comprises two or more dynamic random-access memory (DRAM) elements. Types of DRAM elements **1032** compatible with certain embodiments described herein include, but are not limited to, DDR, DDR2, DDR3, and synchronous DRAM (SDRAM). For example, in the block diagram of FIG. **12**, the first memory bank **1030** comprises eight 64M×8 DDR2 SDRAM elements **1032**. The volatile memory elements **1032** may comprise other types of memory elements such as static random-access memory (SRAM). In addition, volatile memory elements **1032** having bit widths of 4, 8, 16, 32, as well as other bit widths, are compatible with certain embodiments described herein. Volatile memory elements **1032** compatible with certain embodiments described herein have packaging which include, but are not limited to, thin small-outline package (TSOP), ball-grid-array (BGA), fine-pitch BGA (FBOA), micro-BOA (1.1, BGA), mini-BGA (mBGA), and chip-scale packaging (CSP).

In certain embodiments, the second plurality of non-volatile memory elements **1042** comprises one or more flash memory elements. Types of flash memory elements **1042** compatible with certain embodiments described herein include, but are not limited to, NOR flash, NAND flash, ONE-NAND flash, and multi-level cell (MLC). For example, in the block diagram of FIG. **12**, the second memory bank **1040** comprises 512 MB of flash memory organized as four 128 Mb×8 NAND flash memory elements **1042**. In addition, nonvolatile memory elements **1042** having bit widths of 4, 8, 16, 32, as well as other bit widths, are compatible with certain embodiments described herein. Non-volatile memory elements **1042** compatible with certain embodiments described herein have packaging which include, but are not limited to, thin small-outline package (TSOP), ball-grid-array (BOA), fine-pitch BOA (FBGA), micro-BOA (POA), mini-BGA (mBGA), and chip-scale packaging (CSP).

FIG. **13** is a block diagram of an example memory module **10** with ECC (error-correcting code) having a volatile memory subsystem **1030** with nine volatile memory elements **1032** and a non-volatile memory subsystem **1040** with five non-volatile memory elements **1042** in accordance with certain embodiments described herein. The additional memory element **1032** of the first memory bank **1030** and the additional memory element **1042** of the second memory bank **1040** provide the ECC capability. In certain other embodiments, the volatile memory subsystem **1030** comprises other numbers of volatile memory elements **1032** (e.g., 2, 3, 4, 5, 6, 7, more than 9). In certain embodiments, the non-volatile memory subsystem **1040** comprises other numbers of nonvolatile memory elements **1042** (e.g., 2, 3, more than 5).

Referring to FIG. **12**, in certain embodiments, the logic element **1070** comprises a field-programmable gate array (FPGA). In certain embodiments, the logic element **1070** comprises an FPGA available from Lattice Semiconductor Corporation which includes an internal flash. In certain other embodiments, the logic element **1070** comprises an FPOA available from another vendor. The internal flash can improve the speed of the memory system **1010** and save physical space. Other types of logic elements **1070** compatible with certain embodiments described herein include, but are not limited to, a programmable-logic device (PLD), an application-specific integrated circuit (ASIC), a custom-designed semiconductor device, a complex programmable logic device (CPLD). In certain embodiments, the logic element **1070** is a custom device. In certain embodiments, the logic element **1070** comprises various discrete electrical elements, while in certain other embodiments, the logic element **1070** comprises one or more integrated circuits. FIG. **14** is a block diagram of an example memory module **1010** having a microcontroller unit **1060** and logic element **1070** integrated into a single controller **1062** in accordance with certain embodiments described herein. In certain embodiments, the controller **1062** includes one or more other components. For example, in one embodiment, an FPGA without an internal flash is used and the controller **1062** includes a separate flash memory component which stores configuration information to program the FPGA.

In certain embodiments, the at least one circuit **1052** comprises one or more switches coupled to the volatile memory subsystem **1030**, to the controller **1062**, and to the host computer (e.g., via the interface **1022**, as schematically illustrated by FIGS. **12-14**). The one or more switches are responsive to signals (e.g., from the controller **1062**) to selectively operatively decouple the controller **1062** from the volatile memory subsystem **1030** and to selectively operatively couple the controller **1062** to the volatile memory subsystem **1030**. In addition, in certain embodiments, the at least one circuit **1052** selectively operatively couples and decouples the volatile memory subsystem **1030** and the host system.

In certain embodiments, the volatile memory subsystem **1030** can comprise a registered DIMM subsystem comprising one or more registers **1160** and a plurality of DRAM elements **1180**, as schematically illustrated by FIG. **15**A. In certain such embodiments, the at least one circuit **1052** can comprise one or more switches **1172** coupled to the controller **1062** (e.g., logic element **1070**) and to the volatile memory subsystem **1030** which can be actuated to couple and decouple the controller **1062** to and from the volatile memory subsystem **1030**, respectively. The memory system **1010** further comprises one or more switches **1170** coupled to the one or more registers **1160** and to the plurality of DRAM elements **1180** as schematically illustrated by FIG. **15**A. The one or more switches **1170** can be selectively switched, thereby selectively operatively coupling the volatile memory subsystem **1030** to the host system **1150**. In certain other embodiments, as schematically illustrated by FIG. **15**B, the one or more switches **1174** are also coupled to the one or more registers **1160** and to a power source **1162** for the one or more registers **1160**. The one or more switches **1174** can be selectively switched to turn power on or off to the one or more registers **1160**, thereby selectively operatively coupling the volatile memory subsystem **1030** to the host system **1150**. As schematically illustrated by FIG. **15**C, in certain embodiments the at least one circuit **1052** comprises a dynamic on-die termination (ODT) **1176** circuit of the logic element **1070**. For example, the logic element **1070** can comprise a dynamic ODT circuit **1176** which selectively operatively couples and decouples the logic element **1070** to and from the volatile memory subsystem **1030**, respectively. In addition, and similar to the example embodiment of FIG. **15**A described above, the one or more switches **1170** can be selectively switched, thereby selectively operatively coupling the volatile memory subsystem **1030** to the host system **1150**.

Certain embodiments described herein utilize the non-volatile memory subsystem **1040** as a flash "mirror" to provide backup of the volatile memory subsystem **1030** in the event of certain system conditions. For example, the non-volatile memory subsystem **1040** may backup the volatile memory subsystem **1030** in the event of a trigger condition, such as, for example, a power failure or power reduction or a request from the host system. In one embodiment, the nonvolatile memory subsystem **1040** holds intermediate data results in a noisy system environment when the host computer system is engaged in a long computation. In certain embodiments, a backup may be performed on a regular basis. For example, in one embodiment, the backup may occur every millisecond in response to a trigger condition. In certain embodiments, the trigger condition occurs when the memory system **1010** detects that the system voltage is below a certain threshold voltage. For example, in one embodiment, the threshold voltage is 10 percent below a specified operating voltage. In certain embodiments, a trigger condition occurs when the voltage goes above a certain threshold value, such as, for example, 10 percent above a specified operating voltage. In some embodiments, a trigger condition occurs when the voltage goes below a threshold or above another threshold. In various embodiments, a backup and/or restore operation may occur in reboot and/or non-reboot trigger conditions.

As schematically illustrated by FIGS. **12** and **13**, in certain embodiments, the controller **1062** may comprise a microcontroller unit (MCU) **1060** and a logic element **1070**. In certain embodiments, the MCU **1060** provides memory management for the non-volatile memory subsystem **1040** and controls data transfer between the volatile memory subsystem **30** and the nonvolatile memory subsystem **1040**. The MCU **1060** of certain embodiments comprises a 16-bit microcontroller, although other types of microcontrollers are also compatible with certain embodiments described herein. As schematically illustrated by FIGS. **12** and **13**, the logic element **1070** of certain embodiments is in electrical communication with the non-volatile memory subsystem **1040** and the MCU **1060**. The logic element **1070** can provide signal level translation between the volatile memory elements **1032** (e.g., 1.8V SSTL-2 for DDR2 SDRAM elements) and the non-volatile memory elements **1042** (e.g., 3V TTL for NAND flash memory elements). In certain embodiments, the logic element **1070** is also programmed to perform address/address translation between the volatile memory subsystem **1030** and the non-volatile memory subsystem **1040**. In certain preferred embodiments, 1-NAND type flash are used for the non-volatile memory elements **1042** because of their superior read speed and compact structure.

The memory system **1010** of certain embodiments is configured to be operated in at least two states. The at least two states can comprise a first state in which the controller **1062** and the non-volatile memory subsystem **1040** are operatively decoupled (e.g., isolated) from the volatile memory subsystem **1030** by the at least one circuit **1052** and a second state in which the volatile memory subsystem **1030** is operatively coupled to the controller **1062** to allow data to

be communicated between the volatile memory subsystem **1030** and the nonvolatile memory subsystem **1040** via the controller **1062**. The memory system **1010** may transition from the first state to the second state in response to a trigger condition, such as when the memory system **1010** detects that there is a power interruption (e.g., power failure or reduction) or a system hang-up.

The memory system **1010** may further comprise a voltage monitor **1050**. The voltage monitor circuit **1050** monitors the voltage supplied by the host system via the interface **1022**. Upon detecting a low voltage condition (e.g., due to a power interruption to the host system), the voltage monitor circuit **1050** may transmit a signal to the controller **1062** indicative of the detected condition. The controller **1062** of certain embodiments responds to the signal from the voltage monitor circuit **1050** by transmitting a signal to the at least one circuit **1052** to operatively couple the controller to the volatile memory system **1030**, such that the memory system **1010** enters the second state. For example, the voltage monitor **1050** may send a signal to the MCU **1060** which responds by accessing the data on the volatile memory system **1030** and by executing a write cycle on the nonvolatile memory subsystem **1040**. During this write cycle, data is read from the volatile memory subsystem **1030** and is transferred to the non-volatile memory subsystem **1040** via the MCU **1060**. In certain embodiments, the voltage monitor circuit **1050** is part of the controller **1062** (e.g., part of the MCU **1060**) and the voltage monitor circuit **1050** transmits a signal to the other portions of the controller **1062** upon detecting a power threshold condition.

The isolation or operational decoupling of the volatile memory subsystem **1030** from the non-volatile memory subsystem in the first state can preserve the integrity of the operation of the memory system **1010** during periods of operation in which signals (e.g., data) are transmitted between the host system and the volatile memory subsystem **1030**. For example, in one embodiment during such periods of operation, the controller **1062** and the nonvolatile memory subsystem **1040** do not add a significant capacitive load to the volatile memory system **1030** when the memory system **1010** is in the first state. In certain such embodiments, the capacitive load of the controller **1062** and the non-volatile memory subsystem **1040** do not significantly affect the signals propagating between the volatile memory subsystem **1030** and the host system. This can be particularly advantageous in relatively high-speed memory systems where loading effects can be significant. In one preferred embodiment, the at least one circuit **1052** comprises an FSA1208 Low-Power, Eight-Port, Hi-Speed Isolation Switch from Fairchild Semiconductor. In other embodiments, the at least one circuit **1052** comprises other types of isolation devices.

Power may be supplied to the volatile memory subsystem **1030** from a first power supply (e.g., a system power supply) when the memory system **1010** is in the first state and from a second power supply **1080** when the memory system **1010** is in the second state. In certain embodiments, the memory system **1010** is in the first state when no trigger condition (e.g., a power failure) is present and the memory system **1010** enters the second state in response to a trigger condition. In certain embodiments, the memory system **1010** has a third state in which the controller **1062** is operatively decoupled from the volatile memory subsystem **1030** and power is supplied to the volatile memory subsystem **1030** from a third power supply (not shown). For example, in one embodiment the third power supply may provide power to the volatile memory subsystem **1030** when the memory system **1010** detects that a trigger condition is likely to occur but has not yet occurred.

In certain embodiments, the second power supply **1080** does not comprise a battery. Because a battery is not used, the second power supply **1080** of certain embodiments may be relatively easy to maintain, does not generally need to be replaced, and is relatively environmentally friendly. In certain embodiments, as schematically illustrated by FIGS. **12-14**, the second power supply **1080** comprises a step-up transformer **1082**, a step-down transformer **1084**, and a capacitor bank **1086** comprising one or more capacitors (e.g., double-layer capacitors). In one example embodiment, capacitors may take about three to four minutes to charge and about two minutes to discharge. In other embodiments, the one or more capacitors may take a longer time or a shorter time to charge and/or discharge. For example, in certain embodiments, the second power supply **1080** is configured to power the volatile memory subsystem **1030** for less than thirty minutes. In certain embodiments, the second power supply **1080** may comprise a battery. For example, in certain embodiments, the second power supply **1080** comprises a battery and one or more capacitors and is configured to power the volatile memory subsystem **1030** for no more than thirty minutes.

In certain embodiments, the capacitor bank **1086** of the second power supply **1080** is charged by the first power supply while the memory system **1010** is in the first state. As a result, the second power supply **1080** is fully charged when the memory system **1010** enters the second state. The memory system **1010** and the second power supply **1080** may be located on the same printed circuit board **1020**. In other embodiments, the second power supply **1080** may not be on the same printed circuit board **1020** and may be tethered to the printed circuit board **1020**, for example.

When operating in the first state, in certain embodiments, the step-up transformer **1082** keeps the capacitor bank **1086** charged at a peak value. In certain embodiments, the step-down transformer **1084** acts as a voltage regulator to ensure that regulated voltages are supplied to the memory elements (e.g., 1.8V to the volatile DRAM elements **1032** and 3.0V to the non-volatile flash memory elements **1042**) when operating in the second state (e.g., during power down). In certain embodiments, as schematically illustrated by FIGS. **12-14**, the memory module **1010** further comprises a switch **1090** (e.g., FET switch) that switches power provided to the controller **1062**, the volatile memory subsystem **1030**, and the non-volatile memory subsystem **1040**, between the power from the second power supply **1080** and the power from the first power supply (e.g., system power) received via the interface **1022**. For example, the switch **1090** may switch from the first power supply to the second power supply **1080** when the voltage monitor **1050** detects a low voltage condition. The switch **1090** of certain embodiments advantageously ensures that the volatile memory elements **1032** and non-volatile memory elements **1042** are powered long enough for the data to be transferred from the volatile memory elements **1032** and stored in the non-volatile memory elements **1042**. In certain embodiments, after the data transfer is complete, the switch **1090** then switches back to the first power supply and the controller **1062** transmits a signal to the at least one circuit **1052** to operatively decouple the controller **1062** from the volatile memory subsystem **1030**, such that the memory system **1010** reenters the first state.

When the memory system **1010** re-enters the first state, data may be transferred back from the non-volatile memory

subsystem **1040** to the volatile memory subsystem **1030** via the controller **1062**. The host system can then resume accessing the volatile memory subsystem **1030** of the memory module **1010**. In certain embodiments, after the memory system **1010** enters or re-enters the first state (e.g., after power is restored), the host system accesses the volatile memory subsystem **1030** rather than the non-volatile memory subsystem **1040** because the volatile memory elements **1032** have superior read/write characteristics. In certain embodiments, the transfer of data from the volatile memory bank **1030** to the nonvolatile memory bank **1040**, or from the non-volatile memory bank **1040** to the volatile. memory bank **1030**, takes less than one minute per GB.

In certain embodiments, the memory system **1010** protects the operation of the volatile memory when communicating with the host-system and provides backup and restore capability in the event of a trigger condition such as a power failure. In certain embodiments, the memory system **1010** copies the entire contents of the volatile memory subsystem **1030** into the nonvolatile memory subsystem **1040** on each backup operation. Moreover, in certain embodiments, the entire contents of the non-volatile memory subsystem **1040** are copied back into the volatile memory subsystem **1030** on each restore operation. In certain embodiments, the entire contents of the non-volatile memory subsystem **1040** are accessed for each backup and/or restore operation, such that the non-volatile memory subsystem **1040** (e.g., flash memory subsystem) is used generally uniformly across its memory space and wear-leveling is not performed by the memory system **1010**. In certain embodiments, avoiding wear-leveling can decrease cost and complexity of the memory system **1010** and can improve the performance of the memory system **1010**. In certain other embodiments, the entire contents of the volatile memory subsystem **1030** are not copied into the non-volatile memory subsystem **1040** on each backup operation, but only a partial copy is performed. In certain embodiments, other management capabilities such as bad-block management and error management for the flash memory elements of the non-volatile memory subsystem **1040** are performed in the controller **1062**.

The memory system **1010** generally operates as a write-back cache in certain embodiments. For example, in one embodiment, the host system (e.g., a disk controller) writes data to the volatile memory subsystem **1030** which then writes the data to non-volatile storage which is not part of the memory system **1010**, such as, for example, a hard disk. The disk controller may wait for an acknowledgment signal from the memory system **1010** indicating that the data has been written to the hard disk or is otherwise secure. The memory system **1010** of certain embodiments can decrease delays in the system operation by indicating that the data has been written to the hard disk before it has actually done so. In certain embodiments, the memory system **1010** will still be able to recover the data efficiently in the event of a power outage because of the backup and restore capabilities described herein. In certain other embodiments, the memory system **1010** may be operated as a write-through cache or as some other type of cache.

FIG. **16** schematically illustrates an example power module **1100** of the memory system **1010** in accordance with certain embodiments described herein. The power module **1100** provides power to the various components of the memory system **1010** using different elements based on a state of the memory system **1010** in relation to a trigger condition. In certain embodiments, the power module **1100** comprises one or more of the components described above with respect to FIG. **12**. For example, in certain embodiments, the power module **1100** includes the second power supply **1080** and the switch **1090**.

The power module **1100** provides a plurality of voltages to the memory system **1010** comprising non-volatile and volatile memory subsystems **1030**, **1040**. The plurality of voltages comprises at least a first voltage **1102** and a second voltage **1104**. The power module **1100** comprises an input **1106** providing a third voltage **1108** to the power module **1100** and a voltage conversion element **1120** configured to provide the second voltage **1104** to the memory system **1010**. The power module **1100** further comprises a first power element **1130** configured to selectively provide a fourth voltage **1110** to the conversion element **1120**. In certain embodiments, the first power element **1130** comprises a pulse-width modulation power controller. For example, in one example embodiment, the first power element **1130** is configured to receive a 1.8V input system voltage as the third voltage **1108** and to output a modulated 5V output as the fourth voltage **1110**.

The power module **1100** further comprises a second power element **1140** can be configured to selectively provide a fifth voltage **1112** to the conversion element **1120**. The power module **1100** can be configured to selectively provide the first voltage **1102** to the memory system **1010** either from the conversion element **1120** or from the input **1106**.

The power module **1100** can be configured to be operated in at least three states in certain embodiments. In a first state, the first voltage **1102** is provided to the memory system **1010** from the input **1106** and the fourth voltage **1110** is provided to the conversion element **1120** from the first power element **1130**. In a second state, the fourth voltage **1110** is provided to the conversion element **1120** from the first power element **1130** and the first voltage **1102** is provided to the memory system **1010** from the conversion element **1120**. In the third state, the fifth voltage **1112** is provided to the conversion element **1120** from the second power element **1140** and the first voltage **1104** is provided to the memory system **1010** from the conversion element **1120**.

In certain embodiments, the power module **1100** transitions from the first state to the second state upon detecting that a trigger condition is likely to occur and transitions from the second state to the third state upon detecting that the trigger condition has occurred. For example, the power module **1100** may transition to the second state when it detects that a power failure is about to occur and transitions to the third state when it detects that the power failure has occurred. In certain embodiments, providing the first voltage **1102** in the second state from the first power element **1130** rather than from the input **1106** allows a smoother transition from the first state to the third state. For example, in certain embodiments, providing the first voltage **1102** from the first power element **1130** has capacitive and other smoothing effects. In addition, switching the point of power transition to be between the conversion element **1120** and the first and second power elements **1130**, **1140** (e.g., the sources of the pre-regulated fourth voltage **1110** in the second state and the pre-regulated fifth voltage **1112** in the third state) can smooth out potential voltage spikes.

In certain embodiments, the second power element **1140** does not comprise a battery and may comprise one or more capacitors. For example, as schematically illustrated in FIG. **16**, the second power element **1140** comprises a capacitor array **1142**, a buck-boost converter **1144** which adjusts the voltage for charging the capacitor array and a voltage/current limiter **1146** which limits the charge current to the capacitor array **1142** and stops charging the capacitor array **1142** when it has reached a certain charge voltage. In one

example embodiment, the capacitor array **1142** comprises two 50 farad capacitors capable of holding a total charge of 4.6V. For example, in one example embodiment, the buck-boost converter **1144** receives a 1.8V system voltage (first voltage **1108**) and boosts the voltage to 4.3V which is outputted to the voltage current limiter **1146**. The voltage/current limiter **1146** limits the current going to the capacitor array **1142** to 1 A and stops charging the array **1142** when it is charged to 4.3V. Although described with respect to certain example embodiments, one of ordinary skill will recognize from the disclosure herein that the second power element **1140** may include alternative embodiments. For example, different components and/or different value components may be used. For example, in other embodiments, a pure boost converter may be used instead of a buck-boost converter. In another embodiment, only one capacitor may be used instead of a capacitor array **1142**.

The conversion element **1120** can comprise one or more buck converters and/or one or more buck-boost converters. The conversion element **1120** may comprise a plurality of sub-blocks **1122**, **1124**, **1126** as schematically illustrated by FIG. **16**, which can provide more voltages in addition to the second voltage **1104** to the memory system **1010**. The sub-blocks may comprise various converter circuits such as buck-converters, boost converters, and buck-boost converter circuits for providing various voltage values to the memory system **1010**. For example, in one embodiment, sub-block **1122** comprises a buck converter, sub-block **1124** comprises a dual buck converter, and sub-block **1126** comprises a buck-boost converter as schematically illustrated by FIG. **16**. Various other components for the sub-blocks **1122**, **1124**, **1126** of the conversion element **1120** are also compatible with certain embodiments described herein. In certain embodiments, the conversion element **1120** receives as input either the fourth voltage **1110** from the first power element **1130** or the fifth voltage **1112** from the second power element **1140**, depending on the state of the power module **1100**, and reduces the input to an appropriate amount for powering various components of the memory system. For example, the buck-converter of sub-block **1122** can provide 1.8V at 2 A for about 60 seconds to the volatile memory elements **1032** (e.g., DRAM), the non-volatile memory elements **1042** (e.g., flash), and the controller **1062** (e.g., an FPGA) in one embodiment. The sub-block **1124** can provide the second voltage **1104** as well as another reduced voltage **1105** to the memory system **1010**. In one example embodiment, the second voltage **1104** is 2.5V and is used to power the at least one circuit **1052** (e.g., isolation device) and the other reduced voltage **1105** is 1.2V and is used to power the controller **1062** (e.g., FPGA). The subblock **1126** can provide yet another voltage **1107** to the memory system **1010**. For example, the voltage **1107** may be 3.3V and may be used to power both the controller **1062** and the at least one circuit **1052**.

Although described with respect to certain example embodiments, one of ordinary skill will recognize from the disclosure herein that the conversion element **1120** may include alternative embodiments. For example, there may be more or less sub-blocks which may comprise other types of converters (e.g., pure boost converters) or which may produce different voltage values. In one embodiment, the volatile memory elements **1032** and nonvolatile memory elements **1042** are powered using independent voltages and are not both powered using the first voltage **1102**.

FIG. **17** is a flowchart of an example method **1200** of providing a first voltage **1102** and a second voltage **1104** to a memory system **1010** including volatile and nonvolatile memory subsystems **1030**, **1040**. While the method **1200** is described herein by reference to the memory system **1010** schematically illustrated by FIGS. **12-15**, other memory systems are also compatible with embodiments of the method **1200**. During a first condition, the method **1200** comprises providing the first voltage **1102** to the memory system **1010** from an input power supply **1106** and providing the second voltage **1104** to the memory system **1010** from a first power subsystem in operational block **1210**. For example, in one embodiment, the first power subsystem comprises the first power element **1130** and the voltage conversion element **1120** described above with respect to FIG. **16**. In other embodiments, other first power subsystems are used.

The method **1200** further comprises detecting a second condition in operational block **1220**. In certain embodiments, detecting the second condition comprises detecting that a trigger condition is likely to occur. During the second condition, the method **1200** comprises providing the first voltage **1102** and the second voltage **1104** to the memory system **1010** from the first power subsystem in an operational block **1230**. For example, referring to FIG. **16**, a switch **1148** can be toggled to provide the first voltage **1102** from the conversion element **1120** rather than from the input power supply.

The method **1200** further comprises charging a second power subsystem in operational block **1240**. In certain embodiments, the second power subsystem comprises the second power element **1140** or another power supply that does not comprise a battery. For example, in one embodiment, the second power subsystem comprises the second power element **1140** and the voltage conversion element **1120** described above with respect to FIG. **16**. In other embodiments, some other second power subsystem is used.

The method **1200** further comprises detecting a third condition in an operational block **1250** and during the third condition, providing the first voltage **1102** and the second voltage **1104** to the memory system **1010** from the second power subsystem **1140** in an operational block **1260**. In certain embodiments, detecting the third condition comprises detecting that the trigger condition has occurred. The trigger condition may comprise various conditions described herein. In various embodiments, for example, the trigger condition comprises a power reduction, power failure, or system hang-up. The operational blocks of the method **1200** may be performed in different orders in various embodiments. For example, in certain embodiments, the second power subsystem **1140** is charged before detecting the second condition.

In certain embodiments, the memory system **1010** comprises a volatile memory subsystem **1030** and a non-volatile memory subsystem **1040** comprising at least 100 percent more storage capacity than does the volatile memory subsystem. The memory system **1010** also comprises a controller **1062** operatively coupled to the volatile memory subsystem **1030** and operatively coupled to the non-volatile memory subsystem **1040**. The controller **1062** can be configured to allow data to be communicated between the volatile memory subsystem **1030** and the host system when the memory system **1010** is operating in a first state and to allow data to be communicated between the volatile memory subsystem **1030** and the non-volatile memory subsystem **1040** when the memory system **1010** is operating in a second state.

Although the memory system **1010** having extra storage capacity of the non-volatile memory subsystem **1040** has been described with respect to certain embodiments, alter-

native configurations exist. For example, in certain embodiments, there may be more than 100 percent more storage capacity in the non-volatile memory subsystem **1040** than in the volatile memory subsystem **1030**. In various embodiments, there may be at least 200, 300, or 400 percent more storage capacity in the non-volatile memory subsystem **1040** than in the volatile memory subsystem **1030**. In other embodiments, the non-volatile memory subsystem **1040** includes at least some other integer multiples of the storage capacity of the volatile memory subsystem **1030**. In some embodiments, the non-volatile memory subsystem **1040** includes a non-integer multiple of the storage capacity of the volatile memory subsystem **1030**. In one embodiment, the non-volatile memory subsystem **1040** includes less than 100 percent more storage capacity than does the volatile memory subsystem **1030**.

The extra storage capacity of the non-volatile memory subsystem **1040** can be used to improve the backup capability of the memory system **1010**. In certain embodiments in which data can only be written to portions of the non-volatile memory subsystem **1040** which do not contain data (e.g., portions which have been erased), the extra storage capacity of the nonvolatile memory subsystem **1040** allows the volatile memory subsystem **1030** to be backed up in the event of a subsequent power failure or other trigger event. For example, the extra storage capacity of the non-volatile memory subsystem **1040** may allow the memory system **1010** to backup the volatile memory subsystem **1030** efficiently in the event of multiple trigger conditions (e.g., power failures). In the event of a first power failure, for example, the data in the volatile memory system **1030** is copied to a first, previously erased portion of the nonvolatile memory subsystem **1040** via the controller **1062**. Since the non-volatile memory subsystem **1040** has more storage capacity than does the volatile memory subsystem **1030**, there is a second portion of the non-volatile memory subsystem **1040** which does not have data from the volatile memory subsystem **1030** copied to it and which remains free of data (e.g., erased). Once system power is restored, the controller **1062** of the memory system **1010** restores the data to the volatile memory subsystem **1030** by copying the backed-up data from the non-volatile memory subsystem **40** back to the volatile memory subsystem **1030**. After the data is restored, the memory system **1010** erases the non-volatile memory subsystem **1040**. While the first portion of the non-volatile memory subsystem **1040** is being erased, it may be temporarily unaccessible.

If a subsequent power failure occurs before the first portion of the non-volatile memory subsystem **1040** is completely erased, the volatile memory subsystem **1030** can be backed-up or stored again in the second portion of the non-volatile memory subsystem **1040** as described herein. In certain embodiments, the extra storage capacity of the non-volatile memory subsystem **1040** may allow the memory system **1010** to operate more efficiently. For example, because of the extra storage capacity of the non-volatile memory subsystem **1040**, the memory system **1010** can handle a higher frequency of trigger events that is not limited by the erase time of the non-volatile memory subsystem **1040**.

FIG. **18** is a flowchart of an example method **1300** of controlling a memory system **1010** operatively coupled to a host system and which includes a volatile memory subsystem **1030** and a non-volatile memory subsystem **1040**. In certain embodiments, the non-volatile memory subsystem **1040** comprises at least 100 percent more storage capacity than does the volatile memory subsystem **30** as described herein. While the method **1300** is described herein by reference to the memory system **1010** schematically illustrated by FIGS. **12-14**, the method **1300** can be practiced using other memory systems in accordance with certain embodiments described herein. In an operational block **1310**, the method **1300** comprises communicating data between the volatile memory subsystem **1030** and the host system when the memory system **1010** is in a first mode of operation. The method **1300** further comprises storing a first copy of data from the volatile memory subsystem **1030** to the non-volatile memory subsystem **1040** at a first time when the memory system **1010** is in a second mode of operation in an operational block **1320**.

In an operational block **1330**, the method **1300** comprises restoring the first copy of data from the non-volatile memory subsystem **1040** to the volatile memory subsystem **1030**. The method **1300** further comprises erasing the first copy of data from the non-volatile memory subsystem **1040** in an operational block **1340**. The method further comprises storing a second copy of data from the volatile memory subsystem **1030** to the non-volatile memory subsystem **1040** at a second time when the memory system **1010** is in the second mode of operation in an operational block **1350**. Storing the second copy begins before the first copy is completely erased from the non-volatile memory subsystem **1040**.

In some embodiments, the memory system **1010** enters the second mode of operation in response to a trigger condition, such as a power failure. In certain embodiments, the first copy of data and the second copy of data are stored in separate portions of the nonvolatile memory subsystem **1040**. The method **1300** can also include restoring the second copy of data from the non-volatile memory subsystem **1040** to the volatile memory subsystem **1030** in an operational block **1360**. The operational blocks of method **1300** referred to herein may be performed in different orders in various embodiments. For example, in some embodiments, the second copy of data is restored to the volatile memory subsystem **1030** at operational block **1360** before the first copy of data is completely erased in the operational block **1340**.

FIG. **19** schematically illustrates an example clock distribution topology **1400** of a memory system **1010** in accordance with certain embodiments described herein. The clock distribution topology **1400** generally illustrates the creation and routing of the clock signals provided to the various components of the memory system **1010**. A clock source **1402** such as, for example, a 25 MHz oscillator, generates a clock signal. The clock source **1402** may feed a clock generator **1404** which provides a clock signal **1406** to the controller **1062**, which may be an FPGA. In one embodiment, the clock generator **1404** generates a 125 MHz clock signal **1406**. The controller **1062** receives the clock signal **1406** and uses it to clock the controller **1062** master state control logic. For example, the master state control logic may control the general operation of an FPGA controller **1062**.

The clock signal **1406** can also be input into a clock divider **1410** which produces a frequency-divided version of the clock signal **1406**. In an example embodiment, the clock divider **1410** is a divide by two clock divider and produces a 62.5 MHz clock signal in response to the 125 MHz clock signal **1406**. A non-volatile memory phase-locked loop (PLL) block **1412** can be included (e.g., in the controller **1062**) which distributes a series of clock signals to the non-volatile memory subsystem **1040** and to associated control logic. For example, a series of clock signals **1414** can

be sent from the controller **1062** to the non-volatile memory subsystem **1040**. Another clock signal **1416** can be used by the controller logic which is dedicated to controlling the non-volatile memory subsystem **1040**. For example, the clock signal **1416** may clock the portion of the controller **1062** which is dedicated to generating address and/or control lines for the non-volatile memory subsystem **1040**. A feedback clock signal **1418** is fed back into the non-volatile memory PLL block **1412**. In one embodiment, the PLL block **1412** compares the feedback clock **1418** to the reference clock **1411** and varies the phase and frequency of its output until the reference **1411** and feedback **1418** clocks are phase and frequency matched.

A version of the clock signal **1406** such as the backup clock signal **1408** may be sent from the controller to the volatile memory subsystem **1030**. The clock signal **1408** may be, for example, a differential version of the clock signal **1406**. As described herein, the backup clock signal **1408** may be used to clock the volatile memory subsystem **1030** when the memory system **1010** is backing up the data from the volatile memory subsystem **1030** into the non-volatile memory subsystem **1040**. In certain embodiments, the backup clock signal **1408** may also be used to clock the volatile memory subsystem **1030** when the memory system **1010** is copying the backed-up data back into the volatile memory subsystem **1030** from the nonvolatile memory subsystem **1040** (also referred to as restoring the volatile memory subsystem **1030**). The volatile memory subsystem **1030** may normally be run at a higher frequency (e.g., DRAM running at 400 MHz) than the nonvolatile memory subsystem **1040** (e.g., flash memory running at 62.5 MHz) when communicating with the host system (e.g., when no trigger condition is present). However, in certain embodiments the volatile memory subsystem **1030** may be operated at a reduced frequency (e.g., at twice the frequency of the non-volatile memory subsystem **1040**) without introducing significant delay into the system during backup operation and/or restore operations. Running the volatile memory subsystem **1030** at the reduced frequency during a backup and/or restore operation may advantageously reduce overall power consumption of the memory system **1010**.

In one embodiment, the backup clock **1408** and the volatile memory system clock signal **1420** are received by a multiplexer **1422**, as schematically illustrated by FIG. **19**. The multiplexer **1422** can output either the volatile memory system clock signal **1420** or the backup clock signal **1408** depending on the backup state of the memory system **1010**. For example, when the memory system **1010** is not performing a backup or restore operation and is communicating with the host system (e.g., normal operation), the volatile memory system clock signal **1420** may be provided by the multiplexer **422** to the volatile memory PLL block **1424**. When the memory system **1010** is performing a backup (or restore) operation, the backup clock signal **1408** may be provided.

The volatile memory PLL block **1424** receives the volatile memory reference clock signal **1423** from the multiplexer **1422** and can generate a series of clock signals which are distributed to the volatile memory subsystem **1030** and associated control logic. For example, in one embodiment, the PLL block **1424** generates a series of clock signals **1426** which clock the volatile memory elements **1032**. A clock signal **1428** may be used to clock control logic associated with the volatile memory elements, such as one or more registers (e.g., the one or more registers of a registered DIMM). Another clock signal **1430** may be sent to the controller **1062**. A feedback clock signal **1432** is fed back into the volatile memory PLL block **1424**. In one embodiment, the PLL block **1424** compares the feedback clock signal **1432** to the reference clock signal **1423** and varies the phase and frequency of its output until the reference clock signal **1423** and the feedback clock signal **1432** clocks are phase and frequency matched.

The clock signal **1430** may be used by the controller **1062** to generate and distribute clock signals which will be used by controller logic which is configured to control the volatile memory subsystem **1030**. For example, control logic in the controller **1062** may be used to control the volatile memory subsystem **1030** during a backup or restore operation. The clock signal **1430** may be used as a reference clock signal for the PLL block **1434** which can generate one or more clocks **1438** used by logic in the controller **1062**. For example, the PLL block **1434** may generate one or more clock signals **1438** used to drive logic circuitry associated with controlling the volatile memory subsystem **1030**. In certain embodiments, the PLL block **1434** includes a feedback clock signal **1436** and operates in a similar manner to other PLL blocks described herein.

The clock signal **1430** may be used as a reference clock signal for the PLL block **1440** which may generate one or more clock signals used by a sub-block **1442** to generate one or more other clock signals **1444**. In one embodiment, for example, the volatile memory subsystem **1030** comprises DDR2 SDRAM elements and the sub-block **1442** generates one or more DDR2 compatible clock signals **1444**. A feedback clock signal **1446** is fed back into the PLL block **1440**. In certain embodiments, the PLL block **1440** operates in a similar manner to other PLL blocks described herein.

While described with respect to the example embodiment of FIG. **19**, various alternative clock distribution topologies are possible. For example, one or more of the clock signals have a different frequency in various other embodiments. In some embodiments, one or more of the clocks shown as differential signals are single ended signals. In one embodiment, the volatile memory subsystem **1030** operates on the volatile memory clock signal **1420** and there is no backup clock signal **1408**. In some embodiments, the volatile memory subsystem **1030** is operated at a reduced frequency during a backup operation and not during a restore operation. In other embodiments, the volatile memory subsystem **1030** is operated at a reduced frequency during a restore operation and not during a backup operation.

FIG. **20** is a flowchart of an example method **1500** of controlling a memory system **1010** operatively coupled to a host system. Although described with respect to the memory system **1010** described herein, the method **1500** is compatible with other memory systems. The memory system **1010** may include a clock distribution topology **1400** similar to the one described above with respect to FIG. **19** or another clock distribution topology. The memory system **1010** can include a volatile memory subsystem **30** and a non-volatile memory subsystem **1040**.

In an operational block **1510**, the method **1500** comprises operating the volatile memory subsystem **1030** at a first frequency when the memory system **1010** is in a first mode of operation in which data is communicated between the volatile memory subsystem **1030** and the host system. In an operational block **1520**, the method **1500** comprises operating the non-volatile memory subsystem **1040** at a second frequency when the memory system **1010** is in a second mode of operation in which data is communicated between the volatile memory subsystem **1030** and the non-volatile memory subsystem **1040**. The method **1500** further comprises operating the volatile memory subsystem **1030** at a

third frequency in an operational block **1530** when the memory system **1010** is in the second mode of operation. In certain embodiments, the memory system **1010** is not powered by a battery when it is in the second mode of operation. The memory system **1010** may switch from the first mode of operation to the second mode of operation in response to a trigger condition. The trigger condition may be any trigger condition described herein such as, for example, a power failure condition. In certain embodiments, the second mode of operation includes both backup and restore operations as described herein. In other embodiments, the second mode of operation includes backup operations but not restore operations. In yet other embodiments, the second mode of operation includes restore operations but not backup operations.

The third frequency can be less than the first frequency. For example, the third frequency can be approximately equal to the second frequency. In certain embodiments, the reduced frequency operation is an optional mode. In yet other embodiments, the first, second and/or third frequencies are configurable by a user or by the memory system **1010**.

FIG. **21** schematically illustrates an example topology of a connection to transfer data slices from two DRAM segments **1630**, **1640** of a volatile memory subsystem **1030** of a memory system **1010** to a controller **1062** of the memory system **1010**. While the example of FIG. **21** shows a topology including two DRAM segments **1630**, **1640** for the purposes of illustration, each address location of the volatile memory subsystem **1030** comprises more than the two segments in certain embodiments. The data lines **1632**, **1642** from the first DRAM segment **1630** and the second DRAM segment **1640** of the volatile memory subsystem **1030** are coupled to switches **1650**, **1652** which are coupled to the controller **1062** (e.g., logic element **1070**) of the memory system **1010**. The chip select lines **1634**, **1644** and the self-refresh lines **1636**, **1646** (e.g., CKe signals) of the first and second DRAM segments **1630**, **1640**, respectively, are coupled to the controller **1062**. In certain embodiments, the controller **1062** comprises a buffer (not shown) which is configured to store data from the volatile memory subsystem **1030**. In certain embodiments, the buffer is a first-in, first out buffer (FIFO). In certain embodiments, data slices from each DRAM segment **1630**, **1640** comprise a portion of the volatile memory subsystem data bus. In one embodiment, for example, the volatile memory subsystem **1030** comprises a 72-bit data bus (e.g., each data word at each addressable location is 72 bits wide and includes, for example, 64 bits of accessible SDRAM and 8 bits of ECC), the first data slice from the first DRAM segment **1630** may comprise 40 bits of the data word, and the second data slice from the second DRAM segment **1640** may comprise the remaining 32 bits of the data word. Certain other embodiments comprise data buses and/or data slices of different sizes.

In certain embodiments, the switches **1650**, **1652** can each be selectively switched to selectively operatively couple the data lines **1632**, **1642**, respectively from the first and second DRAM segments **1630**, **1640** to the controller **1062**. The chip select lines **1634**, **1644** enable the first and second DRAM segments **1630**, **1640**, respectively, of the volatile memory subsystem **1030**, and the self-refresh lines **1636**, **1646** toggle the first and second DRAM segments **1630**, **1640**, respectively, from self-refresh mode to active mode. In certain embodiments, the first and second DRAM segments **1630**, **1640** maintain stored information but are not accessible when they are in self-refresh mode, and maintain stored information and are accessible when they are in active mode.

In certain embodiments, when the memory system **1010** is backing up the volatile memory system **1030**, data slices from only one of the two DRAM segments **1630**, **1640** at a time are sent to the controller **1062**. For example, when the first slice is being written to the controller **1062** during a back-up, the controller **1062** sends a signal via the CKe line **1636** to the first DRAM segment **1630** to put the first DRAM segment **1630** in active mode. In certain embodiments, the data slice from the first DRAM segment **1630** for multiple words (e.g., a block of words) is written to the controller **1062** before writing the second data slice from the second DRAM segment **1640** to the controller **1062**. While the first data slice is being written to the controller **1062**, the controller **1062** also sends a signal via the CKe line **1646** to put the second DRAM segment **1640** in self-refresh mode. Once the first data slice for one word or for a block of words is written to the controller **1062**, the controller **1062** puts the first DRAM segment **1630** into self-refresh mode by sending a signal via the CKe line **1636** to the first DRAM segment **1640**. The controller **1062** also puts the second DRAM segment **1640** into active mode by sending a signal via the CKe line **1646** to the DRAM segment **1640**. The second slice for a word or for a block of words is written to the controller **1062**. In certain embodiments, when the first and second data slices are written to the buffer in the controller **1062**, the controller **1062** combines the first and second data slices **1630**, **1640** into complete words or blocks of words and then writes each complete word or block of words to the non-volatile memory subsystem **1040**. In certain embodiments, this process is called "slicing" the volatile memory subsystem **1030**.

In certain embodiments, the data may be sliced in a restore operation as well as, or instead of, during a backup operation. For example, in one embodiment, the nonvolatile memory elements **1042** write each backed-up data word to the controller **1062** which writes a first slice of the data word to the volatile memory subsystem **1030** and then a second slice of the data word to the volatile memory subsystem **1030**. In certain embodiments, slicing the volatile memory subsystem **1030** during a restore operation may be performed in a manner generally inverse to slicing the volatile memory subsystem **1030** during a backup operation.

FIG. **22** is a flowchart of an example method **1600** of controlling a memory system **1010** operatively coupled to a host system and which includes a volatile memory subsystem **1030** and a non-volatile memory subsystem **1040**. Although described with respect to the memory system **1010** described herein with respect to FIGS. **12-14** and **21**, the method **1600** is compatible with other memory systems. The method **1600** comprises communicating data words between the volatile memory subsystem **1030** and the host system when the memory system **1010** is in a first mode of operation in an operational block **1610**. For example, the memory system **1010** may be in the first mode of operation when no trigger condition has occurred and the memory system is not performing a backup and/or restore operation or is not being powered by a secondary power supply.

In an operational block **1620**, the method further comprises transferring data words from the volatile memory subsystem **1030** to the non-volatile memory subsystem **1040** when the memory system **1010** is in a second mode of operation. In certain embodiments, each data word comprises the data stored in a particular address of the memory system **1010**. The memory system **1010** may enter the second mode of operation, for example, when a trigger condition (e.g., a power failure) occurs. In certain embodiments, transferring each data word comprises storing a first

portion (also referred to as a slice) of the data word in a buffer in an operational block **1622**, storing a second portion of the data word in the buffer in an operational block **1624**, and writing the entire data word from the buffer to the non-volatile memory subsystem **1040** in an operational block **1626**.

In one example embodiment, the data word may be a 72 bit data word (e.g., 64 bits of accessible SDRAM and 8 bits of ECC), the first portion (or "slice") may comprise 40 bits of the data word, and the second portion (or "slice") may comprise the remaining 32 bits of the data word. In certain embodiments, the buffer is included in the controller **1062**. For example, in one embodiment, the buffer is a first-in, first-out buffer implemented in the controller **1062** which comprises an FPGA. The method **1600** may generally be referred to as "slicing" the volatile memory during a backup operation. In the example embodiment, the process of "slicing" the volatile memory during a backup includes bringing the 32-bit slice out of self-refresh, reading a 32-bit block from the slice into the buffer, and putting the 32-bit slice back into self-refresh. The 40-bit slice is then brought out of self-refresh and a 40-bit block from the slice is read into a buffer. Each block may comprise a portion of multiple words. For example, each 32-bit block may comprise 32-bit portions of multiple 72-bit words. In other embodiments, each block comprises a portion of a single word. The 40-bit slice is then put back into self-refresh in the example embodiment. The 32-bit and 40-bit slices are then combined into a 72-bit block by the controller **1062** and ECC detection/correction is performed on each 72-bit word as it is read from the buffer and written into the non-volatile memory subsystem (e.g., flash).

In some embodiments, the entire data word may comprise more than two portions. For example, the entire data word may comprise three portions instead of two and transferring each data word further comprises storing a third portion of each data word in the buffer. In certain other embodiments, the data word may comprise more than three portions.

In certain embodiments, the data may be sliced in a restore operation as well as, or instead of, during a backup operation. For example, in one embodiment, the nonvolatile memory elements **1040** write each backed-up data word to the controller **1062** which writes a first portion of the data word to the volatile memory subsystem **1030** and then a second portion of the data word to the volatile memory **1030**. In certain embodiments, slicing the volatile memory subsystem **1030** during a restore operation may be performed in a manner generally inverse to slicing the volatile memory subsystem **1030** during a backup operation.

The method **1600** can advantageously provide significant power savings and can lead to other advantages. For example, in one embodiment where the volatile memory subsystem **1030** comprises DRAM elements, only the slice of the DRAM which is currently being accessed (e.g., written to the buffer) during a backup is configured in full-operational mode. The slice or slices that are not being accessed may be put in self-refresh mode. Because DRAM in self-refresh mode uses significantly less power than DRAM in full-operational mode, the method **1600** can allow significant power savings. In certain embodiments, each slice of the DRAM includes a separate self-refresh enable (e.g., CKe) signal which allows each slice to be accessed independently.

In addition, the connection between the DRAM elements and the controller **1062** may be as large as the largest slice instead of as large as the data bus. In the example embodiment, the connection between the controller **1062** and the DRAM may be 40 bits instead of 72 bits. As a result, pins on the controller **1062** may be used for other purposes or a smaller controller may be used due to the relatively low number of pin-outs used to connect to the volatile memory subsystem **1030**. In certain other embodiments, the full width of the data bus is connected between the volatile memory subsystem **1030** and the controller **1062** but only a portion of it is used during slicing operations. For example, in some embodiments, memory slicing is an optional mode.

While embodiments and applications have been shown and described, it would be apparent to those skilled in the art having the benefit of this disclosure that many more modifications than mentioned above are possible without departing from the inventive concepts disclosed herein. The invention, therefore, is not to be restricted except in the spirit of the appended claims.

What is claimed is:

1. A memory module comprising:

a printed circuit board (PCB) having an interface configured to fit into a corresponding slot connector of a host system, the interface including a plurality of edge connections configured to couple power, data, address and control signals between the memory module and the host system;

a voltage conversion circuit coupled to the PCB and configured to provide at least three regulated voltages, wherein the voltage conversion circuit includes at least three buck converters each of which is configured to produce a regulated voltage of the at least three regulated voltages;

a plurality of components coupled to the PCB, each component of the plurality of components coupled to at least one regulated voltage of the at least three regulated voltages, the plurality of components including a plurality of synchronous dynamic random access memory (SDRAM) devices and a first circuit that is coupled to the plurality of SDRAM devices and to a first set of edge connections of the plurality of edge connections, wherein the first circuit is coupled to first and second regulated voltages of the at least three regulated voltages, and wherein the plurality of SDRAM devices are coupled to the first regulated voltage of the at least three regulated voltages.

2. The memory module of claim **1**, wherein the first regulated voltage has a first voltage amplitude, and the second regulated voltage has a second voltage amplitude; and wherein a first one of the first and second voltage amplitudes is less than a second one of the first and second voltage amplitudes.

3. The memory module of claim **1**, wherein a third regulated voltage of the at least three regulated voltages has a voltage amplitude of 1.8 volts.

4. The memory module of claim **1**, further comprising:

a voltage monitor circuit coupled to the PCB and to a second set of edge connections of the plurality of edge connections, the voltage monitor circuit configured to monitor an input voltage received from the second set of edge connections, the voltage monitor circuit configured to produce a trigger signal in response to the input voltage having a voltage amplitude below a predetermined threshold voltage, wherein the memory module transitions from a first operable state to a second operable state in response to the trigger signal.

5. The memory module of claim **4**, further comprising:

a controller coupled to the voltage monitor circuit; wherein, in response to the trigger signal, the controller

is configured to perform one or more operations including a write operation to transfer data to non-volatile memory.

6. The memory module of claim **1**, further comprising:

a voltage monitor circuit coupled to the PCB and to a second set of edge connections of the plurality of edge connections, the voltage monitor circuit configured to monitor an input voltage received from the second set of edge connections, the voltage monitor circuit configured to produce a trigger signal in response to the input voltage having a voltage amplitude above a predetermined threshold voltage, wherein the memory module transitions from a first operable state to a second operable state in response to the trigger signal.

7. The memory module of claim **6**, further comprising:

a controller coupled to the voltage monitor circuit; wherein, in response to the trigger signal, the controller is configured to perform one or more operations including a write operation to transfer data to non-volatile memory.

8. The memory module of claim **1**, further comprising:

a controller coupled to the PCB, the controller including a voltage monitor circuit configured to monitor an input voltage received from a second set of edge connections of the plurality of edge connections, wherein, in response to the voltage monitor circuit detecting a power threshold condition, the voltage monitor circuit transmits a signal to one or more portions of the controller.

9. The memory module of claim **8**, wherein the voltage monitor circuit detecting a power threshold condition includes the voltage monitor circuit detecting an amplitude of the input voltage being above a first predetermined threshold voltage, and wherein the first predetermined threshold voltage is above a specified operating voltage.

10. The memory module of claim **9**, wherein the first predetermined threshold voltage is ten percent above the specified operating voltage.

11. The memory module of claim **9**, wherein the voltage monitor circuit detecting a power threshold condition includes the voltage monitor circuit detecting an amplitude of the input voltage being below a second predetermined threshold voltage, wherein the second predetermined threshold voltage is below the specified operating voltage, and wherein the memory module transitions from a first operable state to a second operable state in response to the signal.

12. The memory module of claim **11**, wherein the second predetermined threshold voltage is ten percent below the specified operating voltage.

13. The memory module of claim **8**, wherein the voltage monitor circuit detecting a power threshold condition includes the voltage monitor circuit detecting a power reduction condition or a low voltage condition of the input voltage.

14. The memory module of claim **8**, wherein the voltage monitor circuit detecting a power threshold condition includes the voltage monitor circuit detecting a request by the host system.

15. The memory module of claim **1**, wherein two of the at least three buck converters are configured to operate as a dual-buck converter.

16. A memory module comprising:

a printed circuit board (PCB) having an interface configured to fit into a corresponding slot connector of a host system;

a voltage conversion circuit coupled to the PCB and configured to provide a plurality of regulated voltages, wherein the voltage conversion circuit includes three buck converters each of which is configured to produce a regulated voltage of the plurality of regulated voltages;

a plurality of components coupled to the PCB, each component of the plurality of components coupled to at least one regulated voltage of the plurality of regulated voltages, the plurality of components including a plurality of synchronous dynamic random access memory (SDRAM) devices, the plurality of SDRAM devices coupled to a first regulated voltage of the plurality of regulated voltages; and

a controller coupled to the PCB, the controller including a voltage monitor circuit coupled to an input voltage received from the host system via the interface, the voltage monitor circuit configured to detect an amplitude change in the input voltage, wherein, in response to the voltage monitor detecting an amplitude change in the input voltage, the memory module transitions from a first operable state to a second operable state.

17. The memory module of claim **16**, wherein the voltage monitor circuit is configured to detect the input voltage being above a first predetermined threshold voltage and to detect the input voltage being below a second predetermined threshold voltage.

18. The memory module of claim **16**, wherein, in the first operable state, the voltage conversion circuit provides the first regulated voltage to the plurality of SDRAM devices using a first pre-regulated voltage, and wherein, in the second operable state, the voltage conversion circuit provides the first regulated voltage to the plurality of SDRAM devices using a second pre-regulated voltage.

19. The memory module of claim **18**, wherein, in the first operable state, the first pre-regulated voltage is supplied to the voltage conversion circuit via a circuit, and wherein, in the second operable state, the second pre-regulated voltage is supplied to the voltage conversion circuit via the circuit.

20. The memory module of claim **19**, wherein the circuit includes a first diode having a first input and a first output, the first input is coupled to the first pre-regulated voltage and the first output is coupled to the voltage conversion circuit, and wherein the circuit includes a second diode having a second input and a second output, the second input is coupled to the second pre-regulated voltage and the second output is coupled to the first output and to the voltage conversion circuit.

21. The memory module of claim **16**, further comprising:

a first circuit having a first input, a second input and an output, the first input coupled to a first pre-regulated voltage, the second input coupled to a second pre-regulated voltage, the output coupled to the voltage conversion circuit, wherein the voltage conversion circuit provides the first regulated voltage to the plurality of SDRAM devices using a supply voltage, wherein, in the first operable state, the first circuit provides the supply voltage to the voltage conversion circuit via the output using the first pre-regulated voltage, and wherein, in the second state, the first circuit provides the supply voltage to the voltage conversion circuit via the output using the second pre-regulated voltage.

22. The memory module of claim **21**, wherein the first circuit includes a first diode coupled between the first input and the output and a second diode coupled between the second input and the output.

23. The memory module of claim **16**, wherein the voltage monitor circuit is configured to produce a trigger signal in response to detecting an amplitude change in the input

voltage; and wherein, in response to the trigger signal, the controller is configured to perform one or more operations including a write operation to transfer data to non-volatile memory.

24. A memory module comprising:

a printed circuit board (PCB) having an interface configured to fit into a corresponding slot connector of a host system;

a voltage conversion circuit configured to provide a plurality of regulated voltages, wherein the voltage conversion circuit includes three buck converters each of which is configured to produce a regulated voltage of the plurality of regulated voltages;

a plurality of components each coupled to at least one regulated voltage of the plurality of regulated voltages, the plurality of components including a plurality of synchronous dynamic random access memory (SDRAM) devices, wherein the plurality of SDRAM devices are coupled to a first regulated voltage of the plurality of regulated voltages; and

a controller including a voltage monitor circuit coupled to an input voltage received from the host system via the interface of the PCB, the voltage monitor circuit configured to monitor the input voltage, wherein the controller is configured to perform one or more operations in response to the voltage monitor circuit detecting an amplitude change in the input voltage, and wherein the one or more operations include a write operation to transfer data into non-volatile memory.

25. The memory module of claim **24**, wherein, in response to the voltage monitor circuit detecting an amplitude change in the input voltage, the memory module transitions from a first operable state to a second operable state.

26. The memory module of claim **25**, wherein, in the first operable state, the voltage conversion circuit provides the first regulated voltage to the plurality of SDRAM devices using a first pre-regulated voltage, and wherein, in the second operable state, the voltage conversion circuit provides the first regulated voltage to the plurality of SDRAM devices using a second pre-regulated voltage.

27. The memory module of claim **26**, wherein the first pre-regulated voltage is coupled to the voltage conversion circuit via a first diode.

28. The memory module of claim **27**, wherein the second pre-regulated voltage is coupled to the voltage conversion circuit via a second diode.

29. The memory module of claim **24**, wherein the voltage monitor circuit is configured to detect the input voltage being above a first predetermined threshold voltage or below a second predetermined threshold voltage.

30. The memory module of claim **29**, wherein the first predetermined threshold voltage is above a specified operating voltage, and the second predetermined threshold voltage is below the specified operating voltage.

* * * * *

US008787060B2

(12) **United States Patent**
Lee

(10) **Patent No.: US 8,787,060 B2**
(45) **Date of Patent: Jul. 22, 2014**

(54) **METHOD AND APPARATUS FOR OPTIMIZING DRIVER LOAD IN A MEMORY PACKAGE**

(75) Inventor: **Hyun Lee**, Ladera Ranch, CA (US)

(73) Assignee: **Netlist, Inc.**, Irvine, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 103 days.

(21) Appl. No.: **13/288,850**

(22) Filed: **Nov. 3, 2011**

(65) **Prior Publication Data**

US 2012/0106228 A1 May 3, 2012

Related U.S. Application Data

(60) Provisional application No. 61/409,893, filed on Nov. 3, 2010.

(51) **Int. Cl.**
G11C 5/06 (2006.01)

(52) **U.S. Cl.**
USPC **365/63**; 365/51

(58) **Field of Classification Search**
USPC 365/63, 51, 52
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,243,283	B1	6/2001	Bertin et al.
6,551,857	B2	4/2003	Leedy
7,098,541	B2	8/2006	Adelmann
7,200,021	B2	4/2007	Raghuram
7,254,036	B2	8/2007	Pauley et al.
7,269,042	B2	9/2007	Kinsley et al.
7,286,436	B2	10/2007	Bhakta et al.
7,289,386	B2	10/2007	Bhakta et al.
7,375,970	B2	5/2008	Pauley et al.
7,442,050	B1	10/2008	Bhakta et al.
7,532,537	B2	5/2009	Solomon et al.
7,619,893	B1	11/2009	Yu
7,619,912	B2	11/2009	Bhakta et al.
7,630,202	B2	12/2009	Pauley et al.

(Continued)

FOREIGN PATENT DOCUMENTS

EP	1 816 570 A2	8/2007
WO	WO 2010-138480	12/2010
WO	WO 2011-049710	4/2011
WO	WO 2011-094437	8/2011

OTHER PUBLICATIONS

International Search Report and Written Opinion, PCT/US2011/059209, Jan. 31, 2013.

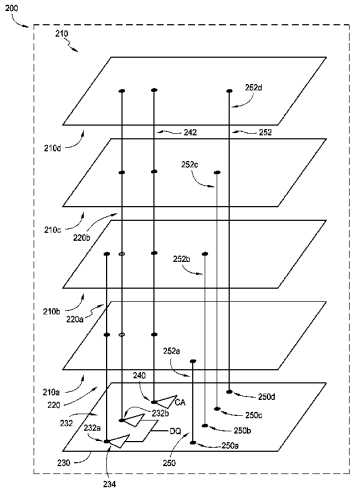
(Continued)

Primary Examiner — Huan Hoang
(74) *Attorney, Agent, or Firm* — Jamie J. Zheng, Esq.

(57) **ABSTRACT**

An apparatus is provided that includes a plurality of array dies and at least two die interconnects. The first die interconnect is in electrical communication with a data port of a first array die and a data port of a second array die and not in electrical communication with data ports of a third array die. The second die interconnect is in electrical communication with a data port of the third array die and not in electrical communication with data ports of the first array die and the second array die. The apparatus includes a control die that includes a first data conduit configured to transmit a data signal to the first die interconnect and not to the second die interconnect, and at least a second data conduit configured to transmit the data signal to the second die interconnect and not to the first die interconnect.

34 Claims, 8 Drawing Sheets



(56) **References Cited**

U.S. PATENT DOCUMENTS

7,633,165	B2	12/2009	Hsu et al.	
7,636,274	B2	12/2009	Solomon et al.	
7,683,459	B2	3/2010	Ma et al.	
7,811,097	B1	10/2010	Bhakta et al.	
7,827,348	B2	11/2010	Lee et al.	
7,830,692	B2	11/2010	Chung et al.	
7,839,645	B2	11/2010	Pauley et al.	
7,864,627	B2	1/2011	Bhakta et al.	
7,881,150	B2	2/2011	Solomon et al.	
7,894,229	B2	2/2011	Lahtinen et al.	
7,894,230	B2	2/2011	Kim	
7,978,721	B2	7/2011	Jeddeloh et al.	
7,990,171	B2	8/2011	Chung et al.	
7,999,367	B2	8/2011	Kang et al.	
8,001,434	B1	8/2011	Lee et al.	
8,019,589	B2	9/2011	Rajan et al.	
8,033,836	B1	10/2011	Bhakta et al.	
2006/0233012	A1	10/2006	Sekiguchi et al.	
2006/0259678	A1 *	11/2006	Gervasi	711/2
2007/0096332	A1	5/2007	Satoh et al.	
2008/0025123	A1	1/2008	Rajan et al.	
2008/0025137	A1 *	1/2008	Rajan et al.	365/239
2008/0094808	A1	4/2008	Kanapathippillai et al.	
2008/0253085	A1	10/2008	Soffer	
2008/0296779	A1	12/2008	Matsui et al.	
2009/0070727	A1	3/2009	Solomon	
2009/0103345	A1	4/2009	McLaren et al.	
2009/0290442	A1	11/2009	Rajan	
2010/0020583	A1	1/2010	Kang et al.	
2010/0090338	A1	4/2010	Lee et al.	
2010/0174858	A1	7/2010	Chen et al.	
2011/0006360	A1	1/2011	Ikebuchi	
2011/0016250	A1	1/2011	Lee et al.	
2011/0016269	A1	1/2011	Lee et al.	
2011/0050320	A1	3/2011	Gillingham	
2011/0108888	A1	5/2011	Or-Bach et al.	
2011/0125982	A1	5/2011	Choi et al.	
2011/0156232	A1	6/2011	Youn et al.	
2011/0169171	A1	7/2011	Marcoux	
2011/0193226	A1 *	8/2011	Kirby et al.	257/738

OTHER PUBLICATIONS

Ahmad et al., "Modeling of peak-to-peak switching noise along a vertical chain of power distribution TSV pairs in a 3D stack of ICs interconnected through TSVs," Norchip Conference, Article No. 5669473, Nov. 15-16, 2010, IEEE Computer Society.

Black et al., "Die Stacking (3D) Microarchitecture," MICRO-39, 39th Annual IEEE/ACM International Symposium on, Dec. 2006, 469-479, Orlando, FL.

Daneshtalab et al., "CMIT—A novel cluster-based topology for 3D stacked architectures," 3D Systems Integration Conference (3DIC), 2010 IEEE International, Nov. 16-18, 2010, pp. 1-5.

Funaya et al., "Cache partitioning strategies for 3-D stacked vector processors," IEEE 3D System Integration Conference, article No. 5751453, Nov. 16-18, 2010, IEEE Computer Society.

Ghosh et al., "Smart Refresh: An Enhancement Memory Controller Design for Reducing Energy in Conventional and 3D Die-Stacked DRAMs," Microarchitecture, 2007, 40th Annual IEEE/ACM International Symposium, Dec. 1-5, 2007, pp. 134-145.

Kang et al., "8 Gb 3-D DDR3 DRAM using through-silicon-via technology," IEEE Journal of Solid-State Circuits, v 45, n. 1, 111-19, Jan. 2010, IEEE, USA.

U.S. Appl. No. 12/422,912, filed Apr. 13, 2009, Hyun Lee, et al.

U.S. Appl. No. 12/422,853, filed Apr. 13, 2009, Hyun Lee, et al.

U.S. Appl. No. 12/815,339, filed Jun. 14, 2010, Hyun Lee.

Kang et al., "Signal integrity and reliability of a new Multi-Stack Package using a Pressure Conductive Rubber," Electrical Design of Advanced Packaging and Systems Symposium, 214-17, Dec. 2008, Seoul, South Korea.

Kawano, "A 3D Packaging Technology for High-Density Stacked DRAM," VLSI Technology, Systems and Applications, 2007, Apr. 23-25, 2007, pp. 1-2.

Kurita et al., "A 3-D packaging technology with highly-parallel memory/logic interconnect," IEICE Transactions on Electronics, v E92-C, No. 12, pp. 1512-1522, 2009, Maruzen Col, Ltd.

Kurita et al., "Vertical Integration of Stacked DRAM and High-Speed Logic Device Using SMAFTI Technology," Advanced Packaging, IEEE Transactions, Aug. 2009, vol. 32 Issue 3, pp. 657-665.

"Posts Tagged '3D Stacking'," from Chip Design Mag., http://chipdesignmag.com/lpd/blog/tag/3d-stacking/ (Printed Oct. 13, 2011).

Russell, Gill, "Intel Micron Hybrid Memory Cube: The Future of Exascale Computing," Bright Side of News. Sep. 19, 2011, <http://www.brightsideofnews.com/news/2011/9/19/intel-micron-hybrid-memory-cube-the-future-of-exascale-computing.aspx> Printed Oct. 13, 2011 in 8 pages.

Loh, "3D-Stacked Memory Architectures for Multi-core Processors," Proceedings of the 35th Annual International Symposium on Computer Architecture (ISCA '08), 453-464, IEEE Computer Society, Washington DC, USA.

Val, "The 3D interconnection applications for mass memories and microprocessors," Proceedings of the Technical Conference, 1991 International Electronic Packaging Conference, 851-60, vol. 2, 1991, Int. Electron. Packaging Soc, Wheaton, IL, USA, 2008.

Weis et al., "Design space exploration for 3D-stacked DRAMs," Proceedings—Design, Automation and Test in Europe Conference and Exhibition, 2011.

Zhang et al. "A Customized Design of DRAM Controller for On-Chip 3D DRAM Stacking" Custom Integrated Circuits Conference (CICC), 2010 IEEE, Sep. 19-22, 2010, pp. 1-4.

* cited by examiner

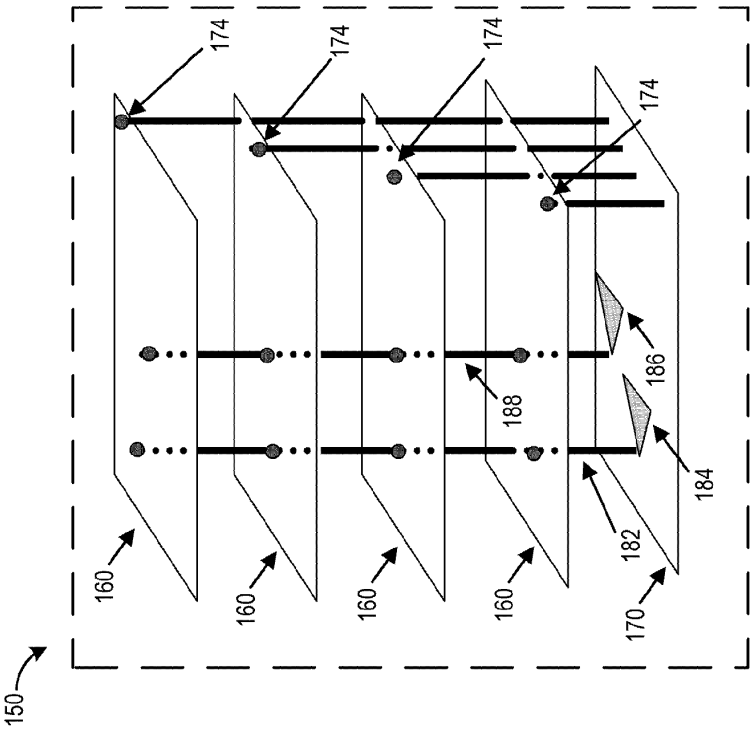

150
160
160
160
160
170
182
184
186
188
174
174
174
174

FIG. 1B

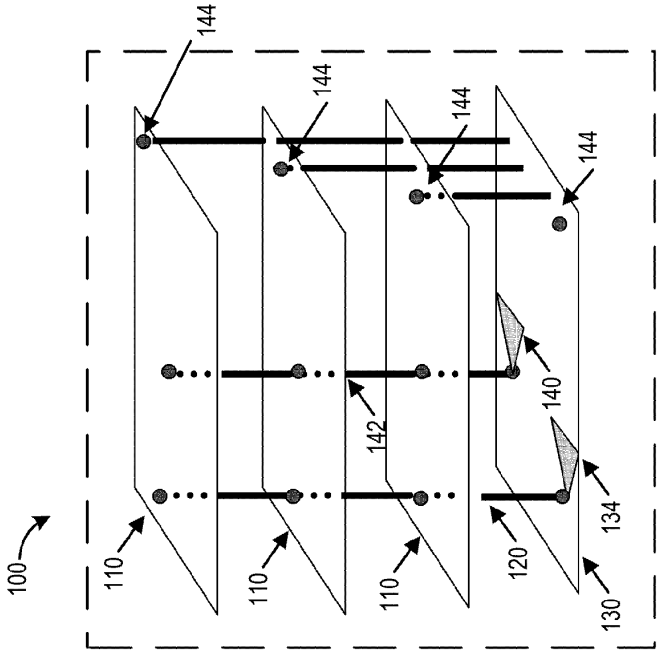

100
110
110
110
120
130
134
140
142
144
144
144
144

FIG. 1A

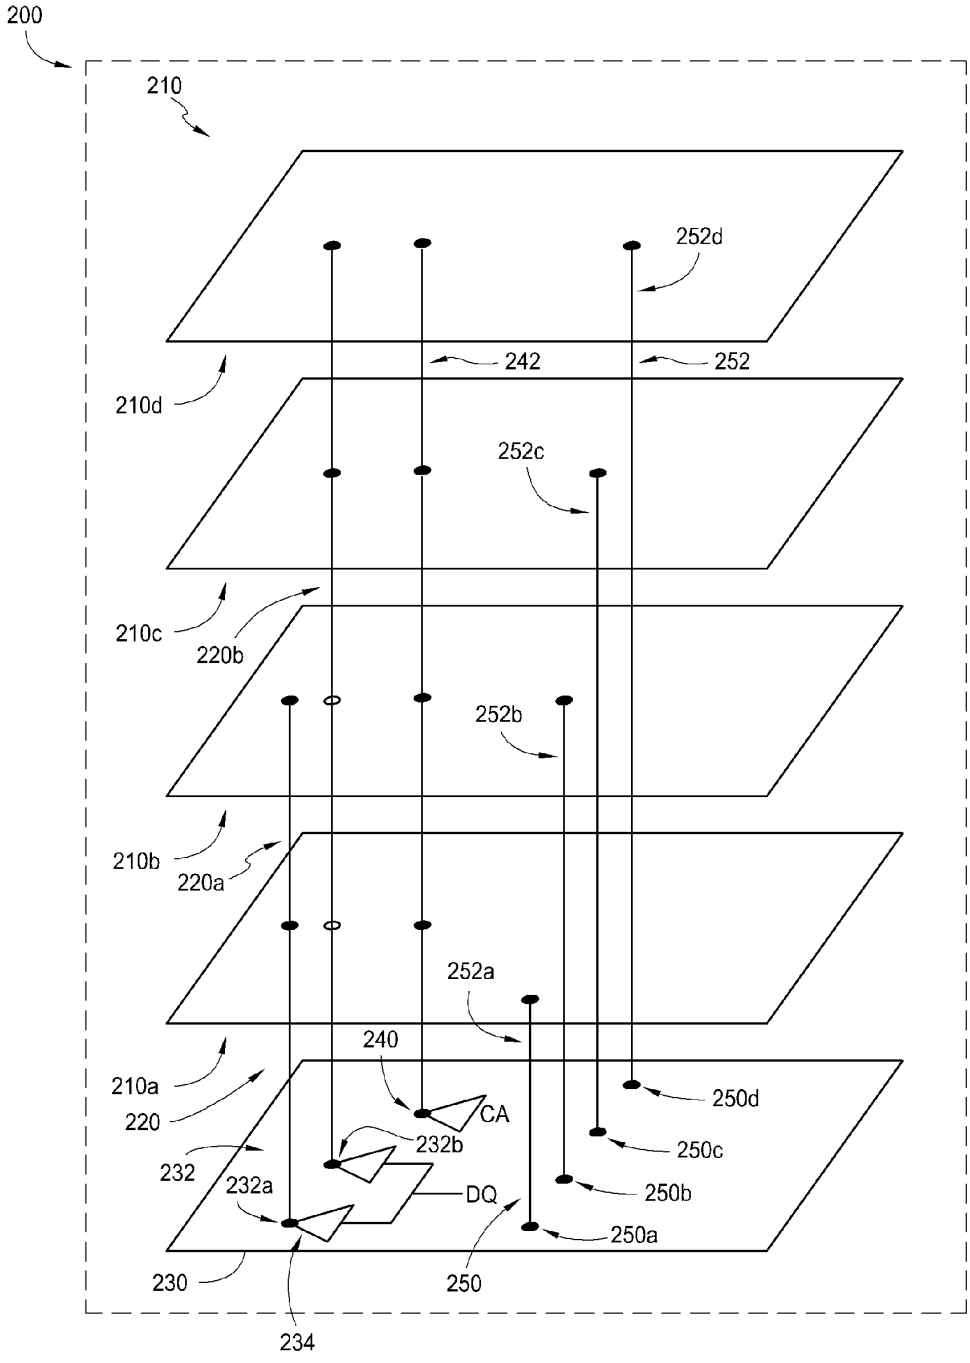

200
210
252d
242
252
210d
252c
210c
220b
252b
210b
220a
252a
240
210a
220
CA
250d
232b
250c
232
232a
DQ
250b
250a
230
250
234

FIG. 2

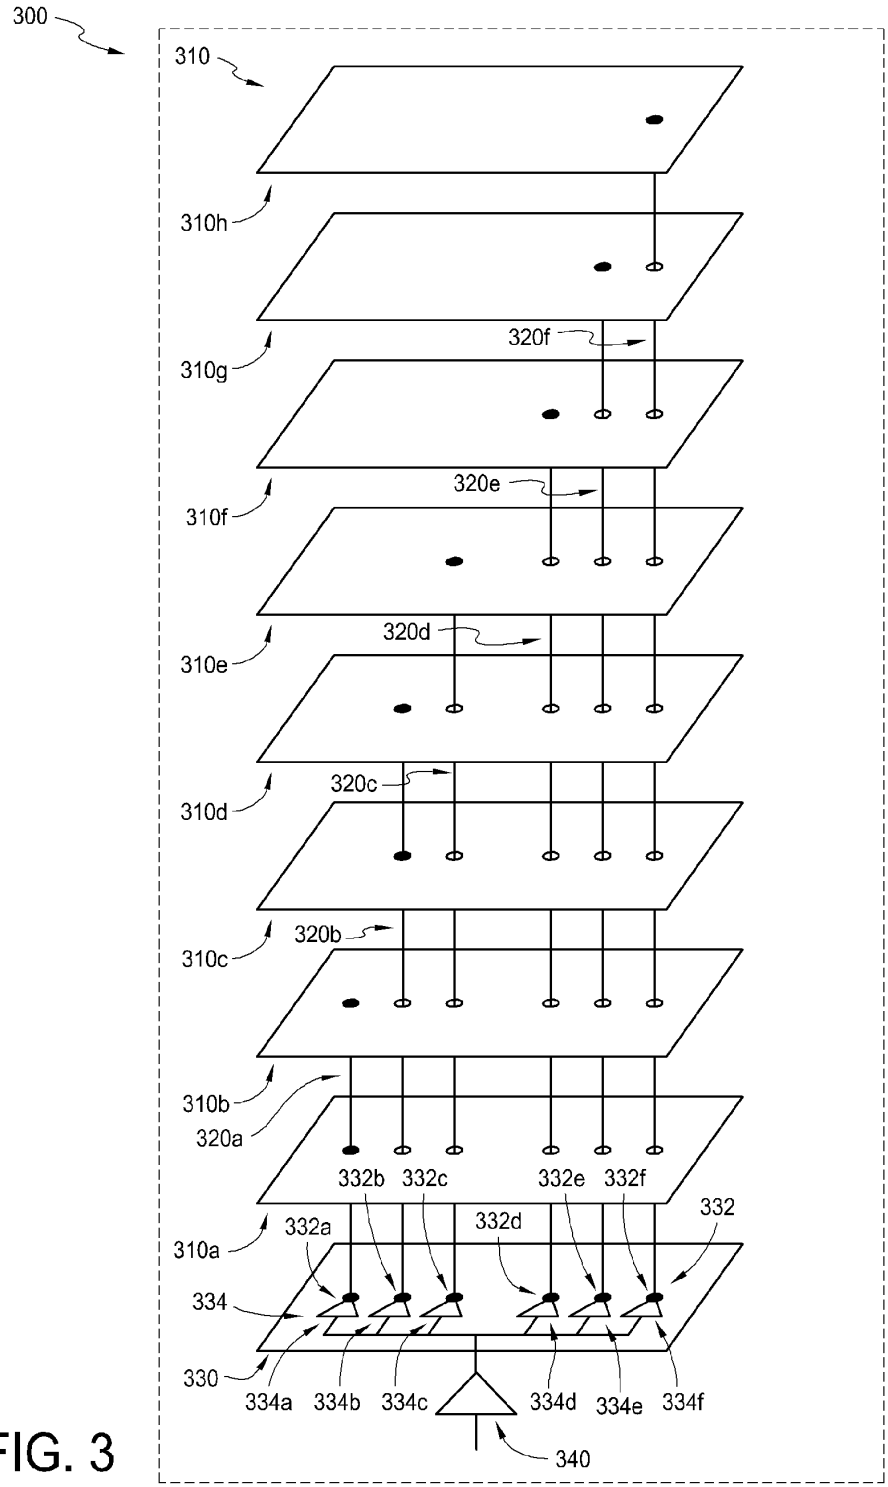

300
310
310h
320f
310g
320e
310f
320d
310e
320c
310d
320b
310c
310b
320a
332b
332c
332e
332f
332
332a
332d
310a
334
330
334a
334b
334c
334d
334e
334f
340

FIG. 3

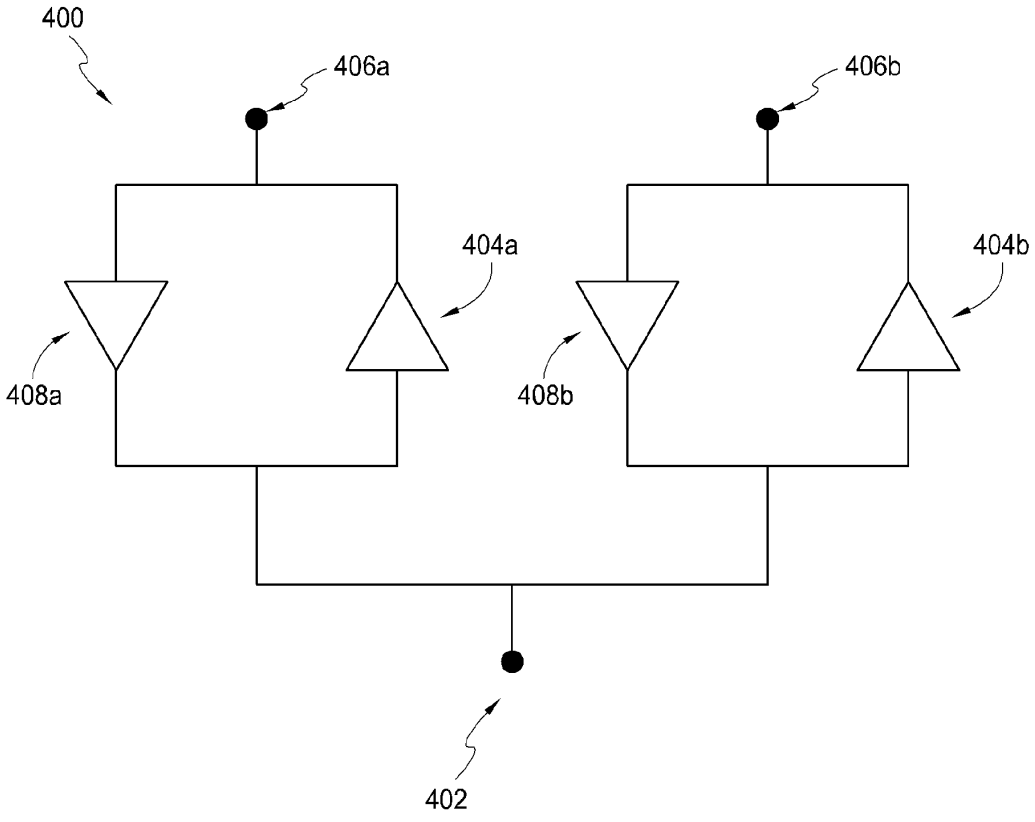


FIG. 4

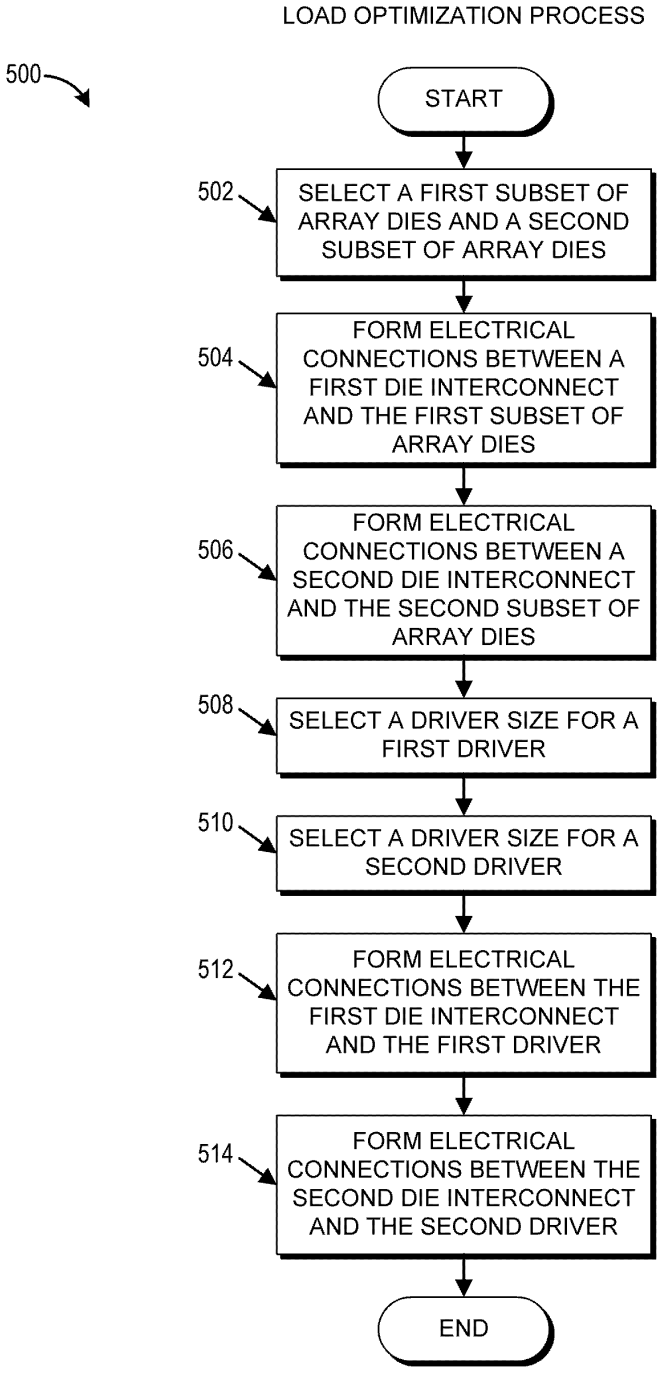

LOAD OPTIMIZATION PROCESS
500
START
502
SELECT A FIRST SUBSET OF ARRAY DIES AND A SECOND SUBSET OF ARRAY DIES
504
FORM ELECTRICAL CONNECTIONS BETWEEN A FIRST DIE INTERCONNECT AND THE FIRST SUBSET OF ARRAY DIES
506
FORM ELECTRICAL CONNECTIONS BETWEEN A SECOND DIE INTERCONNECT AND THE SECOND SUBSET OF ARRAY DIES
508
SELECT A DRIVER SIZE FOR A FIRST DRIVER
510
SELECT A DRIVER SIZE FOR A SECOND DRIVER
512
FORM ELECTRICAL CONNECTIONS BETWEEN THE FIRST DIE INTERCONNECT AND THE FIRST DRIVER
514
FORM ELECTRICAL CONNECTIONS BETWEEN THE SECOND DIE INTERCONNECT AND THE SECOND DRIVER
END

FIG. 5

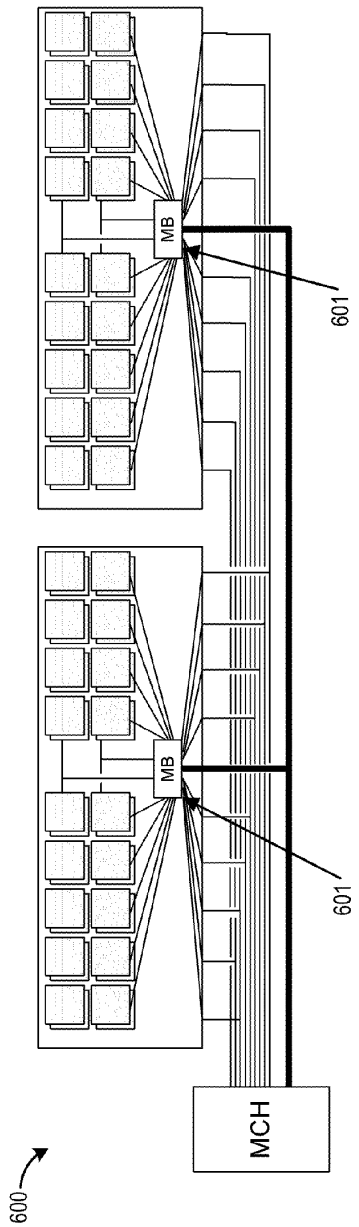


FIG. 6A

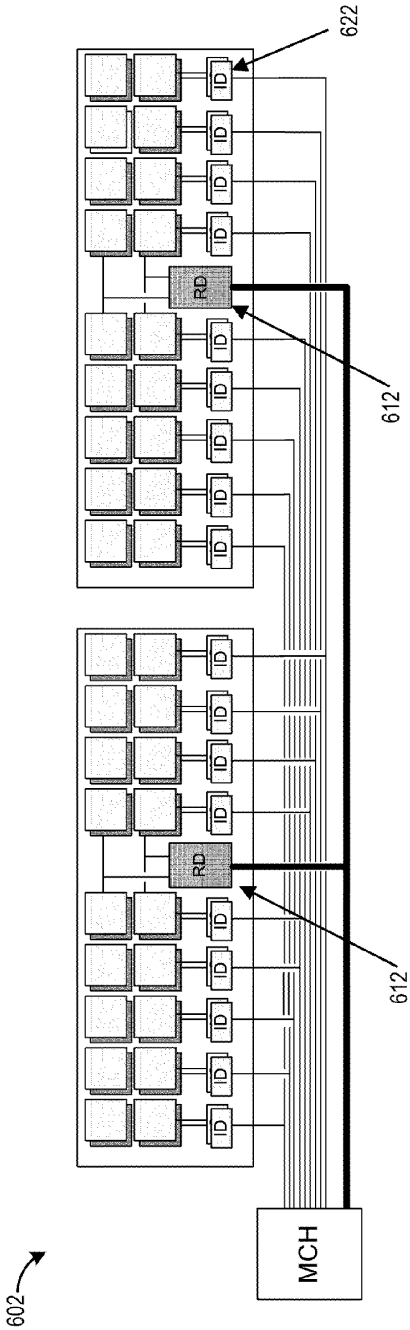


FIG. 6B

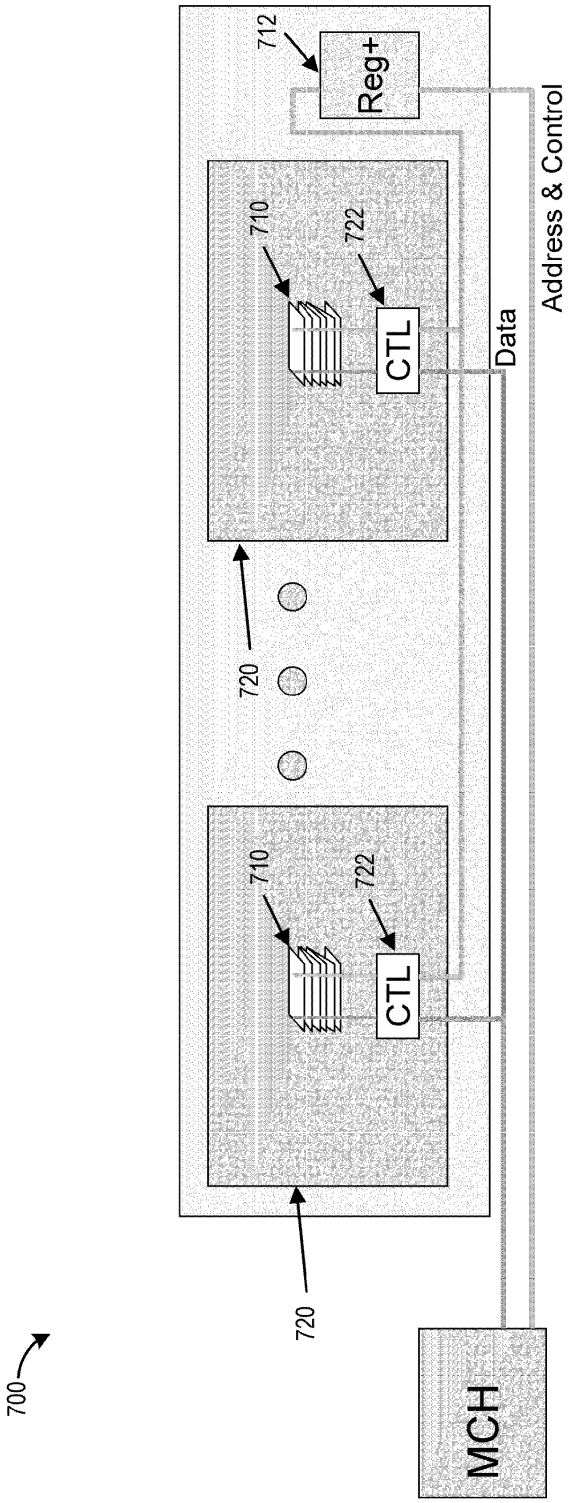

712
Reg+
710
722
CTL
Data
Address & Control
720
710
722
CTL
720
700
MCH

FIG. 7

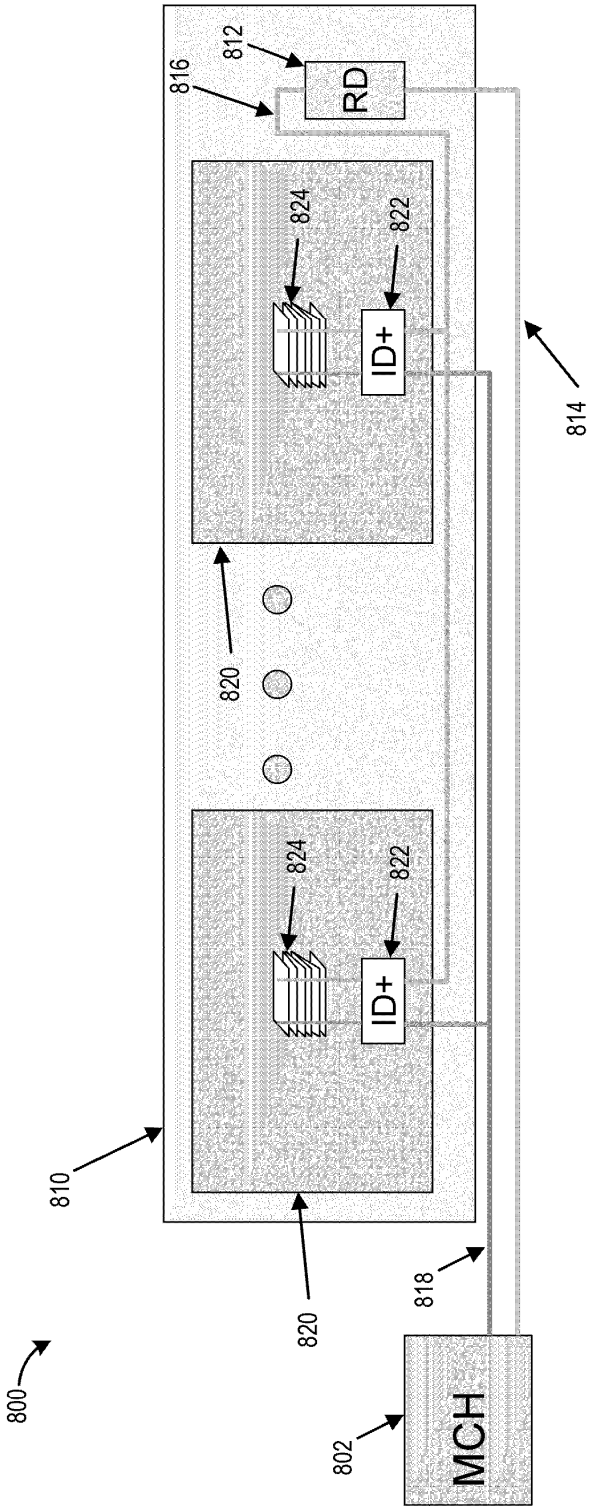


FIG. 8

METHOD AND APPARATUS FOR OPTIMIZING DRIVER LOAD IN A MEMORY PACKAGE

RELATED APPLICATION

This application claims the benefit of priority under 35 U.S.C. §119(e) of U.S. Provisional Patent Application No. 61/409,893, filed on Nov. 3, 2010, and entitled "ARCHITECTURE FOR MEMORY MODULE WITH PACKAGES OF THREE-DIMENSIONAL STACKED (3DS) MEMORY CHIPS," the disclosure of which is hereby incorporated by reference in its entirety.

BACKGROUND

1. Technical Field

The present disclosure relates to memory devices and memory modules. More specifically, the present disclosure relates to systems and methods for reducing the load of drivers of memory packages included on the memory modules.

2. Description of the Related Art

Memory modules may include a number of memory packages. Each memory package may itself include a number of array dies that are packaged together. Each array die may include an individual semiconductor chip that includes a number of memory cells. The memory cell may serve as the basic building block of computer storage representing a single bit of data.

FIGS. **1**A and **1**B schematically illustrate examples of existing memory package designs currently used or proposed to be used to provide the dynamic random-access memory of memory modules. FIG. **1**A schematically illustrates a memory package **100** with three array dies **110** and a control die **130**. The control die **130** is configured to respond to signals received by the memory package **100** by sending appropriate control signals to the array dies **110** and includes a driver **134** for driving data signals to each of the array dies **110** via a corresponding die interconnect **120**. Further, the control die **130** includes a driver **140** for driving command and/or address signals to each of the array dies **110** via another corresponding die interconnect **142**. For simplicity, FIG. **1**A shows only a single driver **134**, die interconnect **120**, driver **140**, and die interconnect **142**. However, additional drivers and die interconnects may be included for each bit the memory package **100** is designed to support. Thus, a 16-bit memory may include 16 pairs of drivers and die interconnects for the data signals and other similar drivers and die interconnects for the command and/or address signals. Each array die **110** also includes a chip select port **144**, with the chip select ports **144** of the array dies **110** configured to receive corresponding chip select signals to enable or select the array dies for data transfer. The array dies **110** are configured to transfer data (e.g. read or write) to or from the selected memory cells identified by the command, address, and chip select signals via the die interconnects.

In some cases, the control die **130** may include memory cells and therefore, also serve as an array die. Thus, as can be seen from FIG. **1**A, the control die **130** may also include a chip select port **144**. Alternatively, the control die **130** and the array dies **110** may be distinct elements and the control die **130** may not include any memory cells.

FIG. **1**B schematically illustrates an example of a memory package **150** that includes four array dies **160** and a control die **170** that does not include memory cells. As can be seen in FIG. **1**B, each array die **160** includes a chip select port **174**. However, because the control die **170** does not also serve as an array die, the control die **170** does not include a chip select port. As with memory package **100**, memory package **150** includes a driver **184** that drives data signals to each of the array dies **160** along a corresponding die interconnect **182**. Further, the memory package **150** includes a driver **186** for driving command and/or address signals to each of the array dies **160** via another die interconnect **188**.

Generally, a load exists on each of the drivers **134**, **140**, **184**, and **186** by virtue of the drivers being in electrical communication with the corresponding die interconnects and the corresponding circuitry of the array dies. Thus, to drive a signal along a die interconnect, a driver typically must be large enough to overcome the load on the driver. However, generally a larger driver not only consumes more space on the control die, but also consumes more power.

SUMMARY

In certain embodiments, an apparatus is provided that comprises a plurality of array dies having data ports. The apparatus further comprises at least a first die interconnect and a second die interconnect. The first die interconnect is in electrical communication with at least one data port of a first array die of the plurality of array dies and at least one data port of a second array die of the plurality of array dies and not in electrical communication with the data ports of at least a third array die of the plurality of array dies. The second die interconnect is in electrical communication with at least one data port of the third array die and not in electrical communication with the data ports of the first array die and the data ports of the second array die. In addition, the apparatus comprises a control die. The control die comprises at least a first data conduit configured to transmit a data signal to the first die interconnect and to not transmit the data signal to the second die interconnect, and at least a second data conduit configured to transmit the data signal to the second die interconnect and to not transmit the data signal to the first die interconnect.

In certain embodiments, an apparatus is provided that comprises a plurality of array dies having ports. The apparatus further comprises at least a first die interconnect and a second die interconnect. The first die interconnect is in electrical communication with at least one port of a first array die of the plurality of array dies and at least one port of a second array die of the plurality of array dies and not in electrical communication with the ports of at least a third array die of the plurality of array dies. The second die interconnect is in electrical communication with at least one port of the third array die and not in electrical communication with the ports of the first array die and the ports of the second array die. In addition, the apparatus comprises a control die. The control die comprises at least a first conduit configured to transmit a signal to the first die interconnect and to not transmit the signal to the second die interconnect, and at least a second conduit configured to transmit the signal to the second die interconnect and to not transmit the signal to the first die interconnect. Moreover, a first load on the first conduit comprises a load of the first die interconnect, a load of the first array die, and a load of the second array die. In addition, a second load on the second conduit comprises a load of the second die interconnect and a load of the third array die.

In certain embodiments, a method is provided for optimizing load in a memory package. The memory package comprises a plurality of array dies, at least a first die interconnect and a second die interconnect, and a control die. The control die comprises at least a first driver and a second driver, the first driver configured to drive a signal along the first die interconnect, and the second driver configured to drive the signal

along the second die interconnect. The method comprises selecting a first subset of array dies of the plurality of array dies and a second subset of array dies of the plurality of array dies. The first subset of array dies and the second subset of array dies are exclusive of one another and are selected to balance a load on the first driver and on the second driver based at least in part on array die loads of array dies of the plurality of array dies and at least in part on die interconnect segment loads of segments of at least the first die interconnect and the second die interconnect. The method further comprises forming electrical connections between the first die interconnect and the first subset of array dies. In addition, the method comprises forming electrical connections between the second die interconnect and the second subset of array dies.

In certain embodiments, an apparatus is provided that comprises a register device configured to receive command/address signals from a memory control hub and to generate data path control signals. The apparatus further comprises a plurality of DRAM packages. Each of the DRAM packages comprises a control die. The control die comprises a plurality of command/address buffers and a data path control circuit configured to control command/address time slots and data bus time slots. The control die is configured to receive data signals from the memory control hub, the data path control signals from the register device, and command/address signals from the register device. Further, each of the DRAM packages comprises a plurality of DDR DRAM dies operatively coupled to the control die to receive the data signals from the control die. Moreover, the memory module is selectively configurable into at least two operational modes. The two operational modes comprise a first operational mode and a second operational mode. In the first operational mode the register device generates the data path control signals, and the control die uses the data path control signals to operate the data path control circuit. In the second operational mode, the control die operates the data path control circuit to provide the command/address signals to the plurality of DDR DRAM dies without decoding the command/address signals.

BRIEF DESCRIPTION OF THE DRAWINGS

Throughout the drawings, reference numbers are re-used to indicate correspondence between referenced elements. The drawings are provided to illustrate certain example embodiments of the inventive subject matter described herein and not to limit the scope thereof.

FIGS. **1**A and **1**B schematically illustrate examples of existing memory package designs.

FIG. **2** schematically illustrates an example embodiment of a memory package in accordance with the present disclosure.

FIG. **3** schematically illustrates another example embodiment of a memory package in accordance with the present disclosure.

FIG. **4** schematically illustrates an example embodiment of a driver structure of a control die in accordance with the present disclosure.

FIG. **5** presents a flowchart for an example embodiment of a load optimization process.

FIGS. **6**A and **6**B schematically illustrate an example of a Load Reduction Dual In-line Memory Module (LRDIMM) and a HyperCloud™ Dual In-line Memory Module (HCDIMM) architecture respectively.

FIG. **7** schematically illustrates an example of a Three-Dimensional Structure Dual In-line Memory Module (3DS-DIMM) architecture.

FIG. **8** schematically illustrates an example embodiment of a memory module architecture in accordance with the present disclosure.

DETAILED DESCRIPTION OF SPECIFIC EMBODIMENTS

In addition to the below, the following U.S. patents are incorporated in their entirety by reference herein: U.S. Pat. Nos. 7,289,386, 7,286,436, 7,442,050, 7,375,970, 7,254,036, 7,532,537, 7,636,274, 7,630,202, 7,619,893, 7,619,912, 7,811,097. Further, the following U.S. patent applications are incorporated in their entirety by reference herein: U.S. patent application Ser. Nos. 12/422,912, 12/422,853, 12/577,682, 12/629,827, 12/606,136, 12/874,900, 12/422,925, 12/504,131, 12/761,179, and 12/815,339.

Certain embodiments of the present disclosure reduce the size of drivers that are configured to drive a signal, such as a data signal, along a die interconnect to one or more array dies. Further, certain embodiments of the present disclosure reduce the power consumption of the drivers.

In certain embodiments, reducing one or both of driver size and driver power consumption may be accomplished by increasing the number of die interconnects and reducing the number of array dies that are in electrical communication with each die interconnect. For example, instead of one die interconnect in electrical communication with four array dies, there may be two die interconnects, each in electrical communication with a different pair of the four array dies.

In certain embodiments, determining the number of die interconnects and the number of array dies in electrical communication with each die interconnect is based, at least in part, on a load of each array die and a load of the die interconnect that is in electrical communication with one or more of the array dies.

In some embodiments, the load contribution from a die interconnect may be negligible compared to the load contribution from the array dies. In such embodiments, determining the number of die interconnects and the number of array dies in electrical communication with each die interconnect may be based, at least in part, on a load of each array die without considering the load of the die interconnect. However, as the physical size of a memory package shrinks, the load of a die interconnect becomes a non-negligible value relative to the load of the array dies. Thus, as memory packages become physically smaller, it becomes more important to consider the load of the die interconnect in determining the number of die interconnects and the number of array dies in electrical communication with each die interconnect. Advantageously, certain embodiments of the present disclosure account for both the loads of the array dies and the loads of the die interconnects on a conduit (e.g., driver) in determining the number of die interconnects to be used and the number of array dies in electrical communication with each die interconnect.

FIG. **2** schematically illustrates an example embodiment of a memory package **200** in accordance with the present disclosure. One example of a memory package that includes array dies and a control die is the Hybrid Memory Cube (HMC). Examples of an HMC compatible with certain embodiments described herein are described by the IDF2011 Intel Developer Forum website, http://www.intel.com/idf/index.htm, which includes presentations and papers from the IDF2011 Intel Developer Forum including the keynote address given by Justin Rattner on Sep. 15, 2011. Additional examples of an HMC compatible with certain embodiments described herein are described by the Hybrid Memory Cube Consortium website, http://www.hybridmemorycube.org.

The memory package **200** can include any type of memory package. For example, the memory package **200** may be a DRAM package, a SDRAM package, a flash memory package, or a DDR SDRAM package (e.g., DDR3, DDR4), to name a few. A memory module (not shown) may include one or more memory packages in accordance with the memory package **200**. Further, the memory package **200** may include input/output terminals (not shown) that are configured to be placed in electrical communication with circuitry of the memory module to transmit signals between the memory package **200** and a Memory Control Hub (MCH) (not shown).

The memory package **200** may include a plurality of array dies **210** (e.g., array dies **210***a*-**210***d*). The plurality of array dies **210** may be sealed within the memory package **200**. Further, circuitry of the array dies **210** may be in electrical communication with the input/output terminals of the memory package **200**. Although generally referred to as array dies herein, the array dies **210** may also be called slave dies or slave chips. Each of the array dies **210***a*-**210***d* may include circuitry (e.g., memory cells) (not shown) for storing data. Examples of array dies compatible with certain embodiments described herein are described by the existing literature regarding the Hybrid Memory Cube (e.g., as cited above). As illustrated in FIG. **2**, the plurality of array dies **210** may be arranged in a stack configuration known as a three-dimensional structure (3DS). Examples of 3DS compatible with certain embodiments described herein are described by the existing literature regarding the Hybrid Memory Cube (e.g., as cited above). However, the structure or layout of the plurality of array dies **210** is not limited as such, and other structures are possible in accordance with the present disclosure. For example, the plurality of array dies **210** may be arranged in a planar structure, or in a structure that combines 3DS with a planar structure. Moreover, while the memory package **200** is illustrated as including four array dies, the memory package **200** is not limited as such and may include any number of array dies. For example, as illustrated in FIG. **3**, the memory package **200** may include eight array dies. As further examples, the memory package **200** may include three or sixteen array dies.

Each of the array dies **210** may include one or more data ports (not shown). The data ports enable electrical communication and data transfer between the corresponding memory circuitry of the array dies **210** and a communication pathway (e.g., a die interconnect).

In the example schematically illustrated in FIG. **2**, the memory package **200** includes a plurality of die interconnects **220** (e.g., die interconnects **220***a*, **220***b*). For example, certain versions of HMC have been reported to have 512 data ports per die, with the corresponding bits of each die all connected to a single die interconnect (e.g., TSV). Examples of die interconnects include, but are not limited to, through-silicon vias (TSV), conducting rods, wire bonds, and pins. (See e.g., U.S. Pat. Nos. 7,633,165 and 7,683,459.) Each of these die interconnects **220** may be coupled to, or in electrical communication with at least one data port of at least one of the array dies **210**. In certain embodiments, at least one of the die interconnects **220** is in electrical communication with at least one data port from each of at least two array dies **210** without being in electrical communication with a data port from at least one array die **210**, which may be in electrical communication with a different die interconnect **220**.

For example, die interconnect **220***a* may be in electrical communication with a data port from array die **210***a* and a data port from array die **210***b* (as illustrated by the darkened circles in FIG. **2**) and not in electrical communication with any data ports from array die **210***c* or any data ports from array die **210***d*. The data ports of array dies **210***a* and **210***b* in electrical communication with the die interconnect **220***a* can be corresponding to the same data bit (e.g., D**0**). Other die interconnects (not shown) can be in electrical communication with other data ports corresponding to other data bits (e.g., D**1**, D**2**, . . .) of array dies **210***a* and **210***b*. These other die interconnects can be electrically isolated from the corresponding data ports of array dies **210***c* and **210***d*.

However, continuing this example, the die interconnect **220***b* may be in electrical communication with a data port from array die **210***c* and a data port from array die **210***d* (as illustrated by the darkened circles in FIG. **2**) (e.g., corresponding to the same data bit, e.g., D**0**) without being in electrical communication with any data ports from array die **210***a* and array die **210***b*. Other die interconnects (not shown) can be in electrical communication with other data ports corresponding to other data bits (e.g., D**1**, D**2**, . . .) of array dies **210***c* and **210***d*. Despite not being in electrical communication with any data ports from array die **210***a* and **210***b*, in some implementations, the die interconnect **220***b* may pass through the array dies **210***a* and **210***b* (as illustrated by the unfilled circles) e.g., through through-holes or vias of array dies **210***a* and **210***b*. For some implementations, each of the array dies **210** may be in electrical communication with corresponding die interconnects **220**, without any of the die interconnects **220** being in electrical communication with all of the array dies **210**. Where existing systems may utilize a single die interconnect to be in electrical communication with the corresponding data ports of each array die (e.g., the data ports corresponding to the same bit), certain embodiments described herein utilize multiple die interconnects to provide electrical communication to the corresponding data ports of the array dies (e.g., the data ports corresponding to the same data bit) with none of the multiple die interconnects in electrical communication with data ports of all the array dies.

In addition to the plurality of array dies **210**, the memory package **200** includes a control die **230**, which may also be called a master die. Examples of master dies compatible with certain embodiments described herein are described by the existing literature regarding the Hybrid Memory Cube (e.g., as cited above). In some embodiments, the control die **230** may be one of the array dies **210**. Alternatively, the control die **230** may be a modified version of one of the array dies **210**. Thus, in some implementations, memory package **200** may include four dies or chips instead of the five illustrated in FIG. **2**. Further, in some embodiments, the control die **230** may comprise a logic layer (e.g., the logic layer of an HMC).

The control die **230** may include a number of data conduits **232**, which includes data conduits **232***a* and **232***b*. Each of these data conduits **232** may be configured to transmit a data signal to a single die interconnect **220**. For example, the data conduit **232***a* may be configured to transmit a data signal to the die interconnect **220***a* without transmitting the data signal to the die interconnect **220***b*. Conversely, the data conduit **232***b* may be configured to transmit the data signal to the die interconnect **220***b* without transmitting the data signal to the die interconnect **220***a* (e.g., data conduit **232***b* is electrically isolated from die interconnect **220***a* and data conduit **232***a* is electrically isolated from die interconnect **220***b*).

In some embodiments, the data conduits **232** may also include one or more drivers **234** as schematically illustrated by FIG. **2**. The drivers **234** may be configured to drive the data signals along the corresponding die interconnects **220**. In some embodiments, a single data conduit **232** or driver **234** may be in electrical communication with multiple die interconnects **220**, each of which may be in electrical communication with different array dies **210**.

Each of the data conduits **232** may be configured to receive the data signal from a common source. For instance, the data conduit **232***a* and the data conduit **232***b* may each receive a substantially similar, if not identical, data signal from the same signal source (e.g. the data signal corresponding to the same data bit). The source of the data signal may include a data path, a driver, a latch, a pin, or any other construct that may provide a data signal to a data conduit.

The data conduits **232** may each be subject to a load. Although not limited as such, this load may be measured as a capacitive load, such as a parasitic capacitance. The load on each of the data conduits **232** may include at least a load of the die interconnect **220** to which the data conduit **232** is coupled or connected as well as a load of each array die **210** with which the die interconnect **220** is in electrical communication via a data port of the die interconnect **220**. Thus, for example, the load of the data conduit **232***a* may include loads of the die interconnect **220***a*, the array die **210***a*, and the array die **210***b*. Similarly, the load of the data conduit **232***b* may include loads of the die interconnect **220***b*, the array die **210***c*, and the array die **210***d*.

Generally, a load that would be on a data conduit that was in electrical communication with a die interconnect that was in electrical communication with at least one data port of each of the array dies **210** can be considered the maximum load for a data conduit. This maximum load is the load of a data conduit of a memory package that does not implement the teachings of this disclosure, but is in accordance with conventional configurations.

In some implementations, the difference between the load of the data conduit **232***a* and the load of the data conduit **232***b* is less than the maximum load for a data conduit as described above. Thus, in some cases, there may exist a degree of balance or equalization between the loads of the data conduits **232***a*, **232***b*. In some implementations, the difference between the load of the data conduit **232***a* and the load of the data conduit **232***b* is zero or substantially zero. In some embodiments, the length of each die interconnect **220**, and the number of array dies **210** in electrical communication with each die interconnect **220** may be selected to maintain the difference between the load of the data conduit **232***a* and the load of the data conduit **232***b* to be at or below a threshold load difference. For example, suppose that the load of each array die **210** is 1, the load of each segment of the die interconnects **220** is 0.25, and that the threshold load difference is 0.5. Using the configuration schematically illustrated in FIG. **2**, the load on the data conduit **232***a* in this example is 2.5 and the load on the data conduit **232***b* in this example is 3. Thus, in this example, the difference between the load of the data conduit **232***a* and the load of the data conduit **232***b* is at the threshold load difference value of 0.5. However, an alternative configuration that places the die interconnect **220***a* in electrical communication with only the array die **210***a*, and the die interconnect **220***b* in electrical communication with the array dies **210***b*-**210***d* would not satisfy the threshold load difference value of 0.5 of the above example. In the alternative configuration, the load on the data conduit **232***a* would be 1.25 and the load on the data conduit **232***b* would be 4. Thus, in the alternative configuration, the difference between the load of the data conduit **232***a* and the load of the data conduit **232***b* is 2.75, which is above the threshold load difference value of 0.5.

For certain embodiments, the load of each data conduit is less than the maximum load as described above. Thus, in some cases, the load of the data conduit **232***a* is less than the maximum load and the load of the data conduit **232***b* is less than the maximum load. Further, in many implementations, the combined load of the data conduit **232***a* and the data conduit **232***b* is less than the maximum load of a single data conduit. In other words, it is possible to design the data conduit **232***a* and the data conduit **232***b* to reduce the overall load compared to a single data conduit that is in electrical communication with a die interconnect that is in electrical communication with at least one data port of each of the array dies **210**. By reducing the overall load compared to the single data conduit, it is possible in many cases to reduce power consumption. Further, it is possible in many cases to maintain signal quality (e.g. maintain signal amplitude, maintain low signal distortion, etc.) while reducing power consumption. Advantageously, in a number of embodiments, by using multiple data conduits instead of a single data conduit, the speed of the memory package **200** can be increased. In some cases, this speed increase can include a reduced latency in accessing array dies **210** and/or operating the memory package **200** at a higher clock frequency.

Each of the die interconnects **220** may include any type of conducting material. For example, the die interconnects **220** may include copper, gold, or a conductive alloy, such as a copper/silver alloy. Further, the die interconnects **220** may include any type of structure for enabling electrical communication between the data conduits **232** and the data ports of the array dies **210**. For example, the die interconnects **220** may include a wire, a conducting rod, or a conducting trace, to name a few. Moreover, the die interconnects **220** may use vias, or through-silicon vias (TSVs) to couple with or to electrically communicate with an array die. For instance, die interconnect **220***a* may be connected with data ports of the array dies **210***a* and **210***b* using vias (illustrated by the filled or darkened circles). Examples of TSVs which may be used with the present disclosure are described further in U.S. Pat. Nos. 7,633,165 and 7,683,459.

In addition, the die interconnects **220** may use via holes to pass through an array die that is not configured to be in electrical communication with the die interconnect. For instance, die interconnect **220***b* may pass through array dies **210***a* and **210***b* using TSVs that do not enable electrical communication between the die interconnect **220***b* and data ports of the array dies **210***a* and **210***b* (illustrated by the unfilled circles). In this way, the array dies **210***a*, **210***b* are not responsive to the data signal beings transmitted by the die interconnect **220***b*. However, the die interconnect **220***b* may be connected with at least one data port of each of the array dies **210***c* and **210***d* using a via (illustrated by the filled or darkened circles). In cases where the die interconnect passes through an array die that is not configured to be in electrical communication with the die interconnect, the TSV may include an insulator or an air gap between the die interconnect and the array die circuitry that is large enough to prevent electrical communication between the die interconnect and the array die circuitry. In certain embodiments, the TSV for array dies that are configured to be in electrical communication with the die interconnect and for array dies that are not configured to be in electrical communication with the die interconnect may be configured the same. However, in such cases, electrical connections leading from the TSV of the array dies that are not configured to be in electrical communication with the die interconnect may not exist or may be stubs. These stubs are not configured to provide electrical communication with the memory cells of the array die.

Although FIG. **2** illustrates a single pair of data conduits **232** corresponding to a single pair of die interconnects **220**, this is only to simplify the drawing figure. The memory package **200** may generally include as many additional data conduits and corresponding die interconnects as the number of

bits the memory package **200** is designed or configured to support per memory address. Thus, if, for example, the memory package **200** is configured to be a 16-bit memory, the memory package **200** may include 16 pairs of data conduits **232** and 16 pairs of corresponding die interconnects **220**. Similarly, if the memory package **200** is configured as a 32- or 64-bit memory, the memory package **200** may include 32 or 64 pairs of data conduits **232** and 32 or 64 pairs of corresponding die interconnects **220**. Generally, the data conduits and die interconnects for each bit are configured identically. Thus, for the memory package **200**, each data conduit is configured to be in electrical communication with a die interconnect that is configured to be in electrical communication with a pair of the array dies **210**. However, it is possible that, in some embodiments, the data conduits and die interconnects of different bits may be configured differently.

In certain embodiments, the same die interconnects and the same corresponding data conduits may be used to transfer data both to and from the array dies **210**. In such embodiments, the die interconnects may be bi-directional. In alternative embodiments, separate die interconnects and corresponding data conduits may be used to transfer data to the array dies **210** and data from the array dies **210** to the control die **230**. Thus, such embodiments may include double the number of die interconnects and data conduits as embodiments that use the same die interconnect to transfer data to and from an array die.

In some embodiments, the control die **230** may include additional command/address conduits **240** and die interconnects **242**, which may be in electrical communication with at least one port of each of the array dies **210**. For simplicity, FIG. **2** shows only a single such conduit **240** and die interconnect **242**. The command/address conduits **240** are configured to provide corresponding signals to the die interconnects **242**. These signals may be command signals, address signals, or may serve as both command and address signals (e.g., may include a memory cell address and a write command or a read command) either simultaneously or based on a determining criterion, such as the edge of a clock signal. The command die **230** may include a command/address conduit **240** and corresponding die interconnect **242** for each bit of the command/address signals that the memory package **200** is configured to support. The number of bits of the command/address signals may be the same or may be different from the number of data bits of the memory package **200**.

In addition, in certain embodiments, the control die **230** may include a plurality of chip select conduits **250** (e.g., chip select conduits **250***a*-**250***d* as shown in FIG. **2**). Further, the control die **230** may include corresponding die interconnects **252** (e.g., die interconnects **252***a*-**252***d*) with one die interconnect **252** in electrical communication with one chip select conduit **250** and one array die **210**. Each of the die interconnects **252** may be in electrical communication with a different array die **210**. For example, the die interconnect **252***a* may be in electrical communication with the array die **210***a* and the die interconnect **252***b* may be in electrical communication with the array die **210***b*. Each of the chip select conduits **250** may be configured to provide a chip select signal to a corresponding array die **210** via a corresponding die interconnect **252**.

In some embodiments, the control die **230** may include additional drivers that are configured to drive the chip select signals along the die interconnects **252**. Alternatively, the chip select signals may be driven by drivers that are external to the control die **230**. For example, a register (not shown) that is part of a memory module that includes the memory package **200** may determine the chip select signals and drive the chip select signals to the array dies **210**. As a second example, the chip select signals may be provided by an MCH. In some embodiments, the control die **230** may determine the array die **210** to select based on, for example, an address signal. In such embodiments, the control die may generate the chip select signals.

FIG. **3** schematically illustrates another example embodiment of a memory package **300** in accordance with the present disclosure. In certain embodiments, some or all of the embodiments described above with respect to the memory package **200** may be applicable to the memory package **300**. However, for ease of illustration and to simplify discussion, certain elements are omitted in FIG. **3**, such as the chip select conduits. Nevertheless, it should be understood that the memory package **300** can include the same or similar elements as described above with respect to the memory package **200**, including, for example, the chip select conduits.

The memory package **300** may include a plurality of array dies **310** (e.g., array dies **310***a*-**310***h*). In the implementation illustrated in FIG. **3**, the memory package includes eight array dies. However, as stated earlier, the memory package **300** may include more or fewer array dies. Each of the array dies **310** may include one or more ports (not shown) that enable electrical communication between the circuitry of the array dies **310** and one or more die interconnects **320**. Each of these die interconnects **320** may be coupled to, or in electrical communication with at least one port of at least one of the array dies **310**. As with the memory package **200**, in certain embodiments, at least one of the die interconnects **320** is in electrical communication with at least one port from each of at least two array dies **310** without being in electrical communication with a port from at least one array die **310**, which may be in electrical communication with a different die interconnect **320**.

In addition to the plurality of array dies **310**, the memory package **300** includes a control die **330**. In some embodiments, the control die **330** may include a number of conduits **332** (e.g., conduits **332***a*-**332***f*). Each of these conduits **332** may be configured to transmit a signal to a single die interconnect **320**.

Further, implementations of the conduits **332** may include one or more drivers **334** (e.g., drivers **334***a*-**334***f*). Each of the drivers **334** may be configured to drive a signal along a corresponding die interconnect **320**. For instance, the driver **334***a* of the conduit **332***a* may be configured to drive a signal along the die interconnect **320***a* to one or more of the array dies **310***a* and **310***b*. As a second example, the driver **334***b* of the conduit **332***b* may be configured to drive a signal along the die interconnect **320***b* to one or more of the array dies **310***c* and **310***d*. Although the die interconnect **320***b* may pass through the array dies **310***a* and **310***b*, because the die interconnect **320***b*, in the example illustrated in FIG. **3**, is not configured to be in electrical communication with the array dies **310***a* and **310***b*, the driver **334***b* does not drive the signal to the array dies **310***a* and **310***b*.

In some embodiments, the signal can be a data signal, a command or address signal, a chip select signal, a supply voltage signal, or a ground voltage signal, to name a few. Further, as the signal is not limited to a data signal, in some embodiments, the conduits **332** may include conduits configured to provide signals other than data signals to the die interconnects **320**. For example, the conduits may include conduits configured to provide a command or address signal, a chip select signal, a supply voltage signal, or a ground voltage signal to one or more die interconnects. Consequently, in some embodiments, the drivers **334** may be configured to drive signals other than data signals.

Generally, the signal that each of the drivers **334** drive to the corresponding die interconnects **320** is from a common source. Thus, each of the drivers **334**, in certain embodiments, is driving the same signal to each corresponding die interconnect **320**.

The size of the drivers **334** is generally related to the load on the driver **334**. In certain embodiments, the load on each driver **334** corresponds to the load of the respective conduit **332**. Although not limited as such, the load may be measured as a capacitive load, such as a parasitic capacitance.

The load on each of the conduits **332** may include at least the loads of the die interconnect **320** with which the conduit **332** is coupled or connected as well as the loads of each array die **310** with which the die interconnect **320** is in electrical communication via a port of the die interconnect **320**. Thus, for example, the loads of the conduit **332***a* may include the loads of the die interconnect **320***a*, the array die **310***a*, and the array die **310***b*. Similarly, the loads of the conduit **332***b* may include the loads of the die interconnect **320***b*, the array die **310***c*, and the array die **310***d*. The loads of both conduits **332***a* and **332***b* include a load of two array dies **310** because the corresponding die interconnects **320** are each configured to be in electrical communication with two array dies **310**. On the other hand, the load of the conduit **332***c*, which may include the loads of the die interconnect **320***c* and the array die **310***e*, includes a load of one array die **310***e* because the corresponding die interconnect **320***c* is configured to be in electrical communication with only one array die **310**.

As previously described with respect to FIG. **2**, a load that would be on a conduit that was in electrical communication with a die interconnect that was in electrical communication with at least one port of each of the array dies **310** can be considered the maximum load for a conduit. This maximum load is the load of a conduit of a memory package that does not implement the teachings of this disclosure, but is used in accordance with conventional configurations.

In some implementations, the difference between the loads of any pair of the conduits **332** is less than the maximum load for a conduit as described above. For instance, the difference between the load of the conduit **332***a* and the load of the conduit **332***b* is less than the maximum load. As a second example, the difference between the load of the conduit **334***f* and any one of the conduits **334***a*-**334***e* is less than the maximum load. Thus, in some cases, the load on each of the conduits **332** may be, at least partially balanced or equalized to reduce or minimize the difference between the load of any pair of the conduits **332**. In some implementations, the difference between the load of a pair of the conduits **332** is zero or substantially zero. In some embodiments, the length of each die interconnect **320**, and the number of array dies **310** in electrical communication with each die interconnect **320** may be selected to maintain the difference between the load of any pair of the conduits **332** to be at or below a threshold load difference.

For certain embodiments, the load of each conduit is less than the maximum load as described above. Thus, for example, the load of the conduit **332***a*, the load of the conduit **332***b*, and the load of the conduit **332***c* are each less than the maximum load defined above.

In certain embodiments, the load associated with each of the array dies **310** may be substantially equivalent. For example, the load of the array die **310***a* may be substantially equal to the load of the array die **310***h*. Thus, the load contribution from the array dies **310** for a specific conduit **332** may be measured as a multiple of the array dies that are in electrical communication with a die interconnect **320** corresponding to a specific conduit **332**. For example, assuming that the load of each of the array dies **310** is then the contribution of the load from the array dies **310** to the conduit **332***a* would be 2L because the die interconnect **320***a* corresponding to the conduit **332***a* is in electrical communication with two array dies **310**, array dies **310***a* and **310***b*. In alternative embodiments, the load of each of the array dies **310***a* may differ. For example, the load of array die **310***a* may be L and the load of the array die **310***b* may be 1.25L.

Similar to the array dies **310**, in some embodiments, a die interconnect **320** can be considered to comprise a plurality of segments, with each segment of a die interconnect **320** contributing a substantially equivalent load to the load of the die interconnect **320**. In this case, the segment of the die interconnect **320** may refer to the portion of the die interconnect between two successive or adjacent dies (array die-to-array die, master die-to-array die, or both) along the die interconnect **320**. Further, the segment may be defined as a portion of the die interconnect **320** between the dies exclusive of a portion of the dies. For example, one segment of the die interconnect **320***a* may extend from the top of the array die **310***a* to the bottom of the array die **310***b*. Alternatively, the segment may be defined to include at least a portion of at least one of the array dies **310**. For example, one segment of the die interconnect **320***a* may extend from the center of array die **310***a* to the center of array die **310***b*. As a second example, one segment of the die interconnect **320***a* may extend from the top of the control die **330** to the top of the array die **310***a*, and therefore may include a portion of the die interconnect **320***a* extending from the bottom of the array die **310***a* to the top of the array die **310***a*. In some implementations, the segments are substantially equal in length to each other. Moreover, the load contribution of each segment may be substantially equal to each other. Alternatively, the segments may be unequal in length and/or may each contribute a different load to the total load of a die interconnect **320**. Further, in some cases, the load contribution of a segment of the die interconnect **320** that is in electrical communication with a port of an array die **310** may differ from the load contribution of a segment of the die interconnect **320** that is not in electrical communication with a port of an array die **310**.

In some cases, the load contribution of each segment of the die interconnect **320** may be measured as a fraction of the load contribution from an array die **310**. For example, the load of one segment of the die interconnect **320***a* may be equivalent to one quarter of the load of an array die **310**. Thus, for example, the load of the conduit **332***a* may be 2.5L assuming a load contribution of L per array die **310** (two in this case) in electrical communication with the die interconnect **320***a* and a load contribution of 0.25L per segment (two in this case) of the die interconnect **320***a*. As a second example, the load of the conduit **332***f* may be 3L assuming the same load values as the previous example and a load contribution from one array die **310***h* and eight die interconnect **320***f* segments. Table 1 specifies the capacitive load values for each conduit **332** assuming, as with the previous two examples, that the load of

each of the array dies **310** is L, and that the load of each segment of the die interconnects **320** is 0.25L. Table 1 also specifies the deviation in load from the conduits having the highest load value, which in this example are conduits **332***b* and **332***f*.

TABLE 1

Conduit	Number of Array Dies	Number of Die Interconnect Segments	Capacitive Load	Deviation from Maximum Load
332a	2	2	2.5 L	0.5
332b	2	4	3 L	0
332c	1	5	2.25 L	0.75
332d	1	6	2.5 L	0.5
332e	1	7	2.75 L	0.25
332f	1	8	3 L	0

As can be seen from Table 1, the maximum load of any conduit **332**, using the example values previously described, is 3L, or rather three times the load of a single array die **310**. Assuming the same example values, the load of a conduit in electrical communication with a die interconnect that itself is in electrical communication with each array die **310** would be 10L. Thus, certain embodiments of the present disclosure enable a reduction in the load of the conduits **332**. Consequently, in some embodiments, the drivers **334** may each be smaller than a single driver that is configured to drive a signal from a conduit along a single die interconnect that is in electrical communication with a port from each of the array dies **310**. Moreover, the drivers **334** may include smaller transistor sizes than a single driver that is configured to drive a signal to each of the array dies **110**.

As can be seen from Table 1, conduits **332***b* and **332***f* have the largest capacitive load of the group of conduits, which is 3L. Conduit **332** has the smallest capacitive load of the group of conduits, which is 2.25L and which is a deviation of only 0.75 from the maximum load. Thus, in some embodiments, each of the drivers **334** may be substantially similar in size. In certain implementations, the drivers **334** may vary in size based on the total capacitive load on each conduit **332**. Thus, the driver **334***f* may be larger than the driver **334***e*. Alternatively, each driver **334** may be substantially equal and may be configured based on the drivers **334** with the largest load, which are drivers **334***b* and **334***f* in the example illustrated in FIG. **3** and Table 1.

In some embodiments, the capacitive load of each conduit **332** or driver **334** can be calculated using formula (1).

$$CL = AD + \frac{S}{M} \quad (1)$$

In formula (1), CL represents the capacitive load of a conduit **332** or driver **334**, AD represents the number of array dies **310** in electrical communication with a die interconnect **320** that is in electrical communication with the conduit **332** and/or driver **334**, S represents the number of die interconnect segments of the die interconnect **320**, and M represents the ratio of the load of an array die **310** to the load of a segment of a die interconnect **320**. Thus, using formula (1) and the example values described above, the load of the driver **332***a*, for example, can be calculated with the following values: AD=2, for two array dies **310**; S=2, for the two segments of the die interconnect **320***a*; and M=4, for the ratio of the load of an array die **310**, L, to the load of a segment of the die interconnect **320***a*, 0.25L. Therefore, as can be seen from Table 1, the capacitive load of the driver **332***a* is 2.5L where L is the load of a single array die.

In some cases, the load on each conduit **332** and/or driver **334** may be evenly balanced. In other words, the load on each conduit **332** and/or driver **334** may be substantially the same. To achieve a balanced load, each of the conduits **332** may be in electrical communication with a combination of array dies **310** and die interconnect **320** segments that results in a load that is substantially equivalent to the loads of the other conduits **332**. In certain embodiments, the load of the conduits **332** is balanced despite each conduit **332** being in electrical communication with a different subset of the array dies **310**. However, in alternative embodiments, the load of each conduit **332** and/or driver **334** may differ. This difference in the load of the conduits **332** and/or drivers **334** may be a design decision (e.g. to maintain a specific number of drivers). Alternatively, or in addition, the load difference between conduits **332** and/or drivers **334** may occur because the loads of the array dies **310** and the die interconnect segments **320** do not allow for perfect or substantially even load balancing.

As previously stated, in some embodiments the loads of the conduits **332** and/or drivers **334** may be balanced to minimize the difference between any pair of conduits **332** and/or drivers **334**. For example, as illustrated in FIG. **3**, a driver in electrical communication with a longer die interconnect, which may consequently include more die interconnect segments, is likely to obtain a larger load contribution from the die interconnect than a driver in electrical communication with a shorter die interconnect (e.g., compare driver **334***f* to **334***a*). Thus, the driver in electrical communication with the longer die interconnect may be in electrical communication with fewer array dies than the driver in electrical communication with the shorter die interconnect (e.g. compare driver **334***f* to **334***a*). The selection of die interconnect length and the selection of array dies to place in electrical communication with the die interconnects may therefore be dependent on the load of each segment of the die interconnects **320** and the loads of the array dies **310**.

Alternatively, or in addition, the loads of the conduits **332** and/or drivers **334** may be balanced to reduce the maximum load of each conduit **332** and/or driver **334**. Further, in some embodiments, the maximum load of each conduit **332** and/or driver **334** may be reduced to maintain the load of each conduit **332** and/or driver **334** to be at or below a threshold load difference.

Although not illustrated in FIG. **3**, each of the conduits **332** may include one or more additional drivers configured to drive a signal received from an array die via a corresponding die interconnect **320**. The additional drivers may drive the signal to a processor, a register, a latch, or any other component or device that may or may not be part of a memory module that includes the memory package **300**. In some implementations, one or more of the die interconnects **320** are configured to be bi-directional, thereby enabling a signal to be driven to the array dies **310** via the die interconnect **320**, and to enable a signal received from the array dies **310** along the same die interconnect **320** to be transmitted to a corresponding conduit **332** of the control die **330**. Alternatively, the memory package **300** comprises one or more die interconnects **320** that are configured to enable a signal to be driven to the array dies **310** without enabling a signal to be received from the array dies **310** (e.g., the die interconnects **320** may not be bi-directional). In certain embodiments where one or more of the die interconnects **320** are not bi-directional, the memory package **300** may include additional die interconnects (not shown) that are in electrical communication with

the one or more additional conduits or drivers and that are configured to enable a signal from the array dies **310** to be transmitted to the one or more additional conduits or drivers.

In addition to the drivers **334** of the control die **330**, the memory package **300** may include one or more pre-drivers **340**. A pre-driver **340**, shown schematically in FIG. **3**, may be large enough to drive a signal (e.g., a data signal) to any of the drivers **332** and subsequently, to any of the array dies **310**. Thus, using the same example values as described above in relation to Table 1, a load of the pre-driver **340** may be at least 10L, which is the load of a conduit in electrical communication with a die interconnect that is in electrical communication with a port from each of the array dies **310**. In some embodiments, the memory package **300** may include any number of additional drivers and/or latches for buffering and/or driving the signal to any of the array dies **310**.

Just as the memory package **300** may include a pre-driver **340** for providing a signal to the conduits **332**, the memory package **300** may include a post-driver (not shown) for driving an output signal from the control die **330**. This post-driver may drive the output signal to additional latches and/or drivers. In some embodiments, the post-driver may drive the signal from the memory package **300** to a bus or other electrical path that is in electrical communication with the memory package **300**.

FIG. **4** illustrates an example embodiment of a driver structure **400** of a control die (e.g. control die **230**) in accordance with the present disclosure. The driver structure **400** is configured for a single data bit. Thus, a control die for a 32-bit memory package will generally include at least 32 instances of the driver structure **400**, one per bit. In some embodiments, the control die **230** may include the driver structure **400** in place of the combination of drivers **232***a* and **232***b*. Similarly, in certain embodiments, the control die **330** may include a structure similar to the driver structure **400** in place of the drivers **334**. In such embodiments, the driver structure **400** would be modified to enable a signal (e.g., data signal) to be driven to the six conduits **332**.

The driver structure **400** can include an input/output port **402** configured to receive or send a signal, such as a data signal. Signals received at the input/output port **402** can be provided to the drivers **404***a* and **404***b*. In turn, these drivers **404***a* and **404***b* can drive the signal to conduits **406***a* and **406***b* respectively, which are in electrical communication with corresponding die interconnects. Each of the die interconnects may be in electrical communication with different array dies in accordance with certain embodiments described herein.

In some embodiments, the driver structure **400** is bi-directional. In such embodiments, operation of the drivers **404***a* and **404***b*, and **408***a* and **408***b* may be controlled or enabled by a control signal (e.g., a directional control signal). This control signal, in some cases, may correspond to one or more of a command/address signal (e.g., read/write) and a chip select signal. In some implementations, the drivers **404***a* and **404***b* may drive a signal to the array die corresponding to the chip select signal, and not to array dies that do not correspond to a chip select signal. For example, suppose the conduit **406***a* is in electrical communication with array dies one and two (not shown), and that the conduit **406***b* is in electrical communication with array dies three and four (not shown). If a chip select signal is received corresponding to array die two, then in some implementations, driver **404***a* may be configured to drive a signal along a die interconnect in electrical communication with the conduit **406***a* to array die one and two. In this example, the driver **404***b* would not drive the signal because the chip select signal does not correspond to either array die three or array die four.

In some embodiments, the drivers **404***a* and **404***b* may drive a signal to all of the array dies. For example, assume the memory package that includes the driver structure **400** also includes four array dies. If the die interconnect in electrical communication with the conduit **406***a* is in electrical communication with two of the array dies, and if the die interconnect in electrical communication with the conduit **406***b* is in electrical communication with the other two array dies, then the drivers **404***a* and **404***b* may drive a signal to all four of the array dies of this example.

Similar to the input/output port **402**, the conduits **406***a* and **406***b* can be configured to receive a signal from the die interconnects that are in electrical communication with the conduits **406***a* and **406***b*. In turn, the signal received at one of the conduits **406***a* and **406***b* can be provided to drivers **408***a* and **408***b* respectively. The drivers **408***a* and **408***b* can each be configured to drive a signal received from a respective data conduit **406***a* and **406***b* to the input/output port **402** and to another component that may be in electrical communication with the input/output port **402**, such as a MCH.

In some embodiments, one or more chip selects, as illustrated in FIG. **2**, may be used to select, determine, or enable the array die that communicates the signal to or from one or more of the conduits **406***a* and **406***b* and the drivers **408***a* and **408***b*. Similarly, in some embodiments, the chip select may select, determine, or enable the array die to receive and/or respond to the signal driven by the drivers **404***a* and **404***b* to the array dies.

It should be noted that the driver structure **400** is a non-limiting example for arranging drivers in a control die. Other driver structures are possible. For example, instead of the drivers **404***a* and **408***a* being in electrical communication with the same conduit **406***a*, each of the drivers **404***a* and **408***a* can be in electrical communication with a separate conduit and consequently a separate die interconnect. In such an example, the drivers **404***a* and **408***a* may still be in electrical communication with the same input/output port **402**, or each driver may be in electrical communication with a separate port, the driver **404***a* in electrical communication with an input port, and the driver **408***a* in communication with an output port.

FIG. **5** presents a flowchart for an example embodiment of a load optimization process **500**. In certain embodiments, the load optimization process **500** may be performed, at least in part, by one or more computing systems. Further, the process **500**, in some embodiments, may be used to optimize one or more loads in a memory package (e.g. memory package **200** or memory package **300**). Optimizing the loads in the memory package can include optimizing the load on one or more conduits and/or drivers. As previously described with respect to FIGS. **2** and **3**, the memory package can include a plurality of array dies, a plurality of die interconnects, and a control die. Furthermore, the control die can include a plurality of drivers, each of which may be configured to drive a signal along a die interconnect.

In some implementations, the process **500** can comprise selecting a first subset of array dies and a second subset of array dies from a plurality of array dies (e.g., array dies **310**), as shown in operational block **502**. Generally, the first subset of array dies and the second subset of array dies may be exclusive of one another. Thus, the first subset of array dies does not include any array dies from the second subset of array dies and vice versa. However, in some cases there may be some overlap between the first subset of array dies and the second subset of array dies. Further, in some cases, at least one of the subsets of array dies includes more than one array die. For instance, the first subset of array dies may include two

array dies, and the second subset of array dies may include one array die, which, depending on the embodiment, may or may not be included in the first subset of array dies.

Further, the first subset of array dies and the second subset of array dies may be selected to balance a load on a first driver and a load on a second driver based, at least in part, on the loads of the array dies **310** and on the loads of the die interconnect segments from a first and a second die interconnect. The first die interconnect may be in electrical communication with the first driver and the second die interconnect may be in electrical communication with the second driver. In some embodiments, the first subset of array dies and the second subset of array dies are selected to balance a load on a first conduit and a load on a second conduit. The load on each driver and/or conduit may be calculated using formula (1) as previously described above. In some embodiments, the first subset of array dies and the second subset of array dies may be selected to balance a load on a first driver and a load on a second driver based, at least in part, on the loads of the array dies **310** without considering the loads of the die interconnect segments from the first die interconnect and the second die interconnect.

The process **500** further comprises forming electrical connections between the first die interconnect and the first subset of array dies in an operational block **504**. The process **500** further comprises forming electrical connections between the second die interconnect and the second subset of array dies in an operational block **506**. Forming the electrical connections places the die interconnects in electrical communication with the respective subsets of array dies. In some embodiments, forming the electrical connections can comprise forming electrical connections between the die interconnects and at least one port from each array die of the respective subsets of array dies.

The process **500** further comprises selecting a driver size for a first driver at block **508** and selecting a driver size for a second driver at block **510**. Selecting the driver size can be based, at least in part, on the calculated load on the driver. Generally, the greater the load on the driver, the larger the driver is selected to drive a signal along, for example, a die interconnect. The size of the driver may be adjusted by the selection of the transistor size and/or number of transistors included in the driver. A larger driver often consumes more power than a smaller driver. Thus, in certain embodiments, balancing the loads on the drivers to reduce the load on each driver can reduce the power consumption of a memory package.

In some embodiments, the size of the first driver and the size of the second driver are both less than the size sufficient for a driver to drive a signal along a die interconnect to each of the array dies **310** (e.g., with less than a predetermined or threshold signal degradation). The threshold signal degradation can be based on any one or more characteristics of a signal. For example, the threshold signal degradation can be based on the amplitude of the signal, the frequency of the signal, the noise distortion included in or introduced into the signal, or the shape of the signal, to name a few.

The process **500** further comprises forming electrical connections between the first die interconnect and the first driver in an operational block **512**. Similarly, the process **500** further comprises forming electrical connections between the second die interconnect and the second driver in an operational block **514**. Forming the electrical connections places the die interconnects in electrical communication with the respective drivers. In some embodiments, forming the electrical connections can comprise forming electrical connections between the die interconnects and a data conduit that includes the respective driver.

Advantageously, certain embodiments of the present disclosure reduce the load on each conduit and on each corresponding driver. In certain embodiments, reducing the load on a driver may increase the speed of data transfer between the array dies and other components in a computer system (e.g. a MCH or a processor). Further, reducing the load on a driver may result in reduced power consumption by the driver and consequently, by the memory package. In addition, certain embodiments of the present disclosure may minimize current switching noise.

Operational Modes

A proposed three-dimensional stacking (3DS) standard for dual in-line memory modules (DIMMs) being considered by the Joint Electron Devices Engineering Council (JEDEC) addresses three major shortcomings in the current JEDEC registered DIMM (RDIMM) and JEDEC load reduced DIMM (LRDIMM) standards (an example LRDIMM structure **600** is schematically illustrated in FIG. **6**A). These shortcomings include:

a) The DIMM density limitation due to the fixed number of chip-select signals received by the DIMM from the system memory controller.

b) The performance loss due to the increased load on the data bus as the DIMM density (e.g., the number of DRAM devices and number of ranks) increases.

c) The upper bound of the DIMM density due to the physical DIMM form factor.

Further, the LRDIMM structure **600** may have timing issues due to signals passing through a single memory buffer **601**. In addition, the increased size of the data path of the LRDIMM architecture **600**, compared to the HCDIMM architecture **602**, an example of which is schematically illustrated in FIG. **6**B, may result in more latency and signal integrity issues.

FIG. **7** schematically illustrates an example memory module **700** that has previously been proposed for the 3DS-DIMM standard. The proposed 3DS Dual In-line Memory Module (3DS-DIMM) schematically illustrated in FIG. **7** attempts to address the above shortcomings using two components, a 3DS register **712** and a controller die **722**. In some implementations, the 3DS-DIMM **700** includes the 3DS register **712** (also known as an enhanced DDR3 register). The 3DS register **712** may include a DDR3 JEDEC standard register with a "rank multiplication" circuit, which increases the number of the output chip-select signals for selecting an array die **710** by decoding one or more higher-order row or column address bits with the incoming chip-select signals from the system controller. The 3DS Register **712** can also include a command/address buffer, a register operational code buffer (RC word buffer), and rank multiplication logic (e.g., supporting rank multiplication of 1-to-2, 1-to-4, and/or 1-to-8). The 3DS register addresses shortcoming (a) of the JEDEC RDIMM and LRDIMM listed above.

Another component of the proposed 3DS-DIMM is a 3DS dynamic random-access memory (DRAM) package **720**. The 3DS DRAM package **720** includes a plurality of stacked DDR DRAM chips **724** or array dies. The 3DS DRAM package **720** has a data I/O load that is equivalent to the data I/O load of a single-die DRAM package, regardless of the actual number of DRAM dies in the 3DS DRAM package **720**. The 3DS DRAM package **720** may comprise a plurality of array dies (e.g., 2, 4 or 8 DDR DRAM dies) and a controller die **722**. The controller die **722** may include data buffers, a data path timing controller, a secondary command/address buffer, a program-

mable secondary 1-to-2 rank decoder, and a data path signal (e.g., ODT, ODT, RTT) controller. Examples of such memory packages include HMC. The 3DS DRAM package **720** addresses the shortcomings (b) and (c) of the JEDEC RDIMM and LRDIMM listed above. However, there are deficiencies in the proposed 3DS-DIMM standard as described below.

The JEDEC DDR3 RDIMM standard contains two major components: a DDR3 register and a plurality of DRAM packages each comprising one or more DRAM chips or dies. The DDR3 register serves as a command/address signal buffer and as a command/control decoder. This DDR3 register holds a set of register control (RC) words, which a computer system configures to ensure the proper operation of the RDIMM. The DDR3 register contains a phase-lock-loop (PLL), which functions as a clock synchronizer for each RDIMM. The DDR3 register outputs a set of buffered command/address signals to all the DRAM packages on the RDIMM, but the data signals are directly fed to the DRAM packages from the system memory controller.

In contrast to the JEDEC DDR3 RDIMM standard, the 3DS-DIMM proposal requires the DRAM package **720** to include the controller die **722** that controls all data paths timing and operations. This arrangement of the DRAM package reduces the load on the data bus, however, it presents three significant shortcomings:

1) The command/address buffer in the 3DS DRAM package introduces clock cycle latency since it needs to provide clock synchronous operation to the DRAM dies in the package.
2) The data path control circuit (e.g., ODT, read/write data direction) in the control die becomes very complicated to support semi-synchronous (fly-by) data path control among all 3DS DRAM packages that are on a DIMM.
3) The variations in the DRAM die timing characteristics within each 3DS DRAM package would be likely to require resynchronization of the data signals during the read/write operations, which increases the read/write latency and the read-to-write and write-to-read turn around time.

The 3DS-DIMM proposal can also be compared to the HyperCloud™ (HC) DIMM architecture of Netlist, Inc. An example of the HCDIMM architecture **602** is illustrated in FIG. **6**B. Further details and embodiments of the HCDIMM architecture is disclosed in U.S. Pat. Nos. 7,289,386, 7,532,537, 7,619,912, and 7,636,274, each of which is incorporated in its entirety by reference herein. One of the main topological differences between the 3DS-DIMM and HCDIMM architectures is that while the 3DS-DIMM architecture uses a control die **722** to buffer the data signals and to decode command/address signals from the 3DS register, certain configurations of the HCDIMM architecture include a plurality of isolation devices (ID), each of which includes data buffers, but no decoding capability. Unlike the HCDIMM architecture, since the command/address signal needs to pass through the control die **722** in the 3DS DRAM package **720**, the 3DS-DIMM proposal presents the same shortcoming (b) of the controller die described above. The data path control signals are generated by the register device (RD) **612** in the HCDIMM architecture **602**, while the data path control signals are generated by the control die **722** in the 3DS DRAM. This aspect of the 3DS-DIMM architecture creates timing critical control paths in 3DS-DIMM architecture.

In certain embodiments described herein, a memory module architecture is proposed that includes a set of device components that support JEDEC 3DS operation with the benefit of RDIMM and HCDIMM architectures (see, e.g., FIG. **8**). This set of device components may comprise two components: a register device **812** (RD), which in some embodiments may be the same or similar RD component as used in HCDIMM architectures, and a plurality of array dies **824**, such as a DDR DRAM Stack Package (DDSP). The DDSP may comprise a DRAM control die that can include command/address buffers and a data path control circuit. Certain embodiments of the architecture described herein differs from the 3DS-DIMM architecture in that instead of the controller die of the 3DS-DIMM (which provides a secondary address buffer, and a second rank multiplication decoder), certain embodiments described herein use an "ID+" die as a control die **822**, which can provide both selective isolation and address pass-through. Selective isolation refers generally to a driver corresponding to an array die driving a signal to the array die in response to a corresponding chip select signal while additional drivers corresponding to additional array dies maintain a previous state (e.g. do not drive the signal). For example, assuming that the memory package **300** implements selective isolation, if a chip select signal is received that corresponds to a selection of array die **310***b*, then the driver **334***a* will drive a signal along the die interconnect **320***a* to the array die **310***b*, and the remaining drivers (e.g., drivers **334***b*-**334***f*) will maintain their state (e.g., not drive the signal). Address pass-through is described in further detail below.

The data path control circuit of certain embodiments may have at least two operational modes: mode-C (HCDIMM/RDIMM Compatible mode) and mode-3DS (3DS DIMM compatible mode). In mode-C, the array dies **824** (e.g. a DDSP) receive the data path control signals from the register device **812**, which can be configured to control a command/address time slot and a data bus time slot. In mode-3DS, the control die **822** internally generates the data path control signals to ensure the proper operation of the data path.

The control die **822** of certain embodiments enables the memory module **800** to work as either a 3DS-DIMM, a RDIMM, or a HCDIMM. However, the memory module **800** may comprise a set of optional control input pins (in addition to the package pins that are included in 3DS-DRAM packages) which in mode-C receive the data path control signals from the register device **812**.

As previously mentioned, FIG. **8** schematically illustrates an example embodiment of a memory module architecture **800** in accordance with the present disclosure. Advantageously, certain embodiments of the memory module architecture **800** address the shortcomings described above without adding complexity or latency, and without causing performance loss. The memory module architecture **800** includes a memory control hub **802** (also known as a memory controller hub, a memory control handler, or a northbridge) and a memory module **810**. As schematically illustrated in FIG. **8**, the memory control hub **802** may communicate directly with one or more components of the memory module **810**. Alternatively, the memory control hub **802** may communicate with one or more intermediary system components (not shown), which in turn communicate with one or more components of the memory module **810**. In some embodiments, the memory control hub **802** may be integrated with another component of the computer system, such as a Central Processing Unit (CPU).

Further, in some embodiments, the memory control hub **802** may communicate with one or more memory modules. Each of the memory modules may be similar or substantially similar to the memory module **810**. Alternatively, some of the memory modules may differ from memory module **810** in configuration, type, or both. For example, some of the memory modules may be capable of operating at different

frequencies, may include a different amount of storage capability, or may be configured for different purposes (e.g. graphics versus non-graphics memory). Although in some cases a system may include memory modules capable of operating at different frequencies, each memory module may be configured to operate at the same frequency when used in a specific device. In certain implementations, the memory control hub **802** may set the operating frequency for the one or more memory modules.

The memory module **810** may include a register device **812**, which is configured to receive command, address, or command and address signals from the memory control hub **802**. For the purpose of simplifying discussion, and not to limit these signals, the signals will be referred to herein as command/address signals.

In some embodiments, the register device **812** receives the command/address signals via a command/address bus **814**. The command/address bus **814**, although illustrated as a single line, may include as many lines or signal conduits as the number of bits of the command/address signal.

Further, the register device **812** may generate data path control signals, which can be provided to a control die **822**, or isolation device, of a memory package **820** via one or more data path control lines **816**. In certain embodiments, the control dies **822** can include some or all of the embodiments described above with respect to the control dies **230** and **330**. Moreover, in certain embodiments, the memory packages **820** can include some or all of the embodiments described above with respect to the memory package **200** and **300**. In general, the memory module **812** may include one or more memory packages **820**.

In certain embodiments, each control die **822**, or isolation device, may be capable of address pass-through. Address pass-through, in some cases, enables the control die **822** to provide an address signal to one or more array dies **824** without decoding the address signal. This is possible, in some implementations, because the address signal received by the control die **822** is not encoded.

Some implementations of the control dies **822** include a plurality of command/address buffers (not shown). These buffers may comprise latches. In certain embodiments, the buffers are configured to hold command/address signals to control the timing of command/address signals. In some cases, controlling the timing of the command/address signals may reduce or slow signal degradation. In some implementations, the control dies **822** include a plurality of data buffers, which may control the timing of data signals to reduce or slow signal degradation. Further, the control dies **822** may include a data path control circuit (not shown) that is configured to control command/address time slots and data bus time slots. Controlling the command/address time slots and the data bus time slots enables the control dies **822** to reduce or prevent signal collisions caused by multiple memory packages **820** sharing the data path control lines **816** and the data bus **818**. In some implementations, the data path control circuit may be separate from the control die **822**.

Each of the control dies **822** may be configured to receive data signals from the memory control hub **802** via the data bus **818**. Further, the control dies **822** may provide data signals to the memory control hub **802** via the data bus **818**. The control dies **822** may also receive data path control signals and/or command/address signals from the register device **812** via the data path control lines **816**.

As illustrated in FIG. **8**, each memory package **820** may include one or more array dies **824** operatively coupled to the control die **822**. The array dies **824** may be configured to receive data signals from the control die **822**. As with the array dies **210** and **310**, the array dies **824** may include any type of Random Access Memory (RAM) die. For example, the array dies **824** may include DDR DRAM, SDRAM, flash memory, or SRAM, to name a few. Further, if the memory module **810** is utilized for graphics, the array dies **824** may include GDDR SDRAM or any other type of graphics memory. In addition, the array dies **824** may be configured in a stack arrangement as illustrated in FIG. **8**. Alternatively, the array dies **824** may be arranged in a planar configuration or a hybrid configuration utilizing both a planar and a stack arrangement.

In some embodiments, the memory module **810** is selectively configurable into two operational modes. In the first operational mode, the register device **812** generates data path control signals provided to the control dies **822** via the data path control lines **816**. The control dies **822** may decode command/address signals included with the data path control signals generated by the register device **812**. In some implementations, the control dies **822** use the data path control signals to operate the data path control circuits of the control dies **822**.

In the second operational mode, the control dies **822** may operate the data path control circuits to provide command/address signals to the array dies **824** without decoding the command/address signals. In this mode, the control dies **822** may use address pass-through to provide received address signals to the array dies **824**.

Other operational modes may also be possible. In some embodiments, the data path control signals generated by the register device **812** may include decoded command/address signals that are decoded from command/address signals received from the memory control hub **802** via the command/address bus **814**.

In some embodiments, the register device **812** may be configured to perform rank multiplication. In addition, or alternatively, the control dies **822** may be configured to perform rank multiplication. Embodiments of rank multiplication are described in greater detail in U.S. Pat. Nos. 7,289,386 and 7,532,537, each of which are incorporated in their entirety by reference herein. In such embodiments, the register device **812** can generate additional chip select signals that are provided to the array dies **824**. For instance, if the memory control hub **802** is configured to recognize a single array die **824** per memory package **820**, but there exists four array dies **824** per memory package **820**, the memory control hub **802** may not provide the correct number of chip select signals to access a specific memory location of the plurality of array dies **824** is memory package **820**. Thus, to access the specific memory location, the register device **812** can determine the array die that includes the specific memory location to be accessed based on the command/address signals received from the memory control hub **802** and can generate the correct chip select signal to access the array die that includes the specific memory location. In certain embodiments, when the memory module **810** is operating in the second operation module as described above, the memory module **810** does not perform rank multiplication.

Terminology

Unless the context clearly requires otherwise, throughout the description and the claims, the words "comprise," "comprising," and the like are to be construed in an inclusive sense, as opposed to an exclusive or exhaustive sense; that is to say, in the sense of "including, but not limited to." The term "coupled" is used to refer to the connection between two elements, the term refers to two or more elements that may be either directly connected, or connected by way of one or more intermediate elements. Additionally, the words "herein,"

"above," "below," and words of similar import, when used in this application, shall refer to this application as a whole and not to any particular portions of this application. Where the context permits, words in the above Detailed Description using the singular or plural number may also include the plural or singular number respectively. The word "or" in reference to a list of two or more items covers all of the following interpretations of the word: any of the items in the list, all of the items in the list, and any combination of the items in the list.

The above detailed description of embodiments of the invention is not intended to be exhaustive or to limit the invention to the precise form disclosed above. While specific embodiments of, and examples for, the invention are described above for illustrative purposes, various equivalent modifications are possible within the scope of the invention, as those skilled in the relevant art will recognize. For example, while processes or blocks are presented in a given order, alternative embodiments may perform routines having steps, or employ systems having blocks, in a different order, and some processes or blocks may be deleted, moved, added, subdivided, combined, and/or modified. Each of these processes or blocks may be implemented in a variety of different ways. Also, while processes or blocks are at times shown as being performed in series, these processes or blocks may instead be performed in parallel, or may be performed at different times.

The teachings of the invention provided herein can be applied to other systems, not necessarily the system described above. The elements and acts of the various embodiments described above can be combined to provide further embodiments.

Conditional language used herein, such as, among others, "can," "might," "may," "e.g.," and the like, unless specifically stated otherwise, or otherwise understood within the context as used, is generally intended to convey that certain embodiments include, while other embodiments do not include, certain features, elements and/or states. Thus, such conditional language is not generally intended to imply that features, elements and/or states are in any way required for one or more embodiments or that one or more embodiments necessarily include logic for deciding, with or without author input or prompting, whether these features, elements and/or states are included or are to be performed in any particular embodiment.

While certain embodiments of the inventions have been described, these embodiments have been presented by way of example only, and are not intended to limit the scope of the disclosure. Indeed, the novel methods and systems described herein may be embodied in a variety of other forms; furthermore, various omissions, substitutions and changes in the form of the methods and systems described herein may be made without departing from the spirit of the disclosure. The accompanying claims and their equivalents are intended to cover such forms or modifications as would fall within the scope and spirit of the disclosure.

What is claimed is:

1. A memory package, comprising:

a plurality of input/output terminals via which the memory package communicates data and control/address signals with one or more external devices;

a plurality of stacked array dies including a first group of array dies and a second group of at least one array die, each array die having data ports;

at least a first die interconnect and a second die interconnect, the first die interconnect in electrical communication with the first group of array dies and not in electrical communication with the second group of at least one array die, the second die interconnect in electrical communication with the second group of at least one array die and not in electrical communication with the first group of array dies; and

a control die comprising at least a first data conduit between the first die interconnect and a first terminal of the plurality of input/output terminals, and at least a second data conduit between the second die interconnect and the first terminal, the first terminal being a data terminal, the control die further comprising a control circuit to control respective states of the first data conduit and the second data conduit in response to control signals received via one or more second terminals of the plurality of terminals.

2. The memory package of claim **1**, wherein the control signals include data path control signals for controlling the first and second data conduits.

3. The memory package of claim **1**, wherein the control circuit is configured to generate data path control signals for controlling the first and second data conduits in response to the received control signals.

4. The memory package of claim **3**, wherein the control signals include command/address signals and wherein the control die is configured to provide the command/address signals to the plurality of stacked array dies.

5. The memory package of claim **1**, wherein the first die interconnect comprises a first through-silicon via and wherein the second die interconnect comprises a second through-silicon via.

6. The memory package of claim **1**, wherein the control die further comprises chip-select conduits, the memory package further comprising:

third die interconnects coupled between respective chip-select conduits and respective ones of the plurality of stacked array dies.

7. The memory package of claim **1**, wherein a first number of array dies in the first group of array dies and a second number of at least one array die in the second group of at least one array die are selected in consideration of a load of the first die interconnect and a load of the second die interconnect so as to reduce a difference between a first load on the first data conduit and a second load on the second data conduit, the first load including a load of the first die interconnect, and a load of the first group of array dies, and the second load including a load of the second die interconnect and a load of the second group of at least one array die.

8. The memory package of claim **1**, wherein the respective states of the first data conduit and the second data conduit are controlled by one or more data path control signals, wherein the control die is configurable to operate in any one of a first mode and a second mode, and wherein:

in the first mode, the control die receives the data path control signals from the one or more external devices; and

in the second mode, the control die generates the data path control signals from at least some of the control/address signals received from the one or more external devices.

9. The memory package of claim **1**, wherein the control die further comprises command/address conduits configured to provide corresponding command/address signals to the array dies, the command/address signals including at least one memory cell address.

10. The memory package of claim **1**, wherein the control die further comprises one or more additional conduits configured to provide one or more of a supply voltage signal and a ground signal to the array dies.

11. A memory package, comprising:
a plurality of input/output terminals via which the memory package communicates data and control/address signals with one or more external devices;
a plurality of array dies arranged in a stack, including a first group of array dies and a second group of at least one array die;
at least a first die interconnect and a second die interconnect, the first die interconnect in electrical communication with the first group of array dies and not in electrical communication with the second group of at least one array die, the second die interconnect in electrical communication with the second group of at least one array die and not in electrical communication with the first group of array dies; and
a control die comprising at least a first data conduit between the first die interconnect and a first terminal of the plurality of input/output terminals, at least a second data conduit between the second die interconnect and the first terminal, and chip select conduits for providing chip select signals to respective array dies;
wherein the control die further comprises a control circuit to control respective states of the first data conduit and the second data conduit to drive a data signal to an array die selected by at least one of the chip-select signals.

12. The memory package of claim 11, wherein the chip select conduits pass through the control die.

13. The memory package of claim 11, wherein the chip select conduits include drivers to drive the chip select signals to the respective array dies.

14. The memory package of claim 11, wherein the first die interconnect comprises one or more through silicon vias and wherein the second die interconnect comprises one or more through silicon vias.

15. The memory package of claim 11, wherein: the first data conduit comprises at least a first driver having a first driver size, and the second data conduit comprises at least a second driver having a second driver size, and wherein the first driver size and the second driver size are both less than a driver size sufficient to drive a signal along a die interconnect in electrical communication with each of the plurality of array dies without significant signal degradation.

16. The memory package of claim 11, wherein the control die is configured to generate the chip select signals from control signals received via second terminals of the plurality of terminals, and wherein the chip select signal is generated using an address signal in the control/address signals.

17. The memory package of claim 11, wherein the control circuit controls the respective states of the first data conduit and the second data conduit in response to at least some of the control/address signals received via second terminals of the plurality of terminals.

18. The memory package of claim 16, wherein the control/address signals comprise: at least one command signal, at least one address signal, and at least one data path control signal.

19. The memory package of claim 11, wherein the control die is configured to generate data path control signals from the control/address signals received via second terminals of the plurality of terminals, and wherein the control circuit controls the respective states of the first data conduit in response to the data path control signals.

20. A method for optimizing load in a memory package comprising a plurality of array dies arranged in a stack and including a first group of array dies and a second group of at least one array die, at least a first die interconnect and a second die interconnect, a control die, and a plurality of input/output terminals via which the memory package communicates data and control/address signals with one or more external devices, the method comprising:
receiving a data signal at a first terminal of the plurality of input/output terminals;
receiving control signals at second terminals of the plurality of input/output terminals;
providing chip select signals to respective array dies through the control die, the chip select signals being related to at least some of the control signals; and
selecting one of a first driver and a second driver in the control die to drive the data signal via a corresponding one of the first die interconnect and the second die interconnect to an array die selected by at least one of the chip select signals, the first die interconnect in electrical communication with the first group of array dies and not in electrical communication with the second group of at least one array die, the second die interconnect in electrical communication with the second group of at least one array die and not in electrical communication with the first group of array dies.

21. The method of claim 20, further comprising:
selecting a first driver size for the first driver based, at least in part, on a load on the first driver; and
selecting a second driver size for the second driver based, at least in part, on a load on the second driver.

22. The method of claim 21, wherein the first driver size and the second driver size are both less than a driver size sufficient to drive a signal along a die interconnect in electrical communication with each of the plurality of array dies without significant signal degradation.

23. The method of claim 20, further comprising generating the chip select signals from at least some of the control signals.

24. The method of claim 20, wherein the control signals include the chip-select signals.

25. The method of claim 20, wherein the first and second die interconnects each comprises at least one through silicon vias.

26. The method of claim 20, wherein the chip select signals pass through through-silicon-vias in the control die.

27. The method of claim 20, further comprising generating data path control signals from at least some of the control signals, the data path control signals being used to select the one of the first driver and the second driver in the control die to drive the data signal.

28. The method of claim 20, wherein the one of the first driver and the second driver is selected using at least some of the control signals.

29. A memory module operable via a memory control hub, comprising:
a register device configured to receive command/address signals from the memory control hub and to generate control signals; and
a plurality of DRAM packages, each DRAM package comprising:
a plurality of data terminals via which the DRAM package communicates data with the memory control hub, and a plurality of control terminals to receive the control signals;
a plurality of array dies arranged in a stack, including a first group of array dies and a second group of at least one array die;
at least a first die interconnect and a second die interconnect, the first die interconnect in electrical communication with the first group of array dies and not in electrical communication with the second group of at least one

array die, the second die interconnect in electrical communication with the second group of at least one array die and not in electrical communication with the first group of array dies; and

a control die comprising at least a first data conduit between the first die interconnect and a first data terminal of the plurality of data terminals, at least a second data conduit between the second die interconnect and the first data terminal, and chip select conduits for providing chip select signals to respective array dies, the chip select signals being related to at least some of the control signals;

wherein the control die further comprises a control circuit to control respective states of the first data conduit and the second data conduit to drive a data signal to an array die selected by at least one of the chip-select signals.

30. The memory module of claim **29**, wherein the register device is further configured to perform rank multiplication by generating the chip select signals, and wherein the control signals include the chip select signals.

31. The memory module of claim **29**, wherein the control signals include data path control signals generated by the register device, the data path control signals being used to control the respective states of the first data conduit and the second data conduit.

32. The memory module of claim **29**, wherein the control die is further configured to perform rank multiplication by generating the chip select signals from at least some of the control signals that include at least one address signal.

33. The memory module of claim **29**, wherein the control signals include command/address signals, and the control die is configured to hold the command/address signals to control timing of the command/address signals.

34. The memory module of claim **29**, wherein the control die is configured to generate data path control signals from at least some of the control signals, the data path control signals being used to control the respective states of the first data conduit and the second data conduit.

* * * * *

US009318160B2

(12) **United States Patent**
Lee

(10) **Patent No.: US 9,318,160 B2**
(45) **Date of Patent: *Apr. 19, 2016**

(54) **MEMORY PACKAGE WITH OPTIMIZED DRIVER LOAD AND METHOD OF OPERATION**

(71) Applicant: **Netlist, Inc.**, Irvine, CA (US)

(72) Inventor: **Hyun Lee**, Ladera Ranch, CA (US)

(73) Assignee: **Netlist, Inc.**, Irvine, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **14/337,168**

(22) Filed: **Jul. 21, 2014**

(65) **Prior Publication Data**

US 2015/0070959 A1 Mar. 12, 2015

Related U.S. Application Data

(63) Continuation of application No. 13/288,850, filed on Nov. 3, 2011, now Pat. No. 8,787,060.

(60) Provisional application No. 61/409,893, filed on Nov. 3, 2010.

(51) **Int. Cl.**
G11C 5/06 (2006.01)
G11C 7/10 (2006.01)

(52) **U.S. Cl.**
CPC ***G11C 5/066*** (2013.01); ***G11C 5/06*** (2013.01); ***G11C 7/1057*** (2013.01); ***G11C 7/1084*** (2013.01)

(58) **Field of Classification Search**
CPC G11C 5/063; G11C 7/18; G11C 5/025; G11C 5/04
USPC 365/63, 51, 52
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,345,412 A 9/1994 Shiratsuchi
5,537,584 A 7/1996 Miyai et al.
(Continued)

FOREIGN PATENT DOCUMENTS

EP 1816570 A2 8/2007
JP 09237492 9/1997
(Continued)

OTHER PUBLICATIONS

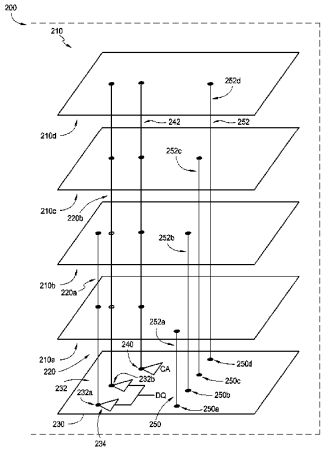
International Search Report and Written Opinion, International Application No. PCT/US14/48517, filed Jul. 28, 2014, dated Oct. 27, 2014.
(Continued)

Primary Examiner — Huan Hoang
(74) *Attorney, Agent, or Firm* — Jamie J. Zheng, Esq.

(57) **ABSTRACT**

An apparatus is provided that includes a plurality of array dies and at least two die interconnects. The first die interconnect is in electrical communication with a data port of a first array die and a data port of a second array die and not in electrical communication with data ports of a third array die. The second die interconnect is in electrical communication with a data port of the third array die and not in electrical communication with data ports of the first array die and the second array die. The apparatus includes a control die that includes a first data conduit configured to transmit a data signal to the first die interconnect and not to the second die interconnect, and at least a second data conduit configured to transmit the data signal to the second die interconnect and not to the first die interconnect.

20 Claims, 8 Drawing Sheets



(56) **References Cited**

U.S. PATENT DOCUMENTS

5,617,559	A	4/1997	Le et al.	
5,649,159	A	7/1997	Le et al.	
5,655,113	A	8/1997	Leung	
5,717,851	A	2/1998	Yishay et al.	
5,724,604	A	3/1998	Moyer	
5,729,716	A	3/1998	Lee et al.	
5,784,705	A	7/1998	Leung	
5,802,541	A	9/1998	Reed	
5,973,392	A	10/1999	Senba et al.	
6,011,710	A	1/2000	Wiggers	
6,173,357	B1	1/2001	Ju	
6,260,127	B1	7/2001	Olarig et al.	
6,446,158	B1	9/2002	Karabatsos	
6,553,449	B1	4/2003	Dodd et al.	
6,618,791	B1	9/2003	Dodd et al.	
6,717,855	B2	4/2004	Underwood	
6,747,887	B2	6/2004	Halbert et al.	
6,788,592	B2	9/2004	Nakata et al.	
6,832,303	B2	12/2004	Tanaka	
6,948,084	B1	9/2005	Manapat et al.	
7,093,066	B2	8/2006	Klein	
7,130,308	B2	10/2006	Haddock et al.	
7,133,960	B1	11/2006	Thompson et al.	
7,379,361	B2	5/2008	Co et al.	
7,464,225	B2	12/2008	Tsern	
7,865,674	B2	1/2011	Gower et al.	
8,089,795	B2	1/2012	Rajan	
8,189,328	B2	5/2012	Kanapathippillai	
8,233,303	B2	7/2012	Best	
8,516,185	B2	8/2013	Lee et al.	
8,689,064	B1	4/2014	Lee et al.	
8,782,350	B2	7/2014	Lee et al.	
8,787,060	B2 *	7/2014	Lee	365/63
2001/0008006	A1	7/2001	Klein	
2002/0038405	A1	3/2002	Leddige et al.	
2002/0048195	A1	4/2002	Klein	
2003/0070052	A1	4/2003	Lai	
2004/0098528	A1	5/2004	Janzen	
2005/0010737	A1	1/2005	Ware et al.	
2005/0257109	A1	11/2005	Averbuj	
2006/0262586	A1	11/2006	Solomon et al.	
2007/0070669	A1	3/2007	Tsern	
2007/0293094	A1	12/2007	Aekins	
2008/0104352	A1	5/2008	Talbot	
2008/0162790	A1	7/2008	Im	
2009/0103387	A1	4/2009	Shau	
2009/0296503	A1	12/2009	Chu et al.	
2011/0085406	A1	4/2011	Solomon et al.	
2011/0125966	A1	5/2011	Amidi et al.	

FOREIGN PATENT DOCUMENTS

JP	10092169	10/1998
JP	2010320270	12/1998
JP	2000285674	10/2000
JP	2000311485	10/2000
JP	2002184176	6/2002
JP	2003007963	1/2003
JP	2008046989	2/2008

OTHER PUBLICATIONS

Inter Partes Review Case No. IPR2014-01029, Petitioner's Request for Rehearing pursuant to 37 C.F.R. § 42.71, filed on Jan. 15, 2015.
Inter Partes Review Case No. IPR2014-01029, Decision Denying Request for Rehearing, Issued on Mar. 3, 2015.
Petition for Inter Partes Review filed on Jun. 24, 2014 for U.S. Pat. No. 8,516,185, IPR Case No. IPR2014-01029, and all associated documents including cited references and expert declarations, available at https://ptabtrials.uspto.gov.
Altera, ACEX iK, Programmable Logic Device Family, Data Sheet, May 2003, Ver 3.4.
Horowitz, "The Art of Electronics," Cambridge Univ. Press, 1989, selected pages.
Huang et al, "An Efficient Parallel Transparent BIST Method for Multiple Embedded Memory Buffers," VLSI Design 2001, p. 379.
Jacob, Bruce L.; "Synchronous DRAM Architectures, Organizations, and Alternative Technologies". University of Maryland, Dec. 10, 2002.
JEDEC Standard No. 21-C Section 4.5.7, 168 Pin Registered SDRAM DIMM Family, Release 7.
JEDEC 21-C, Section 4..6.1, 278 Pin Buffered SDRAM DIMM Family.
JEDEC Standard No. 21-C Section 4.1.2.5, Appendix E, "Specific PD's for Synchronous DRAM (SDRAM)," pp. 1-25.
JEDEC Standard, "Fully Buffered DIMM (FBDIMM): DFx Design for Validation and Test," JESD82-28, Feb. 2008.
Reese, "Introduction to Logic Synthesis using Verilog HDL," Morgan&Claypool Publishers, 2006, pp. 1-28.
Inter Partes Review Case No. IPR2014-01369, Decision Denying Institution of Inter Partes Review 37 C.F.R. § 42.108, issued Mar. 9, 2014.
Inter Partes Review Case No. IPR2014-01029, Petition for Inter Partes Review of U.S. Pat. No. 8,516,185, filed on Jun. 24, 2014.
Inter Partes Review Case No. IPR2014-01029, Exhibit 1008 to Petition for Inter Partes Review, "Declaration of Charles J. Neuhauser, Ph.D. under 37 C.F.R. § 1.68," filed on Jun. 24, 2014.
Inter Partes Review Case No. IPR2014-01029, Supplemental Petition for Inter Partes Review of U.S. Pat. No. 8,516,185, filed on Jul. 23, 2014.
Inter Partes Review Case No. IPR2014-01029, Patent Owner's Preliminary Response pursuant to 37 C.F.R. § 42.107, filed on Oct. 17, 2014.
Inter Partes Review Case No. IPR2014-01029, Decision Denying Institution of Inter Partes Review 37 C.F.R. § 42.108, issued Dec. 16, 2014.
Inter Partes Review Case No. IPR2014-01369, Corrected Petition for Inter Partes Review of Claims 1-19 of U.S. Pat. No. 8,516,185, filed on Sep. 22, 2014.
Inter Partes Review Case No. IPR2014-01369, Exhibit 1008 to Corrected Petition for Inter Partes Review, "Declaration of Dr. Nader Bagherzadeh under 37 C.F.R. § 1.68," filed on Sep. 22, 2014.
Inter Partes Review Case No. IPR2014-01369, Exhibit 1013 to Corrected Petition for Inter Partes Review, "Webster's II New College Dictionary," filed on Sep. 22, 2014.
Inter Partes Review Case No. IPR2014-01369, Exhibit 1014 to Corrected Petition for Inter Partes Review, "Standard Dictionary of Electrical and Electronics Terms," IEEE 1988, filed on Sep. 22, 2014.
Inter Partes Review Case No. IPR2014-00882, Corrected Petition for Inter Partes Review of U.S. Pat. No. 7,881,150. filed on Jul. 8, 2014.
Inter Partes Review Case No. IPR2014-00882, Exhibit 1007 to Petition for Inter Partes Review, "Declaration of Dr. Srinivasan Jagannathan," filed on Jun. 22, 2014.
Inter Partes Review Case No. IPR2014-00883, Corrected Petition for Inter Partes Review of U.S. Pat. No. 8,081,536. filed on Jul. 8, 2014.
Inter Partes Review Case No. IPR2014-00883, Exhibit 1011 to Petition for Inter Partes Review, "Declaration of Dr. Srinivasan Jagannathan," filed on Jun. 21, 2014.
Inter Partes Review Case No. IPR2014-01011, Corrected Petition for Inter Partes Review of U.S. Pat. No. 7,881,150. filed on Jul. 8, 2014.
Inter Partes Review Case No. IPR2014-01011, Exhibit 1007 to Petition for Inter Partes Review, "Declaration of Dr. Srinivasan Jagannathan." filed on Jun. 22, 2014.
McCluskey, Edward J., *Logic Design Principles with Emphasis on Testable Semicustom Circuits*, Prentice Hall, 1986, pp. 104-107 and 119-120.
Inter Partes Review of U.S. Pat. No. 7,881,150, IPR Case No. IPR2014-00882, Patent Owner's Preliminary Response Pursuant to 37 C.F.R. § 42.107, filed Oct. 7, 2014.
Inter Partes Review of U.S. Pat. No. 7,881,150, IPR Case No. IPR2014-00882, Decision—Institution of Inter Partes Review 37 C.F.R. § 42.108, issued Dec. 16, 2014.
Inter Partes Review of U.S. Pat. No. 8,081,536, IPR Case No. IPR2014-00883, Patent Owner's Preliminary Response Pursuant to 37 C.F.R. § 42.107, filed Oct. 7, 2014.

(56) **References Cited**

OTHER PUBLICATIONS

Inter Partes Review of U.S. Pat. No. 8,081,536, IPR Case No. IPR2014-00883, Decision—Institution of Inter Partes Review 37 C.F.R. § 42.108, issued Dec. 16, 2014.
Inter Partes Review of U.S. Pat. No. 7,881,150, IPR Case No. IPR2014-01011, Patent Owner's Preliminary Response Pursuant to 37 C.F.R. § 42.107, filed Oct. 7, 2014.
Inter Partes Review of U.S. Pat. No. 7,881,150, IPR Case No. IPR2014-01011, Decision—Institution of Inter Partes Review 37 C.F.R. § 42.108, issued Dec. 16, 2014.
Inter Partes Review of U.S. Pat. No. 7,881,150, IPR Case No. IPR2014-01011, Exhibit 3001 to Decision—Institution of Inter Partes Review, Excerpts from IEEE Dictionary, issued Dec. 16, 2014.
Inter Partes Review of U.S. Pat. No. 7,881,150, IPR Case No. IPR2014-01011, Exhibit 3002 to Decision—Institution of Inter Partes Review, Excerpts from IEEE Dictionary, issued Dec. 16, 2014.
Inter Partes Review of U.S. Pat. No. 7,881,150, IPR Case No. IPR2014-01011, Exhibit 3003 to Decision—Institution of Inter Partes Review, Excerpts from Oxford English Dictionary, issued Dec. 16, 2014.
Inter Partes Review of U.S. Pat. No. 7,881,150, IPR Case No. IPR2014-01011, Exhibit 3004 to Decision—Institution of Inter Partes Review, Excerpts from Oxford English Dictionary, issued Dec. 16, 2014.
U.S. District Court Northern District of California, Case No. 4:13-CV-05889-YGR, *Netlist, Inc.* v. *Smart Storage Systems, Inc., and Diablo Technologies, Inc.*, Smart Storage Systems, Inc.'s Invalidity Contentions, dated Jun. 6, 2014.
U.S. District Court Northern District of California, Case No. 4:13-CV-05889-YGR, *Netlist, Inc.* v. *Smart Storage Systems, Inc., and Diablo Technologies, Inc.*, Exhibits E.1-E.7 to "Smart Storage Systems, Inc.'s Invalidity Contentions," dated Jun. 6, 2014.
U.S. District Court Northern District of California, Case No. 4:13-CV-05889-YGR, *Netlist, Inc.* v. *Smart Storage Systems, Inc.'s, and Diablo Technologies, Inc.*, Exhibits F.1-F.5 to "Smart Storage Systems, Inc.'s Invalidity Contentions," dated Jun. 6, 2014.
U.S. District Court Northern District of California, Case No. 4:13-CV-05889-YGR, *Netlist, Inc.* v. *Smart Storage Systems, Inc., and Diablo Technologies, Inc.*, Exhibits G.1-G.6 to "Smart Storage Systems, Inc.'s Invalidity Contentions," dated Jun. 6, 2014.
U.S. District Court Northern District of California, Case No. 4:13-CV-05889-YGR, *Netlist, Inc.* v. *Smart Storage Systems, Inc., and Diablo Technologies, Inc.*, Exhibit H to "Smart Storage Systems, Inc.'s Invalidity Contentions," dated Jun. 6, 2014.
U.S. District Court Northern District of California, Case No. 4:13-CV-05889-YGR, *Netlist, Inc.* v. *Smart Storage Systems, Inc., and Diablo Technologies, Inc.*, Diablo Technologies, Inc.'s Invalidity Contentions, dated Jun. 6, 2014.
U.S. District Court Northern District of California, Case No. 4:13-CV-05889-YGR, *Netlist, Inc.* v. *Smart Storage Systems, Inc., and Diablo Technologies, Inc.*, Exhibits D-1 to D6 to "Diablo Technologies, Inc.'s Invalidity Contentions," dated Jun. 6, 2014.
U.S. District Court Northern District of California, Case No. 4:13-CV-05889-YGR, *Netlist, Inc.* v. *Smart Storage Systems, Inc., and Diablo Technologies, Inc.*, Exhibits F-1 to F-5 to "Diablo Technologies, Inc.'s Invalidity Contentions," dated Jun. 6, 2014.
U.S. District Court Northern District of California, Case No. 4:13-CV-05889-YGR, *Netlist, Inc.* v. *Smart Storage Systems, Inc., and Diablo Technologies, Inc.*, Exhibits G-1 to G-6 to "Diablo Technologies, Inc.'s Invalidity Contentions," dated Jun. 6, 2014.
U.S. District Court Northern District of California, Case No. 4:13-CV-05889-YGR, *Netlist, Inc.* v. *Smart Storage Systems, Inc., and Diablo Technologies, Inc.*, Exhibit H to "Diablo Technologies, Inc.'s Invalidity Contentions," dated Jun. 6, 2014.
Anonymous. (Dec. 1996). "Applications Note: Understanding DRAM Operation," IBM, 10 pages.
Behrens, S. "HP Printer Memory Explained", The ZonkPage, Last Updated Jan. 21, 2004. Accessed Feb. 10, 2013, Retrieved from the Internet: URL <http://warshaft.com/hpmem.htm>. 7pp.
Non-Final Office Action, U.S. Appl. No. 13/412,243, Jan. 2, 2014, 20 pages.
Non-final office action, U.S. Appl. No. 13/288,850, Oct. 11, 2013, 24 pages.
Non-final office action, U.S. Appl. No. 13/411,344, Dec. 31, 2013, 28 pages.
Non-final office action, U.S. Appl. No. 13/473,413, Nov. 17, 2011, 46 pages.
Response to non-final office action dated Oct. 14, 2013 for U.S. Appl. No. 13/288,850, filed Jan. 13, 2014, 15 pages.
Response to non-final office action dated Dec. 31, 2013 for U.S. Appl. No. 13/411,344, filed Mar. 31, 2014, 12 pages.
Patent Owner's Response to Office Action mailed Nov. 13, 2012 for Reexamination Control Nos. 95/000,578; 95/000,579, and 95/001,339, filed Jan. 14, 2013, 96 pages.
Patent Owner's Response to Office Action mailed Dec. 19, 2012 for Reexamination Control No. 95/001,758, filed Mar. 19, 2013, 61 pages.
Patent Owner's Response to Office Action mailed Sep. 26, 2013 for Reexamination Control No. 95/001,758, filed Nov. 26, 2013, 85 pages.
Third Party Requester's Comments after Non-Final Action mailed Sep. 26, 2013 for Reexamination Control No. 95/001,758, filed Dec. 26, 2013.
Patent Owner's Appeal Brief for Reexamination Control Nos. 95/000,546 and 95/000,577, filed Oct. 2, 2013. 46 pages.
Patent Trial and Appeal Board Decision on Appeal for Reexamination Control No. 95/001/337, mailed Jan. 16, 2014, 30 pages.
Patent Trial and Appeal Board Decision on Appeal for Reexamination Control No. 95/001/381, mailed Jan. 16, 2014, 24 pages.
Action Closing Prosecution mailed Mar. 27, 2014 for Reexamination Control No. 95/001,758, filed Sep. 14, 2011, 40 pages.
Action Closing Prosecution mailed Mar. 27, 2014 for Reexamination Control No. 95/001,339, filed Jun. 8, 2010, 106 pages.
Notice of Allowance, U.S. Appl. No. 12/504,131, Feb. 12, 2013, 52 pages.
Non-Final Office Action, dated Jan. 2, 2014, for U.S. Appl. No. 13/287,042, filed Nov. 1, 2011, 42 pages.
Response to Non-Final Office Action dated Jan. 2, 2014, filed Apr. 2, 2004, for U.S. Appl. No. 13/287,042, filed Nov. 1, 2011, 12 pages.
Office Action mailed Apr. 2, 2014, for Japanese Patent Application No. 2012-520662 and English translation thereof, 7 pages.

* cited by examiner

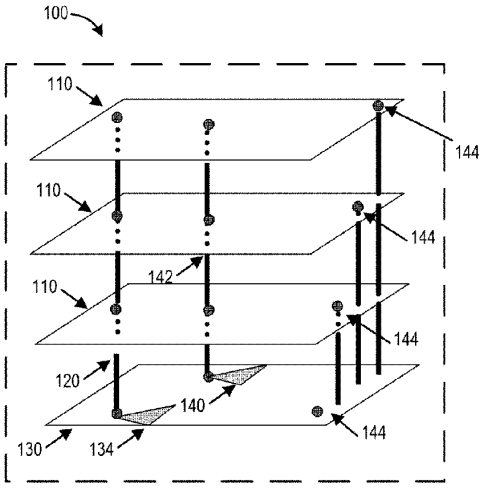


FIG. 1A

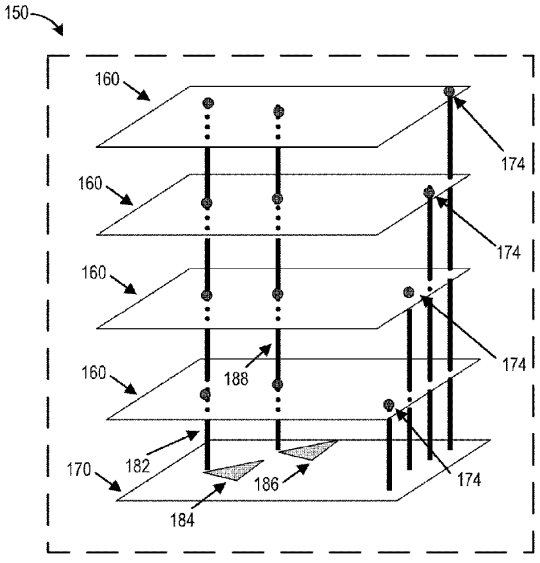


FIG. 1B

Samsung Electronics Co., Ltd.
Ex. 1001, p. 4

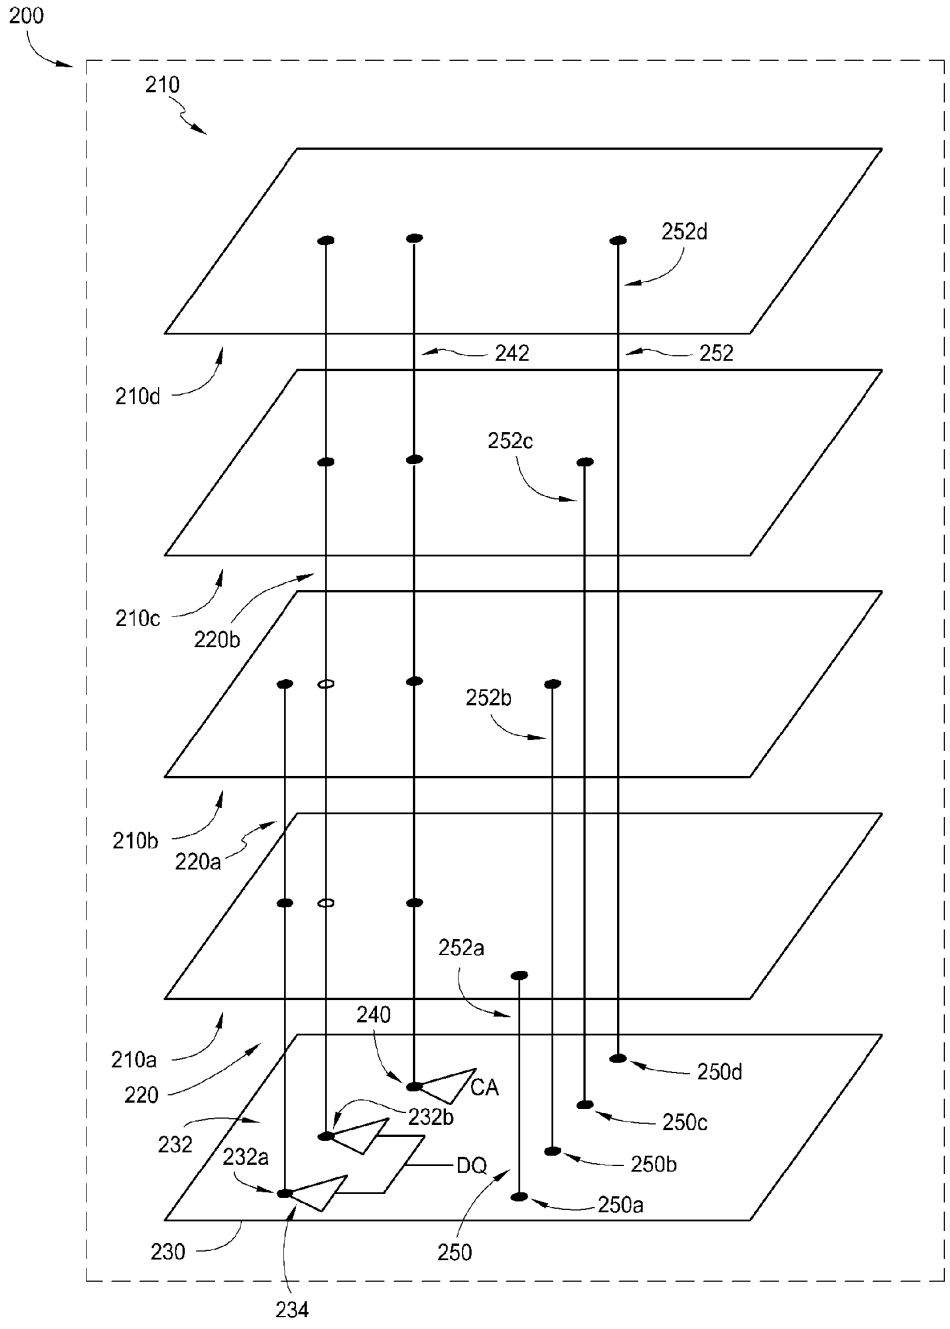

200
210
252d
242
252
210d
252c
210c
220b
252b
210b
220a
252a
240
210a
220
CA
250d
232b
250c
232
232a
DQ
250b
250a
230
250
234

FIG. 2

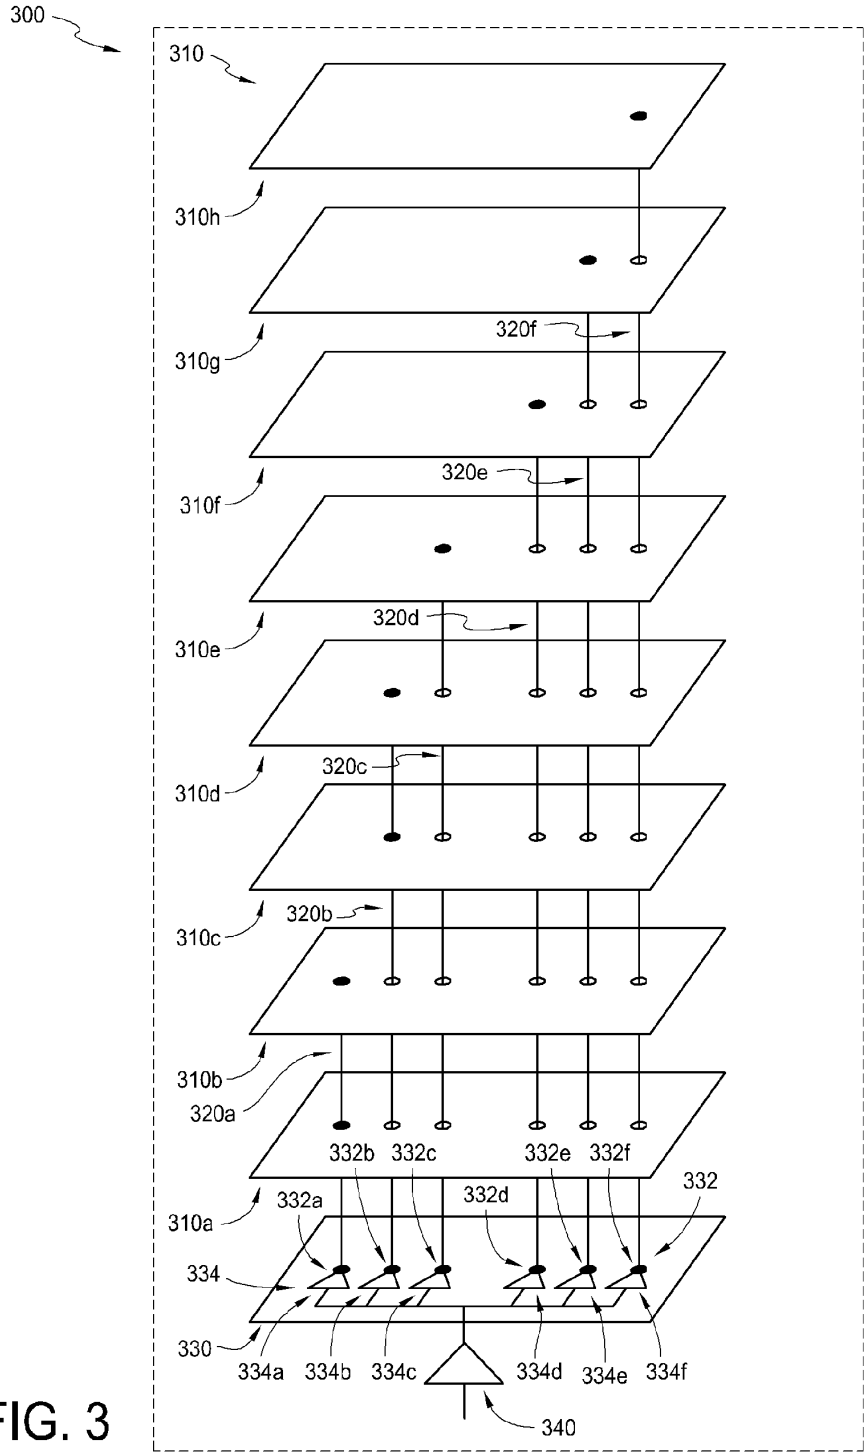

300
310
310h
320f
310g
320e
310f
320d
310e
320c
310d
320b
310c
310b
320a
332b
332c
332e
332f
332
332a
332d
310a
334
330
334a
334b
334c
334d
334e
334f
340

FIG. 3

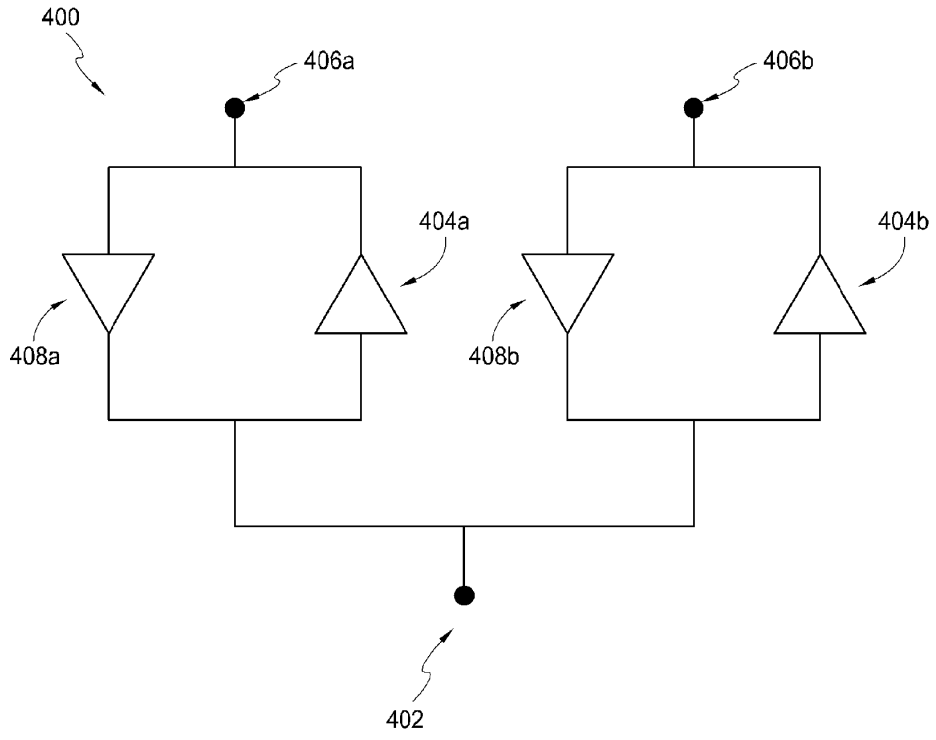


FIG. 4

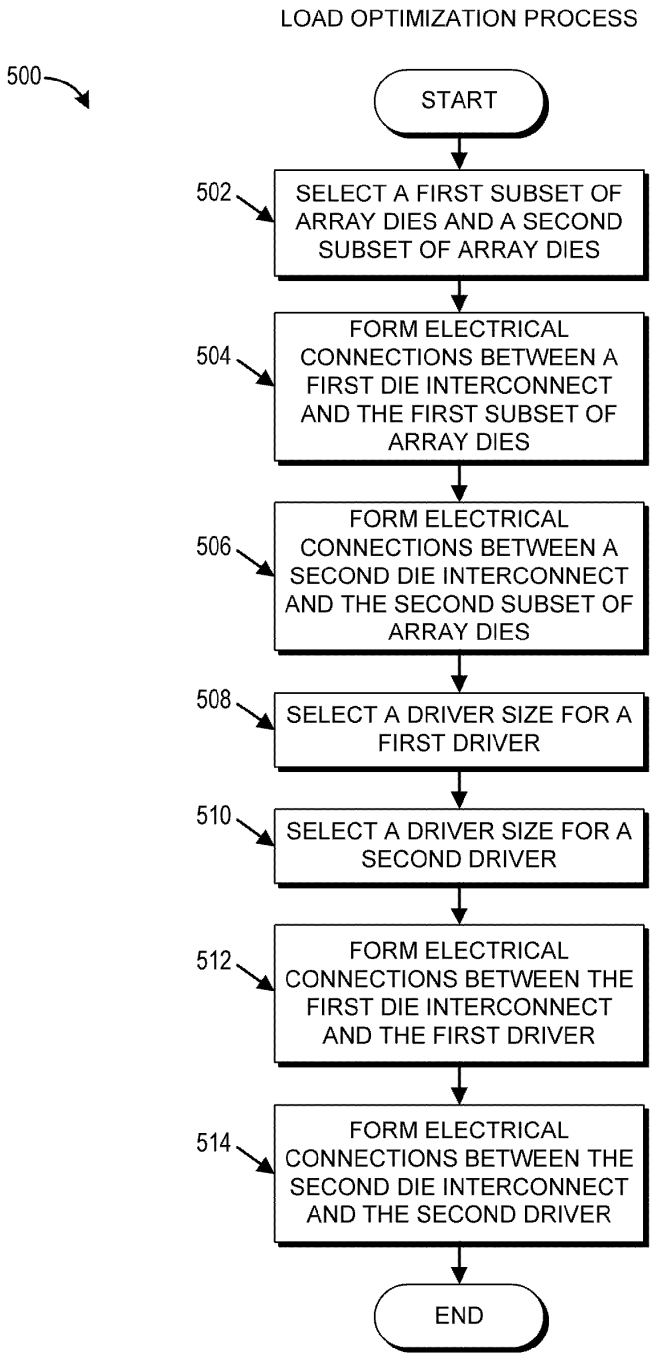

LOAD OPTIMIZATION PROCESS
500
START
502
SELECT A FIRST SUBSET OF ARRAY DIES AND A SECOND SUBSET OF ARRAY DIES
504
FORM ELECTRICAL CONNECTIONS BETWEEN A FIRST DIE INTERCONNECT AND THE FIRST SUBSET OF ARRAY DIES
506
FORM ELECTRICAL CONNECTIONS BETWEEN A SECOND DIE INTERCONNECT AND THE SECOND SUBSET OF ARRAY DIES
508
SELECT A DRIVER SIZE FOR A FIRST DRIVER
510
SELECT A DRIVER SIZE FOR A SECOND DRIVER
512
FORM ELECTRICAL CONNECTIONS BETWEEN THE FIRST DIE INTERCONNECT AND THE FIRST DRIVER
514
FORM ELECTRICAL CONNECTIONS BETWEEN THE SECOND DIE INTERCONNECT AND THE SECOND DRIVER
END

FIG. 5

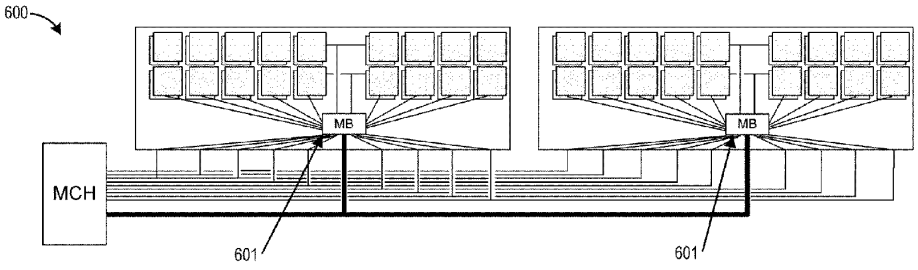


FIG. 6A

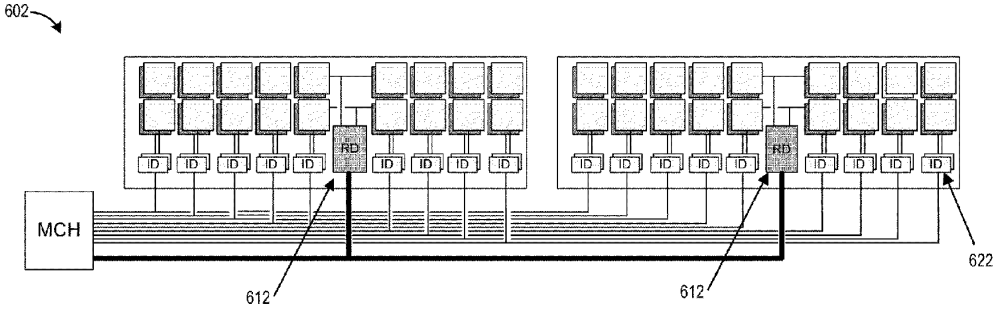


FIG. 6B

Samsung Electronics Co., Ltd.
Ex. 1001, p. 9

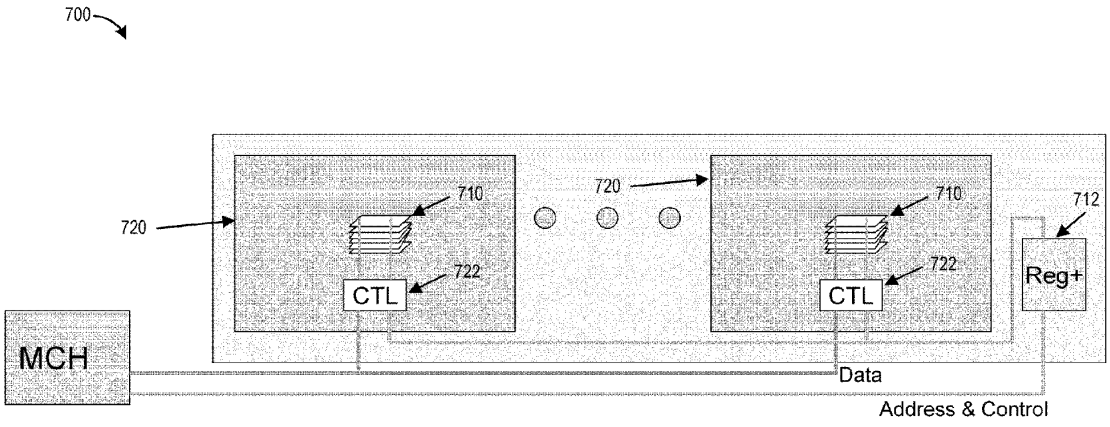


FIG. 7

Samsung Electronics Co., Ltd.
Ex. 1001, p. 10

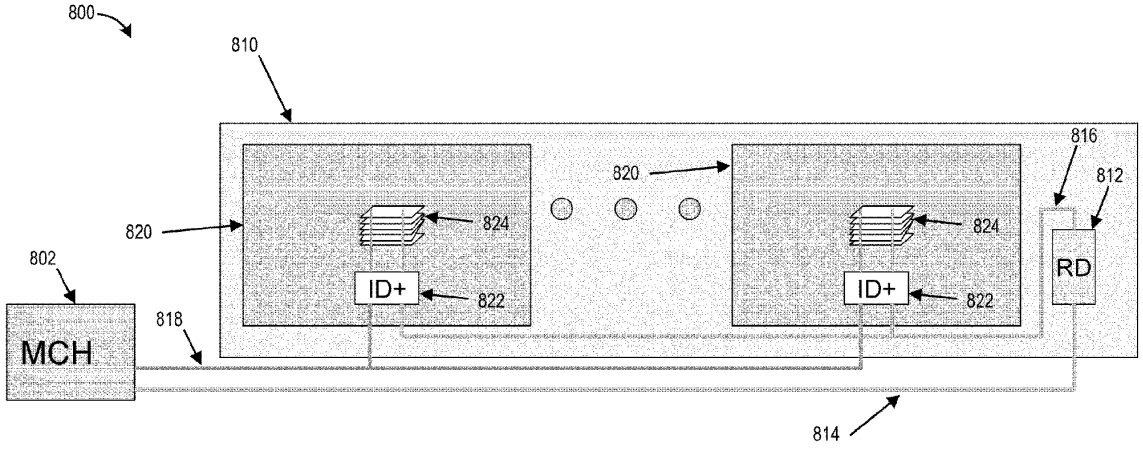


FIG. 8

Samsung Electronics Co., Ltd.
Ex. 1001, p. 11

MEMORY PACKAGE WITH OPTIMIZED DRIVER LOAD AND METHOD OF OPERATION

RELATED APPLICATION

This application is a continuation application of U.S. patent application Ser. No. 13/288,850, filed on Nov. 3, 2011, which claims the benefit of priority under 35 U.S.C. §119(e) of U.S. Provisional Patent Application No. 61/409,893, filed on Nov. 3, 2010, and entitled "ARCHITECTURE FOR MEMORY MODULE WITH PACKAGES OF THREE-DIMENSIONAL STACKED (3DS) MEMORY CHIPS." The disclosure of each of the above applications is hereby incorporated by reference in its entirety.

BACKGROUND

1. Technical Field

The present disclosure relates to memory devices and memory modules. More specifically, the present disclosure relates to systems and methods for reducing the load of drivers of memory packages included on the memory modules.

2. Description of the Related Art

Memory modules may include a number of memory packages. Each memory package may itself include a number of array dies that are packaged together. Each array die may include an individual semiconductor chip that includes a number of memory cells. The memory cell may serve as the basic building block of computer storage representing a single bit of data.

FIGS. **1**A and **1**B schematically illustrate examples of existing memory package designs currently used or proposed to be used to provide the dynamic random-access memory of memory modules. FIG. **1**A schematically illustrates a memory package **100** with three array dies **110** and a control die **130**. The control die **130** is configured to respond to signals received by the memory package **100** by sending appropriate control signals to the array dies **110** and includes a driver **134** for driving data signals to each of the array dies **110** via a corresponding die interconnect **120**. Further, the control die **130** includes a driver **140** for driving command and/or address signals to each of the array dies **110** via another corresponding die interconnect **142**. For simplicity, FIG. **1**A shows only a single driver **134**, die interconnect **120**, driver **140**, and die interconnect **142**. However, additional drivers and die interconnects may be included for each bit the memory package **100** is designed to support. Thus, a 16-bit memory may include 16 pairs of drivers and die interconnects for the data signals and other similar drivers and die interconnects for the command and/or address signals. Each array die **110** also includes a chip select port **144**, with the chip select ports **144** of the array dies **110** configured to receive corresponding chip select signals to enable or select the array dies for data transfer. The array dies **110** are configured to transfer data (e.g. read or write) to or from the selected memory cells identified by the command, address, and chip select signals via the die interconnects.

In some cases, the control die **130** may include memory cells and therefore, also serve as an array die. Thus, as can be seen from FIG. **1**A, the control die **130** may also include a chip select port **144**. Alternatively, the control die **130** and the array dies **110** may be distinct elements and the control die **130** may not include any memory cells.

FIG. **1**B schematically illustrates an example of a memory package **150** that includes four array dies **160** and a control die **170** that does not include memory cells. As can be seen in FIG. **1**B, each array die **160** includes a chip select port **174**. However, because the control die **170** does not also serve as an array die, the control die **170** does not include a chip select port. As with memory package **100**, memory package **150** includes a driver **184** that drives data signals to each of the array dies **160** along a corresponding die interconnect **182**. Further, the memory package **150** includes a driver **186** for driving command and/or address signals to each of the array dies **160** via another die interconnect **188**.

Generally, a load exists on each of the drivers **134**, **140**, **184**, and **186** by virtue of the drivers being in electrical communication with the corresponding die interconnects and the corresponding circuitry of the array dies. Thus, to drive a signal along a die interconnect, a driver typically must be large enough to overcome the load on the driver. However, generally a larger driver not only consumes more space on the control die, but also consumes more power.

SUMMARY

In certain embodiments, an apparatus is provided that comprises a plurality of array dies having data ports. The apparatus further comprises at least a first die interconnect and a second die interconnect. The first die interconnect is in electrical communication with at least one data port of a first array die of the plurality of array dies and at least one data port of a second array die of the plurality of array dies and not in electrical communication with the data ports of at least a third array die of the plurality of array dies. The second die interconnect is in electrical communication with at least one data port of the third array die and not in electrical communication with the data ports of the first array die and the data ports of the second array die. In addition, the apparatus comprises a control die. The control die comprises at least a first data conduit configured to transmit a data signal to the first die interconnect and to not transmit the data signal to the second die interconnect, and at least a second data conduit configured to transmit the data signal to the second die interconnect and to not transmit the data signal to the first die interconnect.

In certain embodiments, an apparatus is provided that comprises a plurality of array dies having ports. The apparatus further comprises at least a first die interconnect and a second die interconnect. The first die interconnect is in electrical communication with at least one port of a first array die of the plurality of array dies and at least one port of a second array die of the plurality of array dies and not in electrical communication with the ports of at least a third array die of the plurality of array dies. The second die interconnect is in electrical communication with at least one port of the third array die and not in electrical communication with the ports of the first array die and the ports of the second array die. In addition, the apparatus comprises a control die. The control die comprises at least a first conduit configured to transmit a signal to the first die interconnect and to not transmit the signal to the second die interconnect, and at least a second conduit configured to transmit the signal to the second die interconnect and to not transmit the signal to the first die interconnect. Moreover, a first load on the first conduit comprises a load of the first die interconnect, a load of the first array die, and a load of the second array die. In addition, a second load on the second conduit comprises a load of the second die interconnect and a load of the third array die.

In certain embodiments, a method is provided for optimizing load in a memory package. The memory package comprises a plurality of array dies, at least a first die interconnect and a second die interconnect, and a control die. The control die comprises at least a first driver and a second driver, the first

driver configured to drive a signal along the first die interconnect, and the second driver configured to drive the signal along the second die interconnect. The method comprises selecting a first subset of array dies of the plurality of array dies and a second subset of array dies of the plurality of array dies. The first subset of array dies and the second subset of array dies are exclusive of one another and are selected to balance a load on the first driver and on the second driver based at least in part on array die loads of array dies of the plurality of array dies and at least in part on die interconnect segment loads of segments of at least the first die interconnect and the second die interconnect. The method further comprises forming electrical connections between the first die interconnect and the first subset of array dies. In addition, the method comprises forming electrical connections between the second die interconnect and the second subset of array dies.

In certain embodiments, an apparatus is provided that comprises a register device configured to receive command/address signals from a memory control hub and to generate data path control signals. The apparatus further comprises a plurality of DRAM packages. Each of the DRAM packages comprises a control die. The control die comprises a plurality of command/address buffers and a data path control circuit configured to control command/address time slots and data bus time slots. The control die is configured to receive data signals from the memory control hub, the data path control signals from the register device, and command/address signals from the register device. Further, each of the DRAM packages comprises a plurality of DDR DRAM dies operatively coupled to the control die to receive the data signals from the control die. Moreover, the memory module is selectively configurable into at least two operational modes. The two operational modes comprise a first operational mode and a second operational mode. In the first operational mode the register device generates the data path control signals, and the control die uses the data path control signals to operate the data path control circuit. In the second operational mode, the control die operates the data path control circuit to provide the command/address signals to the plurality of DDR DRAM dies without decoding the command/address signals.

BRIEF DESCRIPTION OF THE DRAWINGS

Throughout the drawings, reference numbers are re-used to indicate correspondence between referenced elements. The drawings are provided to illustrate certain example embodiments of the inventive subject matter described herein and not to limit the scope thereof.

FIGS. **1**A and **1**B schematically illustrate examples of existing memory package designs.

FIG. **2** schematically illustrates an example embodiment of a memory package in accordance with the present disclosure.

FIG. **3** schematically illustrates another example embodiment of a memory package in accordance with the present disclosure.

FIG. **4** schematically illustrates an example embodiment of a driver structure of a control die in accordance with the present disclosure.

FIG. **5** presents a flowchart for an example embodiment of a load optimization process.

FIGS. **6**A and **6**B schematically illustrate an example of a Load Reduction Dual In-line Memory Module (LRDIMM) and a HyperCloud™ Dual In-line Memory Module (HCDIMM) architecture respectively.

FIG. **7** schematically illustrates an example of a Three-Dimensional Structure Dual In-line Memory Module (3DS-DIMM) architecture.

FIG. **8** schematically illustrates an example embodiment of a memory module architecture in accordance with the present disclosure.

DETAILED DESCRIPTION OF SPECIFIC EMBODIMENTS

In addition to the below, the following U.S. patents are incorporated in their entirety by reference herein: U.S. Pat. Nos. 7,289,386, 7,286,436, 7,442,050, 7,375,970, 7,254,036, 7,532,537, 7,636,274, 7,630,202, 7,619,893, 7,619,912, 7,811,097. Further, the following U.S. patent applications are incorporated in their entirety by reference herein: U.S. patent application Ser. Nos. 12/422,912, 12/422,853, 12/577,682, 12/629,827, 12/606,136, 12/874,900, 12/422,925, 12/504,131, 12/761,179, and 12/815,339.

Certain embodiments of the present disclosure reduce the size of drivers that are configured to drive a signal, such as a data signal, along a die interconnect to one or more array dies. Further, certain embodiments of the present disclosure reduce the power consumption of the drivers.

In certain embodiments, reducing one or both of driver size and driver power consumption may be accomplished by increasing the number of die interconnects and reducing the number of array dies that are in electrical communication with each die interconnect. For example, instead of one die interconnect in electrical communication with four array dies, there may be two die interconnects, each in electrical communication with a different pair of the four array dies.

In certain embodiments, determining the number of die interconnects and the number of array dies in electrical communication with each die interconnect is based, at least in part, on a load of each array die and a load of the die interconnect that is in electrical communication with one or more of the array dies.

In some embodiments, the load contribution from a die interconnect may be negligible compared to the load contribution from the array dies. In such embodiments, determining the number of die interconnects and the number of array dies in electrical communication with each die interconnect may be based, at least in part, on a load of each array die without considering the load of the die interconnect. However, as the physical size of a memory package shrinks, the load of a die interconnect becomes a non-negligible value relative to the load of the array dies. Thus, as memory packages become physically smaller, it becomes more important to consider the load of the die interconnect in determining the number of die interconnects and the number of array dies in electrical communication with each die interconnect. Advantageously, certain embodiments of the present disclosure account for both the loads of the array dies and the loads of the die interconnects on a conduit (e.g., driver) in determining the number of die interconnects to be used and the number of array dies in electrical communication with each die interconnect.

FIG. **2** schematically illustrates an example embodiment of a memory package **200** in accordance with the present disclosure. One example of a memory package that includes array dies and a control die is the Hybrid Memory Cube (HMC). Examples of an HMC compatible with certain embodiments described herein are described by the IDF2011 Intel Developer Forum website, http://www.intel.com/idf/index.htm, which includes presentations and papers from the IDF2011 Intel Developer Forum including the keynote address given by Justin Rattner on Sep. 15, 2011. Additional

examples of an HMC compatible with certain embodiments described herein are described by the Hybrid Memory Cube Consortium website, http://www.hybridmemorycube.org. The memory package **200** can include any type of memory package. For example, the memory package **200** may be a DRAM package, a SDRAM package, a flash memory package, or a DDR SDRAM package (e.g., DDR3, DDR4), to name a few. A memory module (not shown) may include one or more memory packages in accordance with the memory package **200**. Further, the memory package **200** may include input/output terminals (not shown) that are configured to be placed in electrical communication with circuitry of the memory module to transmit signals between the memory package **200** and a Memory Control Hub (MCH) (not shown).

The memory package **200** may include a plurality of array dies **210** (e.g., array dies **210***a*-**210***d*). The plurality of array dies **210** may be sealed within the memory package **200**. Further, circuitry of the array dies **210** may be in electrical communication with the input/output terminals of the memory package **200**. Although generally referred to as array dies herein, the array dies **210** may also be called slave dies or slave chips. Each of the array dies **210***a*-**210***d* may include circuitry (e.g., memory cells) (not shown) for storing data. Examples of array dies compatible with certain embodiments described herein are described by the existing literature regarding the Hybrid Memory Cube (e.g., as cited above). As illustrated in FIG. **2**, the plurality of array dies **210** may be arranged in a stack configuration known as a three-dimensional structure (30S). Examples of 30S compatible with certain embodiments described herein are described by the existing literature regarding the Hybrid Memory Cube (e.g., as cited above). However, the structure or layout of the plurality of array dies **210** is not limited as such, and other structures are possible in accordance with the present disclosure. For example, the plurality of array dies **210** may be arranged in a planar structure, or in a structure that combines 30S with a planar structure. Moreover, while the memory package **200** is illustrated as including four array dies, the memory package **200** is not limited as such and may include any number of array dies. For example, as illustrated in FIG. **3**, the memory package **200** may include eight array dies. As further examples, the memory package **200** may include three or sixteen array dies.

Each of the array dies **210** may include one or more data ports (not shown). The data ports enable electrical communication and data transfer between the corresponding memory circuitry of the array dies **210** and a communication pathway (e.g., a die interconnect).

In the example schematically illustrated in FIG. **2**, the memory package **200** includes a plurality of die interconnects **220** (e.g., die interconnects **220***a*, **220***b*). For example, certain versions of HMC have been reported to have 512 data ports per die, with the corresponding bits of each die all connected to a single die interconnect (e.g., TSV). Examples of die interconnects include, but are not limited to, through-silicon vias (TSV), conducting rods, wire bonds, and pins. (See e.g., U.S. Pat. Nos. 7,633,165 and 7,683,459.) Each of these die interconnects **220** may be coupled to, or in electrical communication with at least one data port of at least one of the array dies **210**. In certain embodiments, at least one of the die interconnects **220** is in electrical communication with at least one data port from each of at least two array dies **210** without being in electrical communication with a data port from at least one array die **210**, which may be in electrical communication with a different die interconnect **220**.

For example, die interconnect **220***a* may be in electrical communication with a data port from array die **210***a* and a data port from array die **210***b* (as illustrated by the darkened circles in FIG. **2**) and not in electrical communication with any data ports from array die **210***c* or any data ports from array die **210***d*. The data ports of array dies **210***a* and **210***b* in electrical communication with the die interconnect **220***a* can be corresponding to the same data bit (e.g., 00). Other die interconnects (not shown) can be in electrical communication with other data ports corresponding to other data bits (e.g., 01, 02, . . .) of array dies **210***a* and **210***b*. These other die interconnects can be electrically isolated from the corresponding data ports of array dies **210***c* and **210***d*.

However, continuing this example, the die interconnect **220***b* may be in electrical communication with a data port from array die **210***c* and a data port from array die **210***d* (as illustrated by the darkened circles in FIG. **2**) (e.g., corresponding to the same data bit, e.g., 00) without being in electrical communication with any data ports from array die **210***a* and array die **210***b*. Other die interconnects (not shown) can be in electrical communication with other data ports corresponding to other data bits (e.g., 01, 02, . . .) of array dies **210***c* and **210***d*. Despite not being in electrical communication with any data ports from array die **210***a* and **210***b*, in some implementations, the die interconnect **220***b* may pass through the array dies **210***a* and **210***b* (as illustrated by the unfilled circles) e.g., through through-holes or vias of array dies **210***a* and **210***b*. For some implementations, each of the array dies **210** may be in electrical communication with corresponding die interconnects **220**, without any of the die interconnects **220** being in electrical communication with all of the array dies **210**. Where existing systems may utilize a single die interconnect to be in electrical communication with the corresponding data ports of each array die (e.g., the data ports corresponding to the same bit), certain embodiments described herein utilize multiple die interconnects to provide electrical communication to the corresponding data ports of the array dies (e.g., the data ports corresponding to the same data bit) with none of the multiple die interconnects in electrical communication with data ports of all the array dies.

In addition to the plurality of array dies **210**, the memory package **200** includes a control die **230**, which may also be called a master die. Examples of master dies compatible with certain embodiments described herein are described by the existing literature regarding the Hybrid Memory Cube (e.g., as cited above). In some embodiments, the control die **230** may be one of the array dies **210**. Alternatively, the control die **230** may be a modified version of one of the array dies **210**. Thus, in some implementations, memory package **200** may include four dies or chips instead of the five illustrated in FIG. **2**. Further, in some embodiments, the control die **230** may comprise a logic layer (e.g., the logic layer of an HMC).

The control die **230** may include a number of data conduits **232**, which includes data conduits **232***a* and **232***b*. Each of these data conduits **232** may be configured to transmit a data signal to a single die interconnect **220**. For example, the data conduit **232***a* may be configured to transmit a data signal to the die interconnect **220***a* without transmitting the data signal to the die interconnect **220***b*. Conversely, the data conduit **232***b* may be configured to transmit the data signal to the die interconnect **220***b* without transmitting the data signal to the die interconnect **220***a* (e.g., data conduit **232***b* is electrically isolated from die interconnect **220***a* and data conduit **232***a* is electrically isolated from die interconnect **220***b*).

In some embodiments, the data conduits **232** may also include one or more drivers **234** as schematically illustrated by FIG. **2**. The drivers **234** may be configured to drive the data signals along the corresponding die interconnects **220**. In some embodiments, a single data conduit **232** or driver **234**

may be in electrical communication with multiple die interconnects **220**, each of which may be in electrical communication with different array dies **210**.

Each of the data conduits **232** may be configured to receive the data signal from a common source. For instance, the data conduit **232***a* and the data conduit **232***b* may each receive a substantially similar, if not identical, data signal from the same signal source (e.g. the data signal corresponding to the same data bit). The source of the data signal may include a data path, a driver, a latch, a pin, or any other construct that may provide a data signal to a data conduit.

The data conduits **232** may each be subject to a load. Although not limited as such, this load may be measured as a capacitive load, such as a parasitic capacitance. The load on each of the data conduits **232** may include at least a load of the die interconnect **220** to which the data conduit **232** is coupled or connected as well as a load of each array die **210** with which the die interconnect **220** is in electrical communication via a data port of the die interconnect **220**. Thus, for example, the load of the data conduit **232***a* may include loads of the die interconnect **220***a*, the array die **210***a*, and the array die **210***b*. Similarly, the load of the data conduit **232***b* may include loads of the die interconnect **220***b*, the array die **210***c*, and the array die **210***d*.

Generally, a load that would be on a data conduit that was in electrical communication with a die interconnect that was in electrical communication with at least one data port of each of the array dies **210** can be considered the maximum load for a data conduit. This maximum load is the load of a data conduit of a memory package that does not implement the teachings of this disclosure, but is in accordance with conventional configurations.

In some implementations, the difference between the load of the data conduit **232***a* and the load of the data conduit **232***b* is less than the maximum load for a data conduit as described above. Thus, in some cases, there may exist a degree of balance or equalization between the loads of the data conduits **232***a*, **232***b*. In some implementations, the difference between the load of the data conduit **232***a* and the load of the data conduit **232***b* is zero or substantially zero. In some embodiments, the length of each die interconnect **220**, and the number of array dies **210** in electrical communication with each die interconnect **220** may be selected to maintain the difference between the load of the data conduit **232***a* and the load of the data conduit **232***b* to be at or below a threshold load difference. For example, suppose that the load of each array die **210** is 1, the load of each segment of the die interconnects **220** is 0.25, and that the threshold load difference is 0.5. Using the configuration schematically illustrated in FIG. **2**, the load on the data conduit **232***a* in this example is 2.5 and the load on the data conduit **232***b* in this example is 3. Thus, in this example, the difference between the load of the data conduit **232***a* and the load of the data conduit **232***b* is at the threshold load difference value of 0.5. However, an alternative configuration that places the die interconnect **220***a* in electrical communication with only the array die **210***a*, and the die interconnect **220***b* in electrical communication with the array dies **210***b*-**210***d* would not satisfy the threshold load difference value of 0.5 of the above example. In the alternative configuration, the load on the data conduit **232***a* would be 1.25 and the load on the data conduit **232***b* would be 4. Thus, in the alternative configuration, the difference between the load of the data conduit **232***a* and the load of the data conduit **232***b* is 2.75, which is above the threshold load difference value of 0.5.

For certain embodiments, the load of each data conduit is less than the maximum load as described above. Thus, in some cases, the load of the data conduit **232***a* is less than the maximum load and the load of the data conduit **232***b* is less than the maximum load. Further, in many implementations, the combined load of the data conduit **232***a* and the data conduit **232***b* is less than the maximum load of a single data conduit. In other words, it is possible to design the data conduit **232***a* and the data conduit **232***b* to reduce the overall load compared to a single data conduit that is in electrical communication with a die interconnect that is in electrical communication with at least one data port of each of the array dies **210**. By reducing the overall load compared to the single data conduit, it is possible in many cases to reduce power consumption. Further, it is possible in many cases to maintain signal quality (e.g. maintain signal amplitude, maintain low signal distortion, etc.) while reducing power consumption. Advantageously, in a number of embodiments, by using multiple data conduits instead of a single data conduit, the speed of the memory package **200** can be increased. In some cases, this speed increase can include a reduced latency in accessing array dies **210** and/or operating the memory package **200** at a higher clock frequency.

Each of the die interconnects **220** may include any type of conducting material. For example, the die interconnects **220** may include copper, gold, or a conductive alloy, such as a copper/silver alloy. Further, the die interconnects **220** may include any type of structure for enabling electrical communication between the data conduits **232** and the data ports of the array dies **210**. For example, the die interconnects **220** may include a wire, a conducting rod, or a conducting trace, to name a few. Moreover, the die interconnects **220** may use vias, or through-silicon vias (TSVs) to couple with or to electrically communicate with an array die. For instance, die interconnect **220***a* may be connected with data ports of the array dies **210***a* and **210***b* using vias (illustrated by the filled or darkened circles). Examples of TSVs which may be used with the present disclosure are described further in U.S. Pat. Nos. 7,633,165 and 7,683,459.

In addition, the die interconnects **220** may use via holes to pass through an array die that is not configured to be in electrical communication with the die interconnect. For instance, die interconnect **220***b* may pass through array dies **210***a* and **210***b* using TSVs that do not enable electrical communication between the die interconnect **220***b* and data ports of the array dies **210***a* and **210***b* (illustrated by the unfilled circles). In this way, the array dies **210***a*, **210***b* are not responsive to the data signal beings transmitted by the die interconnect **220***b*. However, the die interconnect **220***b* may be connected with at least one data port of each of the array dies **210***c* and **210***d* using a via (illustrated by the filled or darkened circles). In cases where the die interconnect passes through an array die that is not configured to be in electrical communication with the die interconnect, the TSV may include an insulator or an air gap between the die interconnect and the array die circuitry that is large enough to prevent electrical communication between the die interconnect and the array die circuitry. In certain embodiments, the TSV for array dies that are configured to be in electrical communication with the die interconnect and for array dies that are not configured to be in electrical communication with the die interconnect may be configured the same. However, in such cases, electrical connections leading from the TSV of the array dies that are not configured to be in electrical communication with the die interconnect may not exist or may be stubs. These stubs are not configured to provide electrical communication with the memory cells of the array die.

Although FIG. **2** illustrates a single pair of data conduits **232** corresponding to a single pair of die interconnects **220**,

this is only to simplify the drawing figure. The memory package **200** may generally include as many additional data conduits and corresponding die interconnects as the number of bits the memory package **200** is designed or configured to support per memory address. Thus, if, for example, the memory package **200** is configured to be a 16-bit memory, the memory package **200** may include 16 pairs of data conduits **232** and 16 pairs of corresponding die interconnects **220**. Similarly, if the memory package **200** is configured as a 32- or 64-bit memory, the memory package **200** may include 32 or 64 pairs of data conduits **232** and 32 or 64 pairs of corresponding die interconnects **220**. Generally, the data conduits and die interconnects for each bit are configured identically. Thus, for the memory package **200**, each data conduit is configured to be in electrical communication with a die interconnect that is configured to be in electrical communication with a pair of the array dies **210**. However, it is possible that, in some embodiments, the data conduits and die interconnects of different bits may be configured differently.

In certain embodiments, the same die interconnects and the same corresponding data conduits may be used to transfer data both to and from the array dies **210**. In such embodiments, the die interconnects may be bi-directional. In alternative embodiments, separate die interconnects and corresponding data conduits may be used to transfer data to the array dies **210** and data from the array dies **210** to the control die **230**. Thus, such embodiments may include double the number of die interconnects and data conduits as embodiments that use the same die interconnect to transfer data to and from an array die.

In some embodiments, the control die **230** may include additional command/address conduits **240** and die interconnects **242**, which may be in electrical communication with at least one port of each of the array dies **210**. For simplicity, FIG. **2** shows only a single such conduit **240** and die interconnect **242**. The command/address conduits **240** are configured to provide corresponding signals to the die interconnects **242**. These signals may be command signals, address signals, or may serve as both command and address signals (e.g., may include a memory cell address and a write command or a read command) either simultaneously or based on a determining criterion, such as the edge of a clock signal. The command die **230** may include a command/address conduit **240** and corresponding die interconnect **242** for each bit of the command/address signals that the memory package **200** is configured to support. The number of bits of the command/address signals may be the same or may be different from the number of data bits of the memory package **200**.

In addition, in certain embodiments, the control die **230** may include a plurality of chip select conduits **250** (e.g., chip select conduits **250***a*-**250***d* as shown in FIG. **2**). Further, the control die **230** may include corresponding die interconnects **252** (e.g., die interconnects **252***a*-**252***d*) with one die interconnect **252** in electrical communication with one chip select conduit **250** and one array die **210**. Each of the die interconnects **252** may be in electrical communication with a different array die **210**. For example, the die interconnect **252***a* may be in electrical communication with the array die **210***a* and the die interconnect **252***b* may be in electrical communication with the array die **210***b*. Each of the chip select conduits **250** may be configured to provide a chip select signal to a corresponding array die **210** via a corresponding die interconnect **252**.

In some embodiments, the control die **230** may include additional drivers that are configured to drive the chip select signals along the die interconnects **252**. Alternatively, the chip select signals may be driven by drivers that are external to the control die **230**. For example, a register (not shown) that is part of a memory module that includes the memory package **200** may determine the chip select signals and drive the chip select signals to the array dies **210**. As a second example, the chip select signals may be provided by an MCH. In some embodiments, the control die **230** may determine the array die **210** to select based on, for example, an address signal. In such embodiments, the control die may generate the chip select signals.

FIG. **3** schematically illustrates another example embodiment of a memory package **300** in accordance with the present disclosure. In certain embodiments, some or all of the embodiments described above with respect to the memory package **200** may be applicable to the memory package **300**. However, for ease of illustration and to simplify discussion, certain elements are omitted in FIG. **3**, such as the chip select conduits. Nevertheless, it should be understood that the memory package **300** can include the same or similar elements as described above with respect to the memory package **200**, including, for example, the chip select conduits.

The memory package **300** may include a plurality of array dies **310** (e.g., array dies **310***a*-**310***h*). In the implementation illustrated in FIG. **3**, the memory package includes eight array dies. However, as stated earlier, the memory package **300** may include more or fewer array dies. Each of the array dies **310** may include one or more ports (not shown) that enable electrical communication between the circuitry of the array dies **310** and one or more die interconnects **320**. Each of these die interconnects **320** may be coupled to, or in electrical communication with at least one port of at least one of the array dies **310**. As with the memory package **200**, in certain embodiments, at least one of the die interconnects **320** is in electrical communication with at least one port from each of at least two array dies **310** without being in electrical communication with a port from at least one array die **310**, which may be in electrical communication with a different die interconnect **320**.

In addition to the plurality of array dies **310**, the memory package **300** includes a control die **330**. In some embodiments, the control die **330** may include a number of conduits **332** (e.g., conduits **332***a*-**332***f*). Each of these conduits **332** may be configured to transmit a signal to a single die interconnect **320**.

Further, implementations of the conduits **332** may include one or more drivers **334** (e.g., drivers **334***a*-**334***f*). Each of the drivers **334** may be configured to drive a signal along a corresponding die interconnect **320**. For instance, the driver **334***a* of the conduit **332***a* may be configured to drive a signal along the die interconnect **320***a* to one or more of the array dies **310***a* and **310***b*. As a second example, the driver **334***b* of the conduit **332***b* may be configured to drive a signal along the die interconnect **320***b* to one or more of the array dies **310***c* and **310***d*. Although the die interconnect **320***b* may pass through the array dies **310***a* and **310***b*, because the die interconnect **320***b*, in the example illustrated in FIG. **3**, is not configured to be in electrical communication with the array dies **310***a* and **310***b*, the driver **334***b* does not drive the signal to the array dies **310***a* and **310***b*.

In some embodiments, the signal can be a data signal, a command or address signal, a chip select signal, a supply voltage signal, or a ground voltage signal, to name a few. Further, as the signal is not limited to a data signal, in some embodiments, the conduits **332** may include conduits configured to provide signals other than data signals to the die interconnects **320**. For example, the conduits may include conduits configured to provide a command or address signal, a chip select signal, a supply voltage signal, or a ground

voltage signal to one or more die interconnects. Consequently, in some embodiments, the drivers **334** may be configured to drive signals other than data signals.

Generally, the signal that each of the drivers **334** drive to the corresponding die interconnects **320** is from a common source. Thus, each of the drivers **334**, in certain embodiments, is driving the same signal to each corresponding die interconnect **320**.

The size of the drivers **334** is generally related to the load on the driver **334**. In certain embodiments, the load on each driver **334** corresponds to the load of the respective conduit **332**. Although not limited as such, the load may be measured as a capacitive load, such as a parasitic capacitance.

The load on each of the conduits **332** may include at least the loads of the die interconnect **320** with which the conduit **332** is coupled or connected as well as the loads of each array die **310** with which the die interconnect **320** is in electrical communication via a port of the die interconnect **320**. Thus, for example, the loads of the conduit **332***a* may include the loads of the die interconnect **320***a*, the array die **310***a*, and the array die **310***b*. Similarly, the loads of the conduit **332***b* may include the loads of the die interconnect **320***b*, the array die **310***c*, and the array die **310***d*. The loads of both conduits **332***a* and **332***b* include a load of two array dies **310** because the corresponding die interconnects **320** are each configured to be in electrical communication with two array dies **310**. On the other hand, the load of the conduit **332***c*, which may include the loads of the die interconnect **320***c* and the array die **310***e*, includes a load of one array die **310***e* because the corresponding die interconnect **320***c* is configured to be in electrical communication with only one array die **310**.

As previously described with respect to FIG. **2**, a load that would be on a conduit that was in electrical communication with a die interconnect that was in electrical communication with at least one port of each of the array dies **310** can be considered the maximum load for a conduit. This maximum load is the load of a conduit of a memory package that does not implement the teachings of this disclosure, but is used in accordance with conventional configurations.

In some implementations, the difference between the loads of any pair of the conduits **332** is less than the maximum load for a conduit as described above. For instance, the difference between the load of the conduit **332***a* and the load of the conduit **332***b* is less than the maximum load. As a second example, the difference between the load of the conduit **334***f* and anyone of the conduits **334***a***334***e* is less than the maximum load. Thus, in some cases, the load on each of the conduits **332** may be, at least partially balanced or equalized to reduce or minimize the difference between the load of any pair of the conduits **332**. In some implementations, the difference between the load of a pair of the conduits **332** is zero or substantially zero. In some embodiments, the length of each die interconnect **320**, and the number of array dies **310** in electrical communication with each die interconnect **320** may be selected to maintain the difference between the load of any pair of the conduits **332** to be at or below a threshold load difference.

For certain embodiments, the load of each conduit is less than the maximum load as described above. Thus, for example, the load of the conduit **332***a*, the load of the conduit **332***b*, and the load of the conduit **332***c* are each less than the maximum load defined above.

In certain embodiments, the load associated with each of the array dies **310** may be substantially equivalent. For example, the load of the array die **310***a* may be substantially equal to the load of the array die **310***h*. Thus, the load contribution from the array dies **310** for a specific conduit **332** may be measured as a multiple of the array dies that are in electrical communication with a die interconnect **320** corresponding to a specific conduit **332**. For example, assuming that the load of each of the array dies **310** is 'L', then the contribution of the load from the array dies **310** to the conduit **332***a* would be 2 L because the die interconnect **320***a* corresponding to the conduit **332***a* is in electrical communication with two array dies **310**, array dies **310***a* and **310***b*. In alternative embodiments, the load of each of the array dies **310***a* may differ. For example, the load of array die **310***a* may be L and the load of the array die **310***b* may be 1.25 L.

Similar to the array dies **310**, in some embodiments, a die interconnect **320** can be considered to comprise a plurality of segments, with each segment of a die interconnect **320** contributing a substantially equivalent load to the load of the die interconnect **320**. In this case, the segment of the die interconnect **320** may refer to the portion of the die interconnect between two successive or adjacent dies (array die-to-array die, master die-to-array die, or both) along the die interconnect **320**. Further, the segment may be defined as a portion of the die interconnect **320** between the dies exclusive of a portion of the dies. For example, one segment of the die interconnect **320***a* may extend from the top of the array die **310***a* to the bottom of the array die **310***b*. Alternatively, the segment may be defined to include at least a portion of at least one of the array dies **310**. For example, one segment of the die interconnect **320***a* may extend from the center of array die **310***a* to the center of array die **310***b*. As a second example, one segment of the die interconnect **320***a* may extend from the top of the control die **330** to the top of the array die **310***a*, and therefore may include a portion of the die interconnect **320***a* extending from the bottom of the array die **310***a* to the top of the array die **310***a*. In some implementations, the segments are substantially equal in length to each other. Moreover, the load contribution of each segment may be substantially equal to each other. Alternatively, the segments may be unequal in length and/or may each contribute a different load to the total load of a die interconnect **320**. Further, in some cases, the load contribution of a segment of the die interconnect **320** that is in electrical communication with a port of an array die **310** may differ from the load contribution of a segment of the die interconnect **320** that is not in electrical communication with a port of an array die **310**.

In some cases, the load contribution of each segment of the die interconnect **320** may be measured as a fraction of the load contribution from an array die **310**. For example, the load of one segment of the die interconnect **320***a* may be equivalent to one quarter of the load of an array die **310**. Thus, for example, the load of the conduit **332***a* may be 2.5 L assuming a load contribution of L per array die **310** (two in this case) in electrical communication with the die interconnect **320***a* and a load contribution of 0.25 L per segment (two in this case) of the die interconnect **320***a*. As a second example, the load of the conduit **332***f* may be 3 L assuming the same load values as the previous example and a load contribution from one array die **310***h* and eight die interconnect **320***f* segments. Table 1 specifies the capacitive load values for each conduit **332** assuming, as with the previous two examples, that the load of each of the array dies **310** is L, and that the load of each segment of the die interconnects **320** is 0.25 L. Table 1 also specifies the deviation in load from the conduits having the highest load value, which in this example are conduits **332***b* and **332***f*.

TABLE 1

Conduit	Number of Array Dies	Number of Die Interconnect Segments	Capacitive Load	Deviation from Maximum Load
332a	2	2	2.5 L	0.5
332b	2	4	3 L	0
332c	1	5	2.25 L	0.75
332d	1	6	2.5 L	0.5
332e	1	7	2.75 L	0.25
332f	1	8	3 L	0

As can be seen from Table 1, the maximum load of any conduit **332**, using the example values previously described, is 3 L, or rather three times the load of a single array die **310**. Assuming the same example values, the load of a conduit in electrical communication with a die interconnect that itself is in electrical communication with each array die **310** would be 10 L. Thus, certain embodiments of the present disclosure enable a reduction in the load of the conduits **332**. Consequently, in some embodiments, the drivers **334** may each be smaller than a single driver that is configured to drive a signal from a conduit along a single die interconnect that is in electrical communication with a port from each of the array dies **310**. Moreover, the drivers **334** may include smaller transistor sizes than a single driver that is configured to drive a signal to each of the array dies **110**.

As can be seen from Table 1, conduits **332***b* and **332***f* have the largest capacitive load of the group of conduits, which is 3 L. Conduit **332** has the smallest capacitive load of the group of conduits, which is 2.25 L and which is a deviation of only 0.75 from the maximum load. Thus, in some embodiments, each of the drivers **334** may be substantially similar in size. In certain implementations, the drivers **334** may vary in size based on the total capacitive load on each conduit **332**. Thus, the driver **334***f* may be larger than the driver **334***e*. Alternatively, each driver **334** may be substantially equal and may be configured based on the drivers **334** with the largest load, which are drivers **334***b* and **334***f* in the example illustrated in FIG. **3** and Table 1.

In some embodiments, the capacitive load of each conduit **332** or driver **334** can be calculated using formula (1).

$$CL=AD+S/M \quad (1)$$

In formula (1), CL represents the capacitive load of a conduit **332** or driver **334**, AD represents the number of array dies **310** in electrical communication with a die interconnect **320** that is in electrical communication with the conduit **332** and/or driver **334**, **8** represents the number of die interconnect segments of the die interconnect **320**, and M represents the ratio of the load of an array die **310** to the load of a segment of a die interconnect **320**. Thus, using formula (1) and the example values described above, the load of the driver **332***a*, for example, can be calculated with the following values: AD=2, for two array dies **310**; 8=2, for the two segments of the die interconnect **320***a*; and M=4, for the ratio of the load of an array die **310**, L, to the load of a segment of the die interconnect **320***a*, 0.25 L. Therefore, as can be seen from Table 1, the capacitive load of the driver **332***a* is 2.5 L where L is the load of a single array die.

In some cases, the load on each conduit **332** and/or driver **334** may be evenly balanced. In other words, the load on each conduit **332** and/or driver **334** may be substantially the same. To achieve a balanced load, each of the conduits **332** may be in electrical communication with a combination of array dies **310** and die interconnect **320** segments that results in a load that is substantially equivalent to the loads of the other conduits **332**. In certain embodiments, the load of the conduits **332** is balanced despite each conduit **332** being in electrical communication with a different subset of the array dies **310**. However, in alternative embodiments, the load of each conduit **332** and/or driver **334** may differ. This difference in the load of the conduits **332** and/or drivers **334** may be a design decision (e.g. to maintain a specific number of drivers). Alternatively, or in addition, the load difference between conduits **332** and/or drivers **334** may occur because the loads of the array dies **310** and the die interconnect segments **320** do not allow for perfect or substantially even load balancing.

As previously stated, in some embodiments the loads of the conduits **332** and/or drivers **334** may be balanced to minimize the difference between any pair of conduits **332** and/or drivers **334**. For example, as illustrated in FIG. **3**, a driver in electrical communication with a longer die interconnect, which may consequently include more die interconnect segments, is likely to obtain a larger load contribution from the die interconnect than a driver in electrical communication with a shorter die interconnect (e.g., compare driver **334***f* to **334***a*). Thus, the driver in electrical communication with the longer die interconnect may be in electrical communication with fewer array dies than the driver in electrical communication with the shorter die interconnect (e.g. compare driver **334***f* to **334***a*). The selection of die interconnect length and the selection of array dies to place in electrical communication with the die interconnects may therefore be dependent on the load of each segment of the die interconnects **320** and the loads of the array dies **310**.

Alternatively, or in addition, the loads of the conduits **332** and/or drivers **334** may be balanced to reduce the maximum load of each conduit **332** and/or driver **334**. Further, in some embodiments, the maximum load of each conduit **332** and/or driver **334** may be reduced to maintain the load of each conduit **332** and/or driver **334** to be at or below a threshold load difference.

Although not illustrated in FIG. **3**, each of the conduits **332** may include one or more additional drivers configured to drive a signal received from an array die via a corresponding die interconnect **320**. The additional drivers may drive the signal to a processor, a register, a latch, or any other component or device that mayor may not be part of a memory module that includes the memory package **300**. In some implementations, one or more of the die interconnects **320** are configured to be bi-directional, thereby enabling a signal to be driven to the array dies **310** via the die interconnect **320**, and to enable a signal received from the array dies **310** along the same die interconnect **320** to be transmitted to a corresponding conduit **332** of the control die **330**. Alternatively, the memory package **300** comprises one or more die interconnects **320** that are configured to enable a signal to be driven to the array dies **310** without enabling a signal to be received from the array dies **310** (e.g., the die interconnects **320** may not be bidirectional). In certain embodiments where one or more of the die interconnects **320** are not bi-directional, the memory package **300** may include additional die interconnects (not shown) that are in electrical communication with the one or more additional conduits or drivers and that are configured to enable a signal from the array dies **310** to be transmitted to the one or more additional conduits or drivers.

In addition to the drivers **334** of the control die **330**, the memory package **300** may include one or more pre-drivers **340**. A pre-driver **340**, shown schematically in FIG. **3**, may be large enough to drive a signal (e.g., a data signal) to any of the drivers **332** and subsequently, to any of the array dies **310**. Thus, using the same example values as described above in relation to Table 1, a load of the pre-driver **340** may be at least 101, which is the load of a conduit in electrical communica-

tion with a die interconnect that is in electrical communication with a port from each of the array dies **310**. In some embodiments, the memory package **300** may include any number of additional drivers and/or latches for buffering and/or driving the signal to any of the array dies **310**.

Just as the memory package **300** may include a pre-driver **340** for providing a signal to the conduits **332**, the memory package **300** may include a post-driver (not shown) for driving an output signal from the control die **330**. This post-driver may drive the output signal to additional latches and/or drivers. In some embodiments, the post-driver may drive the signal from the memory package **300** to a bus or other electrical path that is in electrical communication with the memory package **300**.

FIG. **4** illustrates an example embodiment of a driver structure **400** of a control die (e.g. control die **230**) in accordance with the present disclosure. The driver structure **400** is configured for a single data bit. Thus, a control die for a 32-bit memory package will generally include at least 32 instances of the driver structure **400**, one per bit. In some embodiments, the control die **230** may include the driver structure **400** in place of the combination of drivers **232***a* and **232***b*. Similarly, in certain embodiments, the control die **330** may include a structure similar to the driver structure **400** in place of the drivers **334**. In such embodiments, the driver structure **400** would be modified to enable a signal (e.g., data signal) to be driven to the six conduits **332**.

The driver structure **400** can include an input/output port **402** configured to receive or send a signal, such as a data signal. Signals received at the input/output port **402** can be provided to the drivers **404***a* and **404***b*. In turn, these drivers **404***a* and **404***b* can drive the signal to conduits **406***a* and **406***b* respectively, which are in electrical communication with corresponding die interconnects. Each of the die interconnects may be in electrical communication with different array dies in accordance with certain embodiments described herein.

In some embodiments, the driver structure **400** is bi-directional. In such embodiments, operation of the drivers **404***a* and **404***b*, and **408***a* and **408***b* may be controlled or enabled by a control signal (e.g., a directional control signal). This control signal, in some cases, may correspond to one or more of a command/address signal (e.g., read/write) and a chip select signal. In some implementations, the drivers **404***a* and **404***b* may drive a signal to the array die corresponding to the chip select signal, and not to array dies that do not correspond to a chip select signal. For example, suppose the conduit **406***a* is in electrical communication with array dies one and two (not shown), and that the conduit **406***b* is in electrical communication with array dies three and four (not shown). If a chip select signal is received corresponding to array die two, then in some implementations, driver **404***a* may be configured to drive a signal along a die interconnect in electrical communication with the conduit **406***a* to array die one and two. In this example, the driver **404***b* would not drive the signal because the chip select signal does not correspond to either array die three or array die four.

In some embodiments, the drivers **404***a* and **404***b* may drive a signal to all of the array dies. For example, assume the memory package that includes the driver structure **400** also includes four array dies. If the die interconnect in electrical communication with the conduit **406***a* is in electrical communication with two of the array dies, and if the die interconnect in electrical communication with the conduit **406***b* is in electrical communication with the other two array dies, then the drivers **404***a* and **404***b* may drive a signal to all four of the array dies of this example.

Similar to the input/output port **402**, the conduits **406***a* and **406***b* can be configured to receive a signal from the die interconnects that are in electrical communication with the conduits **406***a* and **406***b*. In turn, the signal received at one of the conduits **406***a* and **406***b* can be provided to drivers **408***a* and **408***b* respectively. The drivers **408***a* and **408***b* can each be configured to drive a signal received from a respective data conduit **406***a* and **406***b* to the input/output port **402** and to another component that may be in electrical communication with the input/output port **402**, such as a MCH.

In some embodiments, one or more chip selects, as illustrated in FIG. **2**, may be used to select, determine, or enable the array die that communicates the signal to or from one or more of the conduits **406***a* and **406***b* and the drivers **408***a* and **408***b*. Similarly, in some embodiments, the chip select may select, determine, or enable the array die to receive and/or respond to the signal driven by the drivers **404***a* and **404***b* to the array dies.

It should be noted that the driver structure **400** is a non-limiting example for arranging drivers in a control die. Other driver structures are possible. For example, instead of the drivers **404***a* and **408***a* being in electrical communication with the same conduit **406***a*, each of the drivers **404***a* and **408***a* can be in electrical communication with a separate conduit and consequently a separate die interconnect. In such an example, the drivers **404***a* and **408***a* may still be in electrical communication with the same input/output port **402**, or each driver may be in electrical communication with a separate port, the driver **404***a* in electrical communication with an input port, and the driver **408***a* in communication with an output port.

FIG. **5** presents a flowchart for an example embodiment of a load optimization process **500**. In certain embodiments, the load optimization process **500** may be performed, at least in part, by one or more computing systems. Further, the process **500**, in some embodiments, may be used to optimize one or more loads in a memory package (e.g. memory package **200** or memory package **300**). Optimizing the loads in the memory package can include optimizing the load on one or more conduits and/or drivers. As previously described with respect to FIGS. **2** and **3**, the memory package can include a plurality of array dies, a plurality of die interconnects, and a control die. Furthermore, the control die can include a plurality of drivers, each of which may be configured to drive a signal along a die interconnect.

In some implementations, the process **500** can comprise selecting a first subset of array dies and a second subset of array dies from a plurality of array dies (e.g., array dies **310**), as shown in operational block **502**. Generally, the first subset of array dies and the second subset of array dies may be exclusive of one another. Thus, the first subset of array dies does not include any array dies from the second subset of array dies and vice versa. However, in some cases there may be some overlap between the first subset of array dies and the second subset of array dies. Further, in some cases, at least one of the subsets of array dies includes more than one array die. For instance, the first subset of array dies may include two array dies, and the second subset of array dies may include one array die, which, depending on the embodiment, mayor may not be included in the first subset of array dies.

Further, the first subset of array dies and the second subset of array dies may be selected to balance a load on a first driver and a load on a second driver based, at least in part, on the loads of the array dies **310** and on the loads of the die interconnect segments from a first and a second die interconnect. The first die interconnect may be in electrical communication with the first driver and the second die interconnect may be in

electrical communication with the second driver. In some embodiments, the first subset of array dies and the second subset of array dies are selected to balance a load on a first conduit and a load on a second conduit. The load on each driver and/or conduit may be calculated using formula (1) as previously described above. In some embodiments, the first subset of array dies and the second subset of array dies may be selected to balance a load on a first driver and a load on a second driver based, at least in part, on the loads of the array dies **310** without considering the loads of the die interconnect segments from the first die interconnect and the second die interconnect.

The process **500** further comprises forming electrical connections between the first die interconnect and the first subset of array dies in an operational block **504**. The process **500** further comprises forming electrical connections between the second die interconnect and the second subset of array dies in an operational block **506**. Forming the electrical connections places the die interconnects in electrical communication with the respective subsets of array dies. In some embodiments, forming the electrical connections can comprise forming electrical connections between the die interconnects and at least one port from each array die of the respective subsets of array dies.

The process **500** further comprises selecting a driver size for a first driver at block **508** and selecting a driver size for a second driver at block **510**. Selecting the driver size can be based, at least in part, on the calculated load on the driver. Generally, the greater the load on the driver, the larger the driver is selected to drive a signal along, for example, a die interconnect. The size of the driver may be adjusted by the selection of the transistor size and/or number of transistors included in the driver. A larger driver often consumes more power than a smaller driver. Thus, in certain embodiments, balancing the loads on the drivers to reduce the load on each driver can reduce the power consumption of a memory package.

In some embodiments, the size of the first driver and the size of the second driver are both less than the size sufficient for a driver to drive a signal along a die interconnect to each of the array dies **310** (e.g., with less than a predetermined or threshold signal degradation). The threshold signal degradation can be based on anyone or more characteristics of a signal. For example, the threshold signal degradation can be based on the amplitude of the signal, the frequency of the signal, the noise distortion included in or introduced into the signal, or the shape of the signal, to name a few.

The process **500** further comprises forming electrical connections between the first die interconnect and the first driver in an operational block **512**. Similarly, the process **500** further comprises forming electrical connections between the second die interconnect and the second driver in an operational block **514**. Forming the electrical connections places the die interconnects in electrical communication with the respective drivers. In some embodiments, forming the electrical connections can comprise forming electrical connections between the die interconnects and a data conduit that includes the respective driver.

Advantageously, certain embodiments of the present disclosure reduce the load on each conduit and on each corresponding driver. In certain embodiments, reducing the load on a driver may increase the speed of data transfer between the array dies and other components in a computer system (e.g. a MCH or a processor). Further, reducing the load on a driver may result in reduced power consumption by the driver and consequently, by the memory package. In addition, certain embodiments of the present disclosure may minimize current switching noise.

Operational Modes

A proposed three-dimensional stacking (3DS) standard for dual in-line memory modules (DIMMs) being considered by the Joint Electron Devices Engineering Council (JEDEC) addresses three major shortcomings in the current JEDEC registered DIMM (RDIMM) and JEDEC load reduced DIMM (LRDIMM) standards (an example LRDIMM structure **600** is schematically illustrated in FIG. **6**A). These shortcomings include: a) The DIMM density limitation due to the fixed number of chip-select signals received by the DIMM from the system memory controller. b) The performance loss due to the increased load on the data bus as the DIMM density (e.g., the number of DRAM devices and number of ranks) increases. c) The upper bound of the DIMM density due to the physical DIMM form factor.

Further, the LRDIMM structure **600** may have timing issues due to signals passing through a single memory buffer **601**. In addition, the increased size of the data path of the LRDIMM architecture **600**, compared to the HCDIMM architecture **602**, an example of which is schematically illustrated in FIG. **68**, may result in more latency and signal integrity issues.

FIG. **7** schematically illustrates an example memory module **700** that has previously been proposed for the 3DS-DIMM standard. The proposed 3DS Dual In-line Memory Module (3DS-DIMM) schematically illustrated in FIG. **7** attempts to address the above shortcomings using two components, a 3DS register **712** and a controller die **722**. In some implementations, the 3DS-DIMM **700** includes the 3DS register **712** (also known as an enhanced DDR3 register). The 3DS register **712** may include a DDR3 JEDEC standard register with a "rank multiplication" circuit, which increases the number of the output chip-select signals for selecting an array die **710** by decoding one or more higher-order row or column address bits with the incoming chip-select signals from the system controller. The 3DS Register **712** can also include a command/address buffer, a register operational code buffer (RC word buffer), and rank multiplication logic (e.g., supporting rank multiplication of 1-to-2, 1-to-4, and/or 1-to-8). The 3DS register addresses shortcoming (a) of the JEDEC RDIMM and LRDIMM listed above.

Another component of the proposed 3DS-DIMM is a 3DS dynamic random-access memory (DRAM) package **720**. The 3DS DRAM package **720** includes a plurality of stacked DDR DRAM chips **724** or array dies. The 3DS DRAM package **720** has a data I/O load that is equivalent to the data I/O load of a single-die DRAM package, regardless of the actual number of DRAM dies in the 3DS DRAM package **720**. The 3DS DRAM package **720** may comprise a plurality of array dies (e.g., 2, 4 or 8 DDR DRAM dies) and a controller die **722**. The controller die **722** may include data buffers, a data path timing controller, a secondary command/address buffer, a programmable secondary 1-to-2 rank decoder, and a data path signal (e.g., DDT, DDT, RTT) controller. Examples of such memory packages include HMC. The 3DS DRAM package **720** addresses the shortcomings (b) and (c) of the JEDEC RDIMM and LRDIMM listed above. However, there are deficiencies in the proposed 3DS-DIMM standard as described below.

The JEDEC DDR3 RDIMM standard contains two major components: a DDR3 register and a plurality of DRAM packages each comprising one or more DRAM chips or dies. The DDR3 register serves as a command/address signal buffer and as a command/control decoder. This DDR3 register holds

a set of register control (RC) words, which a computer system configures to ensure the proper operation of the RDIMM. The DDR3 register contains a phase-lock-loop (PLL), which functions as a clock synchronizer for each RDIMM. The DDR3 register outputs a set of buffered command/address signals to all the DRAM packages on the RDIMM, but the data signals are directly fed to the DRAM packages from the system memory controller.

In contrast to the JEDEC DDR3 RDIMM standard, the 3DS-DIMM proposal requires the DRAM package **720** to include the controller die **722** that controls all data paths timing and operations. This arrangement of the DRAM package reduces the load on the data bus, however, it presents three significant shortcomings: 1) The command/address buffer in the 3DS DRAM package introduces clock cycle latency since it needs to provide clock synchronous operation to the DRAM dies in the package. 2) The data path control circuit (e.g., DDT, read/write data direction) in the control die becomes very complicated to support semi-synchronous (flyby) data path control among all 3DS DRAM packages that are on a DIMM. 3) The variations in the DRAM die timing characteristics within each 3DS DRAM package would be likely to require resynchronization of the data signals during the read/write operations, which increases the read/write latency and the read-to-write and write-to-read turn around time.

The 3DS-DIMM proposal can also be compared to the HyperCloud™ (HC) DIMM architecture of Netlist, Inc. An example of the HCDIMM architecture **602** is illustrated in FIG. **68**. Further details and embodiments of the HCDIMM architecture is disclosed in U.S. Pat. Nos. 7,289,386, 7,532,537, 7,619,912, and 7,636,274, each of which is incorporated in its entirety by reference herein. One of the main topological differences between the 3DS-DIMM and HCDIMM architectures is that while the 3DS-DIMM architecture uses a control die **722** to buffer the data signals and to decode command/address signals from the 3DS register, certain configurations of the HCDIMM architecture include a plurality of isolation devices (ID), each of which includes data buffers, but no decoding capability. Unlike the HCDIMM architecture, since the command/address signal needs to pass through the control die **722** in the 3DS DRAM package **720**, the 3DSDIMM proposal presents the same shortcoming (b) of the controller die described above. The data path control signals are generated by the register device (RD) **612** in the HCDIMM architecture **602**, while the data path control signals are generated by the control die **722** in the 3DS DRAM. This aspect of the 3DS-DIMM architecture creates timing critical control paths in 3DS-DIMM architecture.

In certain embodiments described herein, a memory module architecture is proposed that includes a set of device components that support JEDEC 3DS operation with the benefit of RDIMM and HCDIMM architectures (see, e.g., FIG. **8**). This set of device components may comprise two components: a register device **812** (RD), which in some embodiments may be the same or similar RD component as used in HCDIMM architectures, and a plurality of array dies **824**, such as a DDR DRAM Stack Package (DDSP). The DDSP may comprise a DRAM control die that can include command/address buffers and a data path control circuit. Certain embodiments of the architecture described herein differs from the 3DS-DIMM architecture in that instead of the controller die of the 3DS-DIMM (which provides a secondary address buffer, and a second rank multiplication decoder), certain embodiments described herein use an "ID+" die as a control die **822**, which can provide both selective isolation and address pass-through. Selective isolation refers generally to a driver corresponding to an array die driving a signal to the array die in response to a corresponding chip select signal while additional drivers corresponding to additional array dies maintain a previous state (e.g. do not drive the signal). For example, assuming that the memory package **300** implements selective isolation, if a chip select signal is received that corresponds to a selection of array die **310***b*, then the driver **334***a* will drive a signal along the die interconnect **320***a* to the array die **310***b*, and the remaining drivers (e.g., drivers **334***b*-**334***f*) will maintain their state (e.g., not drive the signal). Address pass-through is described in further detail below.

The data path control circuit of certain embodiments may have at least two operational modes: mode-C(HCDIMM/RDIMM Compatible mode) and mode-3DS (3DS DIMM compatible mode). In mode-C, the array dies **824** (e.g. a DDSP) receive the data path control signals from the register device **812**, which can be configured to control a command/address time slot and a data bus time slot. In mode-3DS, the control die **822** internally generates the data path control signals to ensure the proper operation of the data path.

The control die **822** of certain embodiments enables the memory module **800** to work as either a 3DS-DIMM, a RDIMM, or a HCDIMM. However, the memory module **800** may comprise a set of optional control input pins (in addition to the package pins that are included in 3DS-DRAM packages) which in mode-C receive the data path control signals from the register device **812**.

As previously mentioned, FIG. **8** schematically illustrates an example embodiment of a memory module architecture **800** in accordance with the present disclosure. Advantageously, certain embodiments of the memory module architecture **800** address the shortcomings described above without adding complexity or latency, and without causing performance loss. The memory module architecture **800** includes a memory control hub **802** (also known as a memory controller hub, a memory control handler, or a northbridge) and a memory module **810**. As schematically illustrated in FIG. **8**, the memory control hub **802** may communicate directly with one or more components of the memory module **810**. Alternatively, the memory control hub **802** may communicate with one or more intermediary system components (not shown), which in turn communicate with one or more components of the memory module **810**. In some embodiments, the memory control hub **802** may be integrated with another component of the computer system, such as a Central Processing Unit (CPU).

Further, in some embodiments, the memory control hub **802** may communicate with one or more memory modules. Each of the memory modules may be similar or substantially similar to the memory module **810**. Alternatively, some of the memory modules may differ from memory module **810** in configuration, type, or both. For example, some of the memory modules may be capable of operating at different frequencies, may include a different amount of storage capability, or may be configured for different purposes (e.g. graphics versus nongraphics memory). Although in some cases a system may include memory modules capable of operating at different frequencies, each memory module may be configured to operate at the same frequency when used in a specific device. In certain implementations, the memory control hub **802** may set the operating frequency for the one or more memory modules.

The memory module **810** may include a register device **812**, which is configured to receive command, address, or command and address signals from the memory control hub **802**. For the purpose of simplifying discussion, and not to limit these signals, the signals will be referred to herein as command/address signals.

In some embodiments, the register device **812** receives the command/address signals via a command/address bus **814**. The command/address bus **814**, although illustrated as a single line, may include as many lines or signal conduits as the number of bits of the command/address signal.

Further, the register device **812** may generate data path control signals, which can be provided to a control die **822**, or isolation device, of a memory package **820** via one or more data path control lines **816**. In certain embodiments, the control dies **822** can include some or all of the embodiments described above with respect to the control dies **230** and **330**. Moreover, in certain embodiments, the memory packages **820** can include some or all of the embodiments described above with respect to the memory package **200** and **300**. In general, the memory module **812** may include one or more memory packages **820**.

In certain embodiments, each control die **822**, or isolation device, may be capable of address pass-through. Address pass-through, in some cases, enables the control die **822** to provide an address signal to one or more array dies **824** without decoding the address signal. This is possible, in some implementations, because the address signal received by the control die **822** is not encoded.

Some implementations of the control dies **822** include a plurality of command/address buffers (not shown). These buffers may comprise latches. In certain embodiments, the buffers are configured to hold command/address signals to control the timing of command/address signals. In some cases, controlling the timing of the command/address signals may reduce or slow signal degradation. In some implementations, the control dies **822** include a plurality of data buffers, which may control the timing of data signals to reduce or slow signal degradation. Further, the control dies **822** may include a data path control circuit (not shown) that is configured to control command/address time slots and data bus time slots. Controlling the command/address time slots and the data bus time slots enables the control dies **822** to reduce or prevent signal collisions caused by multiple memory packages **820** sharing the data path control lines **816** and the data bus **818**. In some implementations, the data path control circuit may be separate from the control die **822**.

Each of the control dies **822** may be configured to receive data signals from the memory control hub **802** via the data bus **818**. Further, the control dies **822** may provide data signals to the memory control hub **802** via the data bus **818**. The control dies **822** may also receive data path control signals and/or command/address signals from the register device **812** via the data path control lines **816**.

As illustrated in FIG. **8**, each memory package **820** may include one or more array dies **824** operatively coupled to the control die **822**. The array dies **824** may be configured to receive data signals from the control die **822**. As with the array dies **210** and **310**, the array dies **824** may include any type of Random Access Memory (RAM) die. For example, the array dies **824** may include DDR DRAM, SDRAM, flash memory, or SRAM, to name a few. Further, if the memory module **810** is utilized for graphics, the array dies **824** may include GDDR SDRAM or any other type of graphics memory. In addition, the array dies **824** may be configured in a stack arrangement as illustrated in FIG. **8**. Alternatively, the array dies **824** may be arranged in a planar configuration or a hybrid configuration utilizing both a planar and a stack arrangement.

In some embodiments, the memory module **810** is selectively configurable into two operational modes. In the first operational mode, the register device **812** generates data path control signals provided to the control dies **822** via the data path control lines **816**. The control dies **822** may decode command/address signals included with the data path control signals generated by the register device **812**. In some implementations, the control dies **822** use the data path control signals to operate the data path control circuits of the control dies **822**.

In the second operational mode, the control dies **822** may operate the data path control circuits to provide command/address signals to the array dies **824** without decoding the command/address signals. In this mode, the control dies **822** may use address pass-through to provide received address signals to the array dies **824**.

Other operational modes may also be possible. In some embodiments, the data path control signals generated by the register device **812** may include decoded command/address signals that are decoded from command/address signals received from the memory control hub **802** via the command/address bus **814**.

In some embodiments, the register device **812** may be configured to perform rank multiplication. In addition, or alternatively, the control dies **822** may be configured to perform rank multiplication. Embodiments of rank multiplication are described in greater detail in U.S. Pat. Nos. 7,289,386 and 7,532,537, each of which are incorporated in their entirety by reference herein. In such embodiments, the register device **812** can generate additional chip select signals that are provided to the array dies **824**. For instance, if the memory control hub **802** is configured to recognize a single array die **824** per memory package **820**, but there exists four array dies **824** per memory package **820**, the memory control hub **802** may not provide the correct number of chip select signals to access a specific memory location of the plurality of array dies **824** is memory package **820**. Thus, to access the specific memory location, the register device **812** can determine the array die that includes the specific memory location to be accessed based on the command/address signals received from the memory control hub **802** and can generate the correct chip select signal to access the array die that includes the specific memory location. In certain embodiments, when the memory module **810** is operating in the second operation module as described above, the memory module **810** does not perform rank multiplication.

TERMINOLOGY

Unless the context clearly requires otherwise, throughout the description and the claims, the words "comprise," "comprising," and the like are to be construed in an inclusive sense, as opposed to an exclusive or exhaustive sense; that is to say, in the sense of "including, but not limited to." The term "coupled" is used to refer to the connection between two elements, the term refers to two or more elements that may be either directly connected, or connected by way of one or more intermediate elements. Additionally, the words "herein," "above," "below," and words of similar import, when used in this application, shall refer to this application as a whole and not to any particular portions of this application. Where the context permits, words in the above Detailed Description using the singular or plural number may also include the plural or singular number respectively. The word "or" in reference to a list of two or more items covers all of the following interpretations of the word: any of the items in the list, all of the items in the list, and any combination of the items in the list.

The above detailed description of embodiments of the invention is not intended to be exhaustive or to limit the invention to the precise form disclosed above. While specific

embodiments of, and examples for, the invention are described above for illustrative purposes, various equivalent modifications are possible within the scope of the invention, as those skilled in the relevant art will recognize. For example, while processes or blocks are presented in a given order, alternative embodiments may perform routines having steps, or employ systems having blocks, in a different order, and some processes or blocks may be deleted, moved, added, subdivided, combined, and/or modified. Each of these processes or blocks may be implemented in a variety of different ways. Also, while processes or blocks are at times shown as being performed in series, these processes or blocks may instead be performed in parallel, or may be performed at different times.

The teachings of the invention provided herein can be applied to other systems, not necessarily the system described above. The elements and acts of the various embodiments described above can be combined to provide further embodiments.

Conditional language used herein, such as, among others, "can," "might," "may," "e.g.," and the like, unless specifically stated otherwise, or otherwise understood within the context as used, is generally intended to convey that certain embodiments include, while other embodiments do not include, certain features, elements and/or states. Thus, such conditional language is not generally intended to imply that features, elements and/or states are in any way required for one or more embodiments or that one or more embodiments necessarily include logic for deciding, with or without author input or prompting, whether these features, elements and/or states are included or are to be performed in any particular embodiment.

While certain embodiments of the inventions have been described, these embodiments have been presented by way of example only, and are not intended to limit the scope of the disclosure. Indeed, the novel methods and systems described herein may be embodied in a variety of other forms; furthermore, various omissions, substitutions and changes in the form of the methods and systems described herein may be made without departing from the spirit of the disclosure. The accompanying claims and their equivalents are intended to cover such forms or modifications as would fall within the scope and spirit of the disclosure.

What is claimed is:

1. A memory package, comprising:

data terminals and control terminals;

stacked array dies including a first group of array dies and a second group of at least one array die;

first die interconnects and second die interconnects, the first die interconnects in electrical communication with the first group of array dies and not in electrical communication with the second group of at least one array die, the second die interconnects in electrical communication with the second group of at least one array die and not in electrical communication with the first group of array dies; and

a control die comprising first data conduits between the first die interconnects and the data terminals, and second data conduits between the second die interconnects and the data terminals, the first data conduit including first drivers each having a first driver size and configured to drive a data signal from a corresponding data terminal to the first group of array dies, the second data conduit including second drivers each having a second driver size and configured to drive a data signal from a corresponding data terminal to the second group of at least one array die, the second driver size being different from the first driver size.

2. The memory package of claim **1**, wherein the second die interconnects are longer than the first die interconnects, and wherein the second driver size is larger than the first driver size.

3. The memory package of claim **1**, wherein the second die interconnects are longer than the first die interconnects, and wherein a number of array dies in the second group of at least one array die is less than a number of array dies in the first group of array dies.

4. The memory package of claim **1**, wherein the first driver size and the second driver size are related to a load on the first driver and a load on the second driver.

5. The memory package of claim **1**, wherein the control die further comprises a control circuit to control respective states of the first data conduits and the second data conduits in response to control signals received via the control terminals.

6. A memory package, comprising:

data terminals and control terminals;

stacked array dies including a first group of array dies and a second group of at least one array die;

first die interconnects and second die interconnects, the first die interconnects in electrical communication with the first group of array dies and not in electrical communication with the second group of at least one array die, the second die interconnects in electrical communication with the second group of at least one array die and not in electrical communication with the first group of array dies; and

wherein the second die interconnects are longer than the first die interconnects, and wherein a number of array dies in the second group of at least one array die is less than a number of array dies in the first group of array dies.

7. The memory package of claim **6**, further comprising:

a control die comprising first data conduits between the first die interconnects and the data terminals, and second data conduits between the second die interconnects and the data terminals, the first data conduits including first drivers each having a first driver size and configured to drive a data signal from a corresponding data terminal to the first group of array dies, the second data conduit including second drivers each having a second driver size and configured to drive a data signal from a corresponding data terminal to the second group of at least one array die, the second driver size being different from the first driver size.

8. The memory package of claim **7**, wherein the second driver size is larger than the first driver size.

9. The memory package of claim **7**, wherein the control die further comprises a control circuit to control respective states of the first data conduits and the second data conduits in response to control signals received via one or more second terminals of the plurality of terminals.

10. A memory module operable in a computer system with a system memory controller, comprising:

a register device configured to receive input command/address signals from the system memory controller and to output control signals; and

a plurality of DRAM packages, each DRAM package comprising:

data terminals via which the DRAM package communicate data signals with the system memory controller, and control terminals via which the DRAM package receive the control signals from the register device;

stacked array dies including a first group of array dies and a second group of at least one array die;

die interconnects including first die interconnects and second die interconnects, the first die interconnects in electrical communication with the first group of array dies and not in electrical communication with the second group of at least one array die, the second die interconnects in electrical communication with the second group of at least one array die and not in electrical communication with the first group of array dies; and

a control die including first data conduits between the first die interconnects and data terminals, and second data conduits between the second die interconnects and the data terminals, the first data conduits including first drivers each having a first driver size, the second data conduits including second drivers each having a second driver size different from the first driver size, the first drivers to drive first write data signals received from the system memory controller to one of the first group of array dies, the second drivers to drive second write data signals received from the system memory controller to one of the second group of at least one array die.

11. The memory module of claim **10**, wherein the control die receives the control signals and further includes a control circuit to control respective states of the first data conduits and the second data conduits in response to the control signals.

12. The memory module of claim **11**, wherein the control signals include data path control signals generated by the register device, the data path control signals being used to control the respective states of the first data conduits and the second data conduits.

13. The memory module of claim **11**, wherein the control die is configured to generate data path control signals from at least some of the control signals, the data path control signals being used to control the respective states of the first data conduits and the second data conduits.

14. The memory module of claim **10**, wherein the control signals include address signals and the control die provides the address signals to the plurality of array dies.

15. The memory module of claim **10**, wherein the input command/address signals include first chip select signals, wherein the register device is configured to perform rank multiplication by generating second chip select signals from at least some of the input command/address signals, the second chip select signals having a number of chip select signals greater than the first chip select signals and equal to a number of array dies in the plurality of array dies, and wherein the DRAM package further comprises chip select die interconnects for conducting the second chip select signals to respective ones of the plurality of array dies.

16. The memory module of claim **10**, wherein the control signals include output command/address signals derived from the input command/address signals, the output command/address signals including first chip select signals, wherein the control die is further configured to perform rank multiplication by generating second chip select signals from at least some of the output command/address signals, the second chip select signals having a number of chip select signals greater than the first chip select signals and equal to a number of array dies in the plurality of array dies, and wherein the DRAM package further comprises chip select die interconnects for conducting the second chip select signals to respective ones of the plurality of array dies.

17. The memory module of claim **10**, wherein the first driver size and the second driver size are related to a first load on the first driver and a second load on the second driver.

18. The memory module of claim **10**, wherein the control signals include command/address signals, and the control die includes buffers to control the timing of the command/address signals.

19. The memory module of claim **10**, wherein the control die includes data buffers to control the timing of the data signals.

20. The memory module of claim **10**, wherein the first group of array dies include a greater number of array dies than the second group of at least one array die.

* * * * *

CERTIFICATE OF SERVICE AND FILING

I certify that on December 16, 2024, I electronically filed the foregoing **NON-CONFIDENTIAL PRINCIPAL BRIEF** of Defendants-Appellants using the Court's CM/ECF filing system. Counsel for Plaintiff-Appellee Netlist, Inc. were electronically served by and through the Court's CM/ECF filing system per Fed. R. App. P. 25 and Fed. Cir. R. 25(e).

/s/ Ruffin B. Cordell
Ruffin B. Cordell

Samsung Electronics Co., Ltd.
Ex. 1077, p. 832

CERTIFICATE OF COMPLIANCE

The **NON-CONFIDENTIAL PRINCIPAL BRIEF** of Defendants-Appellants is submitted in accordance with the type-volume limitation of Fed. Cir. R. 32(b). The brief contains 13,844 words, excluding the parts of the brief exempted by Fed. R. App. P. 32(f) and Fed. Cir. R. 32(b)(2). This brief has been prepared in a proportionally spaced typeface using Microsoft® Word for Microsoft 365 in Times New Roman, 14 Point.

Dated: December 16, 2024

/s/ Ruffin B. Cordell
Ruffin B. Cordell

Samsung Electronics Co., Ltd.
Ex. 1077, p. 833

CERTIFICATE OF CONFIDENTIAL INFORMATION

The **NON-CONFIDENTIAL PRINCIPAL BRIEF** of Defendants-Appellants complies with Federal Circuit Rule 25.1(d)'s limitations on the amount of material that can be marked confidential in motions, responses, replies, or briefs. This brief contains 0 unique words (including numbers) marked confidential, which does not exceed the maximum of 15 words permitted by Fed. Cir. R. 25.1(d)(1)(A) & (C). The Addendum includes the district court opinion which was filed under seal before the district court.

Dated: December 16, 2024

/s/ Ruffin B. Cordell
Ruffin B. Cordell

Samsung Electronics Co., Ltd.
Ex. 1077, p. 834