

Nos. 24-2240, 24-2241

---

IN THE  
**United States Court of Appeals**  
FOR THE FEDERAL CIRCUIT

NETLIST, INC.,  
*Appellant,*

v.

SAMSUNG ELECTRONICS CO., LTD., MICRON TECHNOLOGY, INC.,  
MICRON SEMICONDUCTOR PRODUCTS, INC.,  
MICRON TECHNOLOGY TEXAS, LLC,  
*Appellees.*

On Appeal from the U.S. Patent and Trademark Office,  
Patent Trial and Appeal Board,  
Nos. IPR2022-01427, IPR2022-01428,  
IPR2023-00882, IPR2023-00883

---

**OPENING BRIEF FOR APPELLANT NETLIST, INC.**

---

Elizabeth Kathleen Clarke  
Bonnie K. St. Charles  
MOLOLAMKEN LLP  
300 N. LaSalle Street, Suite 5350  
Chicago, IL 60654  
(312) 450-6700

Sara Margolis  
Catherine Martinez  
MOLOLAMKEN LLP  
430 Park Avenue  
New York, NY 10022  
(212) 607-8160

Jeffrey A. Lamken  
Rayiner Hashem  
Jennifer Elizabeth Fischell  
Lidiya Mishchenko  
Kayvon Ghayoumi  
MOLOLAMKEN LLP  
The Watergate, Suite 500  
600 New Hampshire Avenue, N.W.  
Washington, D.C. 20037  
(202) 556-2000  
jlamken@mololamken.com

*Counsel for Appellant Netlist, Inc.*

## **REPRESENTATIVE PATENT CLAIMS**

### **U.S. Patent No. 8,787,060**

1.
  - [a] A memory package, comprising:
  - [b] a plurality of input/output terminals via which the memory package communicates data and control/address signals with one or more external devices;
  - [c] 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;
  - [d] 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
  - [e] 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.
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.

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.

**U.S. Patent No. 9,318,160**

1. [a] A memory package, comprising:
- [b] data terminals and control terminals;
- [c] stacked array dies including a first group of array dies and a second group of at least one array die;
- [d] 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
- [e] 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.

**UNITED STATES COURT OF APPEALS  
FOR THE FEDERAL CIRCUIT****CERTIFICATE OF INTEREST****Case Number** 24-2240, 24-2241**Short Case Caption** Netlist, Inc. v. Samsung Electronics Co., Ltd.**Filing Party/Entity** Netlist, Inc.**Instructions:**

1. Complete each section of the form and select none or N/A if appropriate.
2. Please enter only one item per box; attach additional pages as needed, and check the box to indicate such pages are attached.
3. In answering Sections 2 and 3, be specific as to which represented entities the answers apply; lack of specificity may result in non-compliance.
4. Please do not duplicate entries within Section 5.
5. Counsel must file an amended Certificate of Interest within seven days after any information on this form changes. Fed. Cir. R. 47.4(c).

I certify the following information and any attached sheets are accurate and complete to the best of my knowledge.

Date: 02/04/2025Signature: /s/ Jeffrey A. LamkenName: Jeffrey A. Lamken

| <b>1. Represented Entities.</b><br>Fed. Cir. R. 47.4(a)(1).                             | <b>2. Real Party in Interest.</b><br>Fed. Cir. R. 47.4(a)(2).                                                                                                                                              | <b>3. Parent Corporations and Stockholders.</b><br>Fed. Cir. R. 47.4(a)(3).                                                                                                                                   |
|-----------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Provide the full names of all entities represented by undersigned counsel in this case. | 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.<br><br><input checked="" type="checkbox"/> None/Not Applicable | 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.<br><br><input checked="" type="checkbox"/> None/Not Applicable |
| Netlist, Inc.                                                                           |                                                                                                                                                                                                            |                                                                                                                                                                                                               |
|                                                                                         |                                                                                                                                                                                                            |                                                                                                                                                                                                               |
|                                                                                         |                                                                                                                                                                                                            |                                                                                                                                                                                                               |
|                                                                                         |                                                                                                                                                                                                            |                                                                                                                                                                                                               |
|                                                                                         |                                                                                                                                                                                                            |                                                                                                                                                                                                               |
|                                                                                         |                                                                                                                                                                                                            |                                                                                                                                                                                                               |
|                                                                                         |                                                                                                                                                                                                            |                                                                                                                                                                                                               |
|                                                                                         |                                                                                                                                                                                                            |                                                                                                                                                                                                               |
|                                                                                         |                                                                                                                                                                                                            |                                                                                                                                                                                                               |
|                                                                                         |                                                                                                                                                                                                            |                                                                                                                                                                                                               |
|                                                                                         |                                                                                                                                                                                                            |                                                                                                                                                                                                               |

☐ Additional pages attached

**4. Legal Representatives.** 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).

☐ None/Not Applicable

☐ Additional pages attached

|                     |                   |               |
|---------------------|-------------------|---------------|
| Irell & Manella LLP | Hong Annita Zhong | Jason Sheasby |
| Blair A. Silver     |                   |               |
|                     |                   |               |

**5. Related Cases.** 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 (file separate notice; see below) ☐ No ☐ N/A (amicus/movant)

If yes, concurrently file a separate Notice of Related Case Information that complies with Fed. Cir. R. 47.5(b). **Please do not duplicate information.** This separate Notice must only be filed with the first Certificate of Interest or, subsequently, if information changes during the pendency of the appeal. Fed. Cir. R. 47.5(b).

**6. Organizational Victims and Bankruptcy Cases.** 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

☐ Additional pages attached

|  |  |  |
|--|--|--|
|  |  |  |
|  |  |  |

## **TABLE OF CONTENTS**

|                                                                                                                                        | <u>Page</u> |
|----------------------------------------------------------------------------------------------------------------------------------------|-------------|
| JURISDICTIONAL STATEMENT .....                                                                                                         | 1           |
| STATEMENT OF ISSUES .....                                                                                                              | 1           |
| STATEMENT OF THE CASE.....                                                                                                             | 1           |
| I. Technological Background.....                                                                                                       | 1           |
| A. Conventional Server Memory Modules .....                                                                                            | 2           |
| B. HBM/3DS Memory and the Problem of Load.....                                                                                         | 4           |
| C. Netlist Invents a New Architecture To Reduce Load .....                                                                             | 6           |
| D. Netlist Is Awarded the '060 and '160 Patents.....                                                                                   | 12          |
| II. Procedural History .....                                                                                                           | 14          |
| A. Asserted Prior Art.....                                                                                                             | 14          |
| 1. Kim.....                                                                                                                            | 15          |
| 2. Rajan .....                                                                                                                         | 21          |
| 3. Riho .....                                                                                                                          | 22          |
| 4. Wyman .....                                                                                                                         | 22          |
| B. The Final Written Decisions .....                                                                                                   | 23          |
| 1. '060 Patent Claim 1: Electrically Connecting Die<br>Interconnects to Array Die Data Ports in Only One Group<br>(Kim and Rajan)..... | 23          |
| 2. '060 Patent Claim 7: Considering Die-Interconnect Load<br>and Reducing Load Differences (Kim, Rajan, and Riho).....                 | 28          |

|     |                                                                                                                                                                                    |    |
|-----|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 3.  | '060 Patent Claim 15: Differently Sized Drivers (Kim, Rajan, and Wyman) .....                                                                                                      | 30 |
|     | SUMMARY OF ARGUMENT .....                                                                                                                                                          | 32 |
|     | STANDARD OF REVIEW .....                                                                                                                                                           | 34 |
|     | ARGUMENT .....                                                                                                                                                                     | 34 |
| I.  | The Board's Determination that the Prior Art Teaches a First and Second Group of Array Dies, Each Group Sharing Separate Interconnects as in the Claims, Cannot Be Sustained ..... | 34 |
| A.  | The Board's Motivation To Modify Kim—By Connecting Multiple Array Dies to Each TSV—Defies Precedent and the APA .....                                                              | 37 |
| 1.  | The Board Legally Erred in Overlooking that Kim Taught Away from Connecting Multiple Array Dies to Each TSV .....                                                                  | 38 |
| 2.  | The Board's Analysis Defies Precedent by Proposing Modifications that Eviscerate Kim's Invention and Purpose .....                                                                 | 43 |
| 3.  | The Board's Generic Assertion of Motivation Was Insufficient .....                                                                                                                 | 47 |
| 4.  | The Board's Effort To Downplay Data Collisions Is No Answer .....                                                                                                                  | 49 |
| B.  | The Board's Expectation-of-Success Analysis Was Deficient .....                                                                                                                    | 51 |
| II. | The Board's Determination that Riho Teaches Claim 7 of the '060 Patent Cannot Be Sustained .....                                                                                   | 55 |
| A.  | The Board Misconstrued Claim 7 .....                                                                                                                                               | 56 |
| B.  | Reversal Is Warranted Because Riho Does Not Teach Claim 7 .....                                                                                                                    | 60 |
| C.  | The Board's Determination that Skilled Artisans Would Have Combined Riho with Kim Is Unreasoned and Unsupported .....                                                              | 62 |

|                  |                                                                                      |    |
|------------------|--------------------------------------------------------------------------------------|----|
| III.             | The Board’s Decision on Claim 15 of the ’060 Patent Cannot Be Sustained.....         | 64 |
| IV.              | The Board’s Unreasoned Analysis of Other Claims and Limitations Requires Remand..... | 68 |
| CONCLUSION ..... |                                                                                      | 69 |

## TABLE OF AUTHORITIES

| CASES                                                                                                      | Page(s)       |
|------------------------------------------------------------------------------------------------------------|---------------|
| <i>ActiveVideo Nets., Inc. v. Verizon Commc'ns, Inc.</i> ,<br>694 F.3d 1312 (Fed. Cir. 2012) .....         | 47, 48        |
| <i>AstraZeneca AB v. Mylan Pharms. Inc.</i> ,<br>19 F.4th 1325 (Fed. Cir. 2021) .....                      | 38, 42        |
| <i>Auris Health, Inc. v. Intuitive Surgical Operations, Inc.</i> ,<br>32 F.4th 1154 (Fed. Cir. 2022) ..... | 50            |
| <i>Belden Inc. v. Berk-Tek LLC</i> ,<br>805 F.3d 1064 (Fed. Cir. 2015) .....                               | 48            |
| <i>Chemours Co. FC, LLC v. Daikin Indus., Ltd.</i> ,<br>4 F.4th 1370 (Fed. Cir. 2021) .....                | <i>passim</i> |
| <i>Cont'l Circuits LLC v. Intel Corp.</i> ,<br>915 F.3d 788 (Fed. Cir. 2019) .....                         | 57, 58        |
| <i>Corephotonics, Ltd. v. Apple Inc.</i> ,<br>84 F.4th 990 (Fed. Cir. 2023) .....                          | 53            |
| <i>Del. Dep't of Natural Res. &amp; Env't Control v. EPA</i> ,<br>785 F.3d 1 (D.C. Cir. 2015) .....        | 64            |
| <i>DePuy Spine, Inc. v. Medtronic Sofamor Danek, Inc.</i> ,<br>567 F.3d 1314 (Fed. Cir. 2009) .....        | 26, 38, 42    |
| <i>Endo Pharms. Inc. v. Actavis LLC</i> ,<br>922 F.3d 1365 (Fed. Cir. 2019) .....                          | 43            |
| <i>Fanduel, Inc. v. Interactive Games LLC</i> ,<br>966 F.3d 1334 (Fed. Cir. 2020) .....                    | 68            |
| <i>In re Fritch</i> ,<br>972 F.2d 1260 (Fed. Cir. 1992) .....                                              | 48, 49        |
| <i>Gen. Elec. Co. v. Raytheon Techs. Corp.</i> ,<br>983 F.3d 1334 (Fed. Cir. 2020) .....                   | 50            |

|                                                                                                         |               |
|---------------------------------------------------------------------------------------------------------|---------------|
| <i>Google LLC v. EcoFactor, Inc.</i> ,<br>92 F.4th 1049 (Fed. Cir. 2024) .....                          | 56, 59        |
| <i>In re Gordon</i> ,<br>733 F.2d 900 (Fed. Cir. 1984) .....                                            | 43            |
| <i>Henny Penny Corp. v. Frymaster LLC</i> ,<br>938 F.3d 1324 (Fed. Cir. 2019) .....                     | 67            |
| <i>In re Hodges</i> ,<br>882 F.3d 1107 (Fed. Cir. 2018) .....                                           | 61, 62        |
| <i>HTC Corp. v. Cellular Commc'ns Equip., LLC</i> ,<br>877 F.3d 1361 (Fed. Cir. 2017) .....             | 34, 56        |
| <i>Icon Health &amp; Fitness, Inc. v. Strava, Inc.</i> ,<br>849 F.3d 1034 (Fed. Cir. 2017) .....        | 63, 68, 69    |
| <i>Intel Corp. v. Koninklijke Philips, N.V.</i> ,<br>No. 22-2034, 2024 WL 725243 (Fed. Cir. 2024) ..... | 46            |
| <i>KSR Int'l Co. v. Teleflex Inc.</i> ,<br>550 U.S. 398 (2007) .....                                    | 65            |
| <i>In re Lee</i> ,<br>277 F.3d 1338 (Fed. Cir. 2002) .....                                              | 37            |
| <i>Life Techs, Inc. v. Clonetech Labs., Inc.</i> ,<br>224 F.3d 1320 (Fed. Cir. 2000) .....              | 54            |
| <i>Mars, Inc. v. Coin Acceptors, Inc.</i> ,<br>514 F. Supp. 2d 624 (D.N.J. 2007) .....                  | 39            |
| <i>Medtronic, Inc. v. Teleflex Innovations, S.à.r.l.</i> ,<br>69 F.4th 1341 (Fed. Cir. 2023) .....      | 43, 46, 47    |
| <i>Natera, Inc. v. NeoGenomics Labs., Inc.</i> ,<br>106 F.4th 1369 (Fed. Cir. 2024) .....               | 49            |
| <i>In re Nuvasive, Inc.</i> ,<br>842 F.3d 1376 (Fed. Cir. 2016) .....                                   | <i>passim</i> |

|                                                                                                         |               |
|---------------------------------------------------------------------------------------------------------|---------------|
| <i>Pers. Web Techs., LLC v. Apple, Inc.</i> ,<br>848 F.3d 987 (Fed. Cir. 2017) .....                    | 34, 62        |
| <i>Plas-Pak Indus., Inc. v. Sulzer Mixpac AG</i> ,<br>600 F. App'x 755 (Fed. Cir. 2015) .....           | 43, 46, 52    |
| <i>Polaris Indus., Inc. v. Arctic Cat, Inc.</i> ,<br>882 F.3d 1056 (Fed. Cir. 2018) .....               | 38, 50, 65    |
| <i>Princeton Vanguard, LLC v. Frito-Lay N. Am., Inc.</i> ,<br>786 F.3d 960 (Fed. Cir. 2015) .....       | 50            |
| <i>Provisur Techs., Inc. v. Weber, Inc.</i> ,<br>50 F.4th 117 (Fed. Cir. 2022) .....                    | 42, 43, 47    |
| <i>In re Rouffet</i> ,<br>149 F.3d 1350 (Fed. Cir. 1998) .....                                          | 47            |
| <i>TQ Delta, LLC v. CISCO Sys., Inc.</i> ,<br>942 F.3d 1352 (Fed. Cir. 2019) .....                      | 51, 62        |
| <i>In re Van Os</i> ,<br>844 F.3d 1359 (Fed. Cir. 2017) .....                                           | 37            |
| <i>VirnetX Inc. v. Cisco Sys., Inc.</i> ,<br>776 F. App'x 698 (Fed. Cir. 2019) .....                    | 63, 69        |
| <i>Virtek Vision Int'l ULC v. Assembly Guidance Sys., Inc.</i> ,<br>97 F.4th 882 (Fed. Cir. 2024) ..... | <i>passim</i> |
| <i>W.L. Gore &amp; Assocs., Inc. v. Garlock, Inc.</i> ,<br>721 F.2d 1540 (Fed. Cir. 1983) .....         | 67            |

## STATUTES AND RULES

|                                |    |
|--------------------------------|----|
| 28 U.S.C. § 1295(a)(4)(A)..... | 1  |
| 35 U.S.C. § 141(c) .....       | 1  |
| 35 U.S.C. § 314 .....          | 1  |
| 35 U.S.C. § 316(e).....        | 68 |
| 35 U.S.C. § 318(a).....        | 1  |

|                       |   |
|-----------------------|---|
| 35 U.S.C. § 319 ..... | 1 |
|-----------------------|---|

# **OTHER AUTHORITIES**

|                                                                   |        |
|-------------------------------------------------------------------|--------|
| <i>Webster's New International Dictionary</i> (2d ed. 1934) ..... | 39, 58 |
|-------------------------------------------------------------------|--------|

## **STATEMENT OF RELATED CASES**

No appeal in or from this proceeding was previously before this or any other appellate court. The following cases may directly affect or be directly affected by this Court's decision in this case: *Netlist, Inc. v. Samsung Electronics Co., Ltd., et al.*, No. 2:21-cv-00463 (E.D. Tex.); *Netlist, Inc. v. Micron Tech., Inc., et al.*, No. 2:22-cv-00203 (E.D. Tex.); and *Netlist, Inc. v. Samsung Electronics Co., Ltd.*, No. 24-2203 (Fed. Cir.) (appeal from No. 2:21-cv-00463 (E.D. Tex.)). In addition, on September 12, 2024, this Court ordered that this appeal and the following related appeals be treated as companion cases and assigned to the same merits panel: Nos. 2024-1707, 2024-1859, 2024-1863, and 2024-2203. ECF No. 15.

## **JURISDICTIONAL STATEMENT**

The Board had jurisdiction under 35 U.S.C. §§ 314, 318(a). It issued final written decisions on April 1, 2024. Appx1; Appx69. Director review was denied on June 17, 2024. Appx144-46. Netlist timely appealed on August 19, 2024. Appx11867-71; Appx1418-20. This Court has jurisdiction under 28 U.S.C. § 1295(a)(4)(A) and 35 U.S.C. §§ 141(c), 319.

## **STATEMENT OF ISSUES**

1. Whether the Board’s obviousness determination regarding claim 1 of the ’060 patent and similar claims contravenes the Court’s precedents and the APA’s substantial-evidence and reasoned-decisionmaking requirements.

2. Whether the Board’s obviousness determination regarding dependent claim 7 of the ’060 patent rests on an erroneous claim construction or contravenes the Court’s precedents and the APA.

3. Whether the Board’s obviousness determination regarding claim 15 of the ’060 patent and similar claims contravenes the Court’s precedents and the APA.

## **STATEMENT OF THE CASE**

### **I. TECHNOLOGICAL BACKGROUND**

“Cloud computing” and internet-based “artificial intelligence” created enormous demand for servers that can process large amounts of data. To perform computations involving vast amounts of data, servers require large quantities of

memory. But efforts to increase memory capacity while maintaining high speeds and acceptable power consumption confronted technical barriers.

To increase server memory capacity and speed, the industry began developing “high-bandwidth memory” (“HBM”) and “three-dimensional stack” (“3DS”) memory, which involve stacking multiple silicon memory “dies” into a single memory “package.” But early approaches to HBM and 3DS memory faced problems with capacitive load—a physical phenomenon that limits the number of memory dies that can be stacked, reduces performance, and results in excessive power usage.

Netlist is a pioneer in the field of memory technology. Since 2000, Netlist, a public company based in California, has developed, manufactured, and supplied cutting-edge memory products to companies such as IBM, Dell, HP, and Apple. As relevant here, Netlist developed an innovative architecture for HBM and 3DS memory that solved the load problem and enabled a larger number of memory dies to be stacked in a high-speed memory package with reasonable power consumption.

#### **A. Conventional Server Memory Modules**

Conventionally, server memory is contained in “memory modules,” like the one below. Appx10309. A memory module includes dynamic random-access memory (“DRAM”) packages and other circuitry mounted on a printed circuit board. Appx4595; Appx7409(221:2-5).



Appx10310. In the memory module above, eighteen memory packages (black rectangles) are organized into two rows. The module contains other chips and circuitry, along with an edge connector (gold strip at the bottom). The edge connector plugs into the server's motherboard, electrically connecting the memory module to the "memory controller" in the server's CPU. Appx10309-12.

Conventional memory packages include a single silicon "die" (or layer) with billions of memory "cells" that each store one data "bit" (a 0 or 1). Appx4391. That silicon die is placed in a plastic or ceramic package—the visible part of the memory package (black rectangles above)—and mounted on the circuit board. Appx4595; Appx4391.

The data-storage capacity of silicon memory dies is limited by their physical size and manufacturing technology. Conventionally, each memory package had a single memory die or layer. As a result, increasing memory capacity required either increasing the number of memory modules (above), or increasing the number of memory packages (black rectangles) on each module. Appx185(10:41-55).

## B. HBM/3DS Memory and the Problem of Load

1. Faced with ever-increasing demands for servers with more memory capacity, memory-module manufacturers sought to incorporate *multiple* memory dies (or layers) into each “memory package.” Appx181(1:23-30).




FIG. 1B

Appx173(Fig. 1B) (annotated). As shown above, manufacturers stacked memory dies (160, blue) into a three-dimensional structure and placed that stack into the memory package. Appx181(1:23-30). “[D]ie interconnects” (vertical black lines, 182 and 188)—for example, “through-silicon vias,” or “TSVs”— were punched through the silicon dies. Appx181(2:2-7); Appx183(5:51-54). The die inter-

connects transmitted command and data signals between the CPU's memory controller and individual memory dies. Appx181(1:35-49).

Figure 1B above shows a so-called “multi-drop” configuration. Appx6113; Appx6391(¶34); Appx6559-60(¶299). The stack includes four memory dies (160, blue) and one “control die” (170, red). Appx181(1:62-66). Die interconnects originate at the control die and pass through all memory dies to “enable electrical communication and data transfer” between the memory dies and control die. Appx183(5:41-45); Appx181(1:63-2:7). As each die interconnect passes through each memory die (layer), it connects to a “data port[ ]” on that die, creating a path for electrical communication between that die and the control die. Appx181(2:19-20).<sup>1</sup>

2. The control die includes “drivers” (triangles 184, 186) for each die interconnect. Appx181(2:4-11). Drivers push electrical signals through the interconnects to each memory die. Appx181(2:4-9); Appx7416(228:21-24). To do so, the driver must overcome “load”—a physical phenomenon that distorts high-frequency electrical signals on interconnects and limits how far they can travel along the interconnect. Load reduces signal quality and data speeds, and can cause data loss. Appx4493; Appx184(8:8-12); Appx5173-74; Appx5197. In the multi-drop

---

<sup>1</sup> Interconnects 174 carry “chip select” signals that go to only one die. Appx181(1:63-67).

configuration above, *each* die interconnect (182, 188) electrically communicates with the control die and *all* the memory dies (each layer). Each additional memory die that is connected to a die interconnect increases the load. Appx6512(¶229); Appx5173-74; Appx5197. Load also increases as the die interconnect itself becomes longer. Appx186(12:16-21); Appx187(13:7-18).

The multi-drop structure, in which each die interconnect is relatively *longer* and in electrical communication with *every* memory die, results in high loads on drivers. To “overcome the load,” powerful drivers are required to drive signals to the memory dies. Appx181(2:8-15); Appx183(5:46-51). Sufficiently powerful drivers, however, occupy more “space on the control die,” “consume[ ] more power,” and produce more heat. Appx181(2:12-15); Appx4341.


### **C. Netlist Invents a New Architecture To Reduce Load**

Around 2009-2010, Netlist engineers developed an improved architecture for HBM and 3DS memory that solved the load problem with three key improvements.

1. To reduce the load on each driver, Netlist “reduc[ed] the number of” memory dies—“array dies” in Netlist’s invention—“that are in electrical communication with each die interconnect,” and “increas[ed] the number of die interconnects.” Appx182(4:22-29). Specifically, it divided the array dies into separate groups, and electrically connected each die interconnect *only* to data ports on array

dies in a particular group. Thus, die interconnects were no longer in communication with *all* the array dies, but only a *subset*.

In Figure 2, below, four array dies (210a-d, green/blue) are stacked on a control die (230, red).



Appx174(Fig. 2, annotated); Appx183(5:46-6:59). Filled-in circles show where an interconnect is electrically connected to an array die's data port, while unfilled circles show where an interconnect passes through a hole in the die. Appx183(5:41-45, 5:63-6:8). Here, one interconnect (220a, green) connects to data ports on *all* array dies in the first group (green), but *none* in the second group (blue). Appx183(6:9-15). Another interconnect (220b, blue) connects to data ports on all array dies in the second group (blue), but none in the first group (green).

Appx183(6:9-15). Thus, the load on each driver reflects the load of *two* array dies plus the load of one die interconnect, rather than all *four* array dies and corresponding die interconnects.<sup>2</sup>

Electrically connecting each die interconnect to data ports on only a subset of dies was counterintuitive. It required more die interconnects, which are expensive, are prone to manufacturing defects or failure, and require valuable space on memory dies for data ports or pass-through holes. Appx184(8:8-18). Increasing the number of interconnects also required additional drivers on the control die. Appx184(8:8-18). But reducing the number of array dies in electrical communication with each interconnect reduced the load on each driver, reducing “power consumption” and “increas[ing]” memory speed. Appx184(8:8-18).

There may be different numbers of array dies in each group. 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,” while a “different die interconnect” communicates with “at least *one* array die.” Appx160(5:55-65) (emphasis added). The number of array dies in each group thus can be selected to balance tradeoffs between reduced load and additional space occupied by die inter-

---

<sup>2</sup> In this embodiment, interconnects for data signals (220a, green, and 220b, blue)—which are especially sensitive to load—are connected to only a subset of array dies. Other interconnects, such as those for “chip select signals” (252a-d), may be configured differently. Appx185(9:28-60, 10:56-58); see p. 5 n.1, *supra*.

connects and drivers. That capability also enables balancing load across different drivers, as explained below.

2. Netlist's invention allows "balancing" the load on the drivers by selecting an appropriate number of array dies in each group to achieve a similar total load on each driver. Appx187(14:4-46). Each driver must overcome load that includes "the load of the array dies" *and* "the load of the die interconnect." Appx182(4:40-45). The load from array dies increases with the number of dies, while the load from the die interconnect increases with the interconnect's length. Appx6512(¶229); Appx184(7:9-21); Appx186(12:16-21). In the image on p. 7, above, interconnects 220a (green) and 220b (blue) are both connected to two array dies. Their corresponding drivers thus experience similar loads from those dies. Appx184(7:9-21). But interconnect 220b, connected to the upper group of dies (blue) is *longer* than interconnect 220a, connected to the lower group of dies (green). As a result, interconnect 220b's driver experiences a higher *total* load and must be more powerful to overcome that load. Appx187(14:22-39); Appx186(11:54-59); Appx186(12:39-51).

To "balance" loads across the drivers coupled to different die interconnects, the invention selects the "number of array dies" in communication with each interconnect based on the interconnect's "length," so as to "maintain the difference

between the load” on different drivers “below a threshold.” Appx184(7:30-43). The image below provides one example:




FIG. 3

Appx175(Fig. 3) (annotated). Driver 332a transmits signals on interconnect 320a, which is electrically connected to two array dies, 310a and 310b (all green). Similarly, driver 332b transmits signals on interconnect 320b, which also is connected to two array dies, 310c and 310d (all blue). And driver 332c transmits signals on interconnect 320c, which is connected to just one array die, 310e (all yellow). The load on the three drivers is thus roughly balanced: Drivers 332a and 332b (green and blue) have more load from their two array dies but less load from their shorter

interconnects, while driver 332c (yellow) has less load from its one array die but more load from its longer interconnect.

The specification teaches how to select the number of array dies in each group to balance load considering the corresponding die interconnect's length. Appx184(7:30-8:18). Table 1 (below) illustrates keeping load differences below a threshold where die interconnects are connected to two array dies, or one array die, depending on length (as measured by the number of "Interconnect Segments"):

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                           |

Appx187(13:7-18). Balancing load between drivers reduces power usage, allowing for smaller drivers while increasing speed, decreasing costs, and enhancing performance. Appx189(18:4-13).

3. Netlist's invention also allows tuning each driver's power to account for differences in load between die interconnects. It may not be possible or desirable to select the number of array dies in each group to produce "substantially equivalent"

load for all die interconnects. Appx187(14:5-20). The specification thus teaches using different “driver size[s]” for different die interconnects, “based . . . on the calculated load on [each] driver.” Appx189(17:35-50). Driver size and power can be tuned by adjusting the “size” and “number” of “transistors” in the driver. Appx189(17:36-50). Using different driver sizes for different die interconnects saves power, because a “larger driver” generally “consumes more power than a smaller driver.” Appx189(17:45-49). It also saves space on the control die, because “generally a larger driver . . . consumes more space.” Appx181(2:14-15).

#### **D. Netlist Is Awarded the '060 and '160 Patents**

Netlist obtained the '060 and '160 patents for its innovative architecture for reducing load in a HBM or 3DS memory. Appx171-94; Appx147-70.

Claim 1 of the '060 patent is representative. It recites:

- 1[a] A memory package, comprising:
- 1[b] a plurality of input/output terminals via which the memory package communicates data and control/address signals with one or more external devices;
- 1[c] 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;
- 1[d] 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

1[e] 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.

Appx192(23:58-24:15) (emphasis added).

Dependent claim 7 adds selecting the number of array dies in each group to balance load:

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.

Appx192(24:37-48) (emphasis added).

Dependent claim 15 depends from claim 11, which is materially the same as claim 1. It adds the requirement of using two sets of drivers of different sizes:

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.

Appx193(25:35-42) (emphasis added). Limitation [e] of claim 1 of the '160 patent is similar. Appx169(23:59-24:3).

## **II. PROCEDURAL HISTORY**

In 2015, four years after introducing its HBM concept, Netlist gave a presentation to Samsung regarding its '060 patent. Appx7408-09(220:17-23, 221:13-19); Appx7528-33; Appx6400(¶54). One year later, Samsung released memory modules incorporating that technology. *See, e.g.*, Appx7503. Netlist sued Samsung for infringement of the '060 and '160 patents. Appx5601-02.

Samsung then filed IPR petitions challenging all claims of both patents. Appx203; Appx11014. Micron filed substantively identical IPR petitions, and the Board joined Micron to Samsung's IPRs. Appx749-50; Appx1413-14. While the IPRs were pending, a jury found that Samsung willfully infringed both patents and awarded over \$100 million in damages. *Netlist Inc. v. Samsung Elecs. Co.*, No. 2:21-cv-463, Dkt. 479 at 4, 6 (E.D. Tex. Apr. 21, 2023).

### **A. Asserted Prior Art**


The Petition asserted four main prior-art references: “Kim,” “Rajan,” “Riho,” and “Wyman.”<sup>3</sup>

---

<sup>3</sup> The petitions are substantively similar. “The Petition” refers to Samsung’s petition for the '060 patent, unless otherwise noted.

# 1. Kim


Kim is a patent application (Pub. No. 2011/0103156) for a “memory apparatus capable of *sharing* a data input/output” or I/O “circuit.” Appx3628(¶ 12) (emphasis added). As shown below, conventional memory dies (green) comprise (a) memory cells that store data (purple), and (b) “write driver[s]” and “read sense amplifier[s]” (yellow) that write data to and read data from the memory cells. Appx3628(¶¶ 7-10); Appx3623.



Appx3623(Fig. 1) (annotated). Memory dies also contain I/O circuits (blue) that “amplif[y]” signals from the memory cells, “buffer” data, and “align[.]” data before transmitting that data to the memory controller across the “DQ” line (DQ0 or DQ1). Appx3628(¶ 10). Conventionally, each memory die contains an I/O circuit,


connected to the write driver and read sense amplifier via a “global input/output line” (GIO1 or GIO2). Appx3628(¶¶ 7-8).

Kim proposed removing the I/O circuit from each memory die and placing it in a separate control die, so multiple memory dies share the *same* I/O circuit. As shown below, Kim’s unique memory dies (100, 200, green) contain drivers and sense amplifiers, but no I/O circuit. The I/O circuitry (1200, blue) is instead placed on a control die, and “shared” between memory dies. Appx3629(¶ 14). Memory dies are connected to the control die through global I/O lines (GIO\_Rank0 and GIO\_Rank1), also called “data . . . lines.” Appx3629-30(¶ 28).



Appx3625(Fig. 3) (annotated).


Kim arranges its memory dies into a stack, as shown below. Kim discloses a memory device with two memory dies (C1 and C2), which Kim calls “ranks,” and a control die (C0). Appx3628(¶¶ 7-10). The memory dies are connected to the control die using TSVs. Appx3628(¶¶ 7-10). In Kim, the global I/O lines in some figures correspond to TSVs in other figures; the global I/O lines in the image above (GIO\_Rank0 and GIO\_Rank1) correspond to “TSV1” and “TSV2” in the image below. Appx3628(¶ 6).



Appx3627(Fig. 5). Kim teaches a *one-to-one correspondence* between data lines (TSVs) and memory dies (ranks): the “plurality of data . . . lines” are “*respectively connected* to the plurality of [memory] chips.” Appx3629(¶ 17) (emphasis added).

Kim explains that removing I/O circuits from each memory die, in favor of a shared I/O circuit, creates a potentially fatal “problem”—data collisions.


Appx3631(¶42). In Figure 4A, below, the upper part depicts a sequence of memory operations, while the lower part depicts the signals transmitted on data lines leading from two memory dies (GIO1 and GIO2) to the shared I/O circuit.



Appx3626(Fig. 4A) (annotated). As shown in the image’s lower part, the data lines (GIO1 and GIO2) between the memory dies and the shared I/O circuit carry the same signal (either “Read Data” or “Write Data”) because they connect to the same shared I/O circuit. Appx3631(¶42). Moreover, Kim’s memory has *write* latency—delay between a write command and when “Write Data” is loaded on the data lines (blue arrow above). Conventionally, memory dies also have *read* latency—delay between a read command and when read data is loaded on the DQ line. See Appx3580(¶86). But Kim’s memory dies have negligible read latency—when a read operation is performed (green arrow above), “Read Data” is “immediately loaded on the data . . . line GIO1 or GIO2,” with no delay. Appx3631(¶43); see Appx6453-54(¶137). Thus, if a “write operation” to one rank (“Rank0”) is followed by a “read operation”

to the other rank (“Rank1”), as shown above, “Read Data” from the read operation may still be loaded on the data lines when “Write Data” from an earlier write operation becomes available. That results in an “unavoidable” “data collision,” shown by the red “X.” Appx3631 (¶42).

Kim solved that problem with a “rank selecting unit 1100” in the “shared I/O section 1000” on the control die, as shown in Kim’s Figure 3:




Appx3625(Fig. 3) (annotated); Appx3630(¶34). GIO\_Rank0 and GIO\_Rank1 no longer connect directly to the shared I/O circuit (1200), but to the rank-selecting unit (1100, pink), which functions like a railroad switch. The rank-selecting unit establishes a connection from GIO\_Rank0 or GIO\_Rank1 to either the write path (1210,



## 2. *Rajan*

Rajan (U.S. Patent No. 8,041,881) is entitled “Memory Device with Emulated Characteristics.” Appx3634. Rajan discloses a memory device with a “buffer chip” and four DRAM chips. Appx3663(5:36-43). As shown below, two DRAM chips (417A and 417B) share one data bus (415A), while the other two DRAM chips (417C and 417D) share another data bus (415B). Both buses are connected to the buffer chip (413).



Appx3646(Fig. 4) (annotated). Rajan teaches conventional “DRAM chips” with a dedicated I/O circuit, unlike Kim’s shared-I/O-circuit design. Appx3628(¶¶7-11); Appx3668(15:3-12).

### 3. *Riho*

Riho is a patent application (Pub. No. 2011/0026293) for a memory device with a control die and a stack of memory dies. Appx3682(¶28). The memory dies connect to the control die through “TSVs.” Appx3682-84(¶¶34, 45). Riho teaches that it is “very important” for the TSVs to be “*equal* in length to each other” to minimize “skew”—differences in how long it takes for signals to traverse different TSVs. Appx3684(¶50) (“equi-length interconnections”).<sup>4</sup>

### 4. *Wyman*

Wyman (U.S. Patent No. 7,969,192) is entitled “Variable Off-Chip Drive.” Appx3695. Wyman teaches that the driver power required to transmit signals on a TSV depends on the TSV’s length. Appx3703(1:45-48). Shorter TSVs require less power, while “longer” ones require more. Appx3705(6:15-28).

Wyman recognizes that using a driver’s full power is wasteful if less power would be sufficient for a given-length TSV. Appx3703(1:22-28). It thus teaches a driver with multiple “tap[s]” (power transistors) that can be enabled selectively to “suppl[y]” only as much power as “necessary” for a particular TSV. Appx3703-04(2:65-3:1). The same driver thus can be used whether “high” or “low drive

---

<sup>4</sup> The parties disputed whether each TSV was “in electrical communication,” within the meaning of the challenged claims, with *every* DRAM chip, or only a *pair* of DRAM chips. Appx11448. Because the Board did not rely on Riho for the relevant limitation, *see* Appx120-26, that dispute is irrelevant on appeal.

capability” is required. Appx3706(7:26-32). Wyman teaches sharing a single driver between multiple TSVs, and touts the “advantag[e]” of doing so. Appx3704-05(4:46-51, 5:26-30).

## **B. The Final Written Decisions**


The Board found all challenged claims obvious. Under Ground 1 of the ’060 patent IPR, it found claims 1-6, 8-14, 16-19, and 29-34 obvious in view of Kim and Rajan. Appx112, 120. Under Ground 2 of the ’060 patent IPR, it found claim 7 of the ’060 patent obvious in view of Kim, Rajan, and Riho. Appx126. And under Ground 3 of the ’060 patent IPR and Ground 1 of the ’160 patent IPR, it found claim 15 and claims 20-28 of the ’060 patent, and all claims of the ’160 patent, obvious in view of Kim, Rajan, and Wyman. Appx50-51; Appx135.

### *1. ’060 Patent Claim 1: Electrically Connecting Die Interconnects to Array Die Data Ports in Only One Group (Kim and Rajan)*

Claim 1 of the ’060 patent requires: (1) a “first group” and “second group” of “array dies”; and (2) a “first die interconnect” and “second die interconnect,” wherein (a) “the first die interconnect” is in electrical communication with all dies in the “first group of array dies” but “not . . . the second group”; and (b) “the second die interconnect” is in electrical communication with all dies in the “second group” but “not . . . the first group.” Appx192(23:65-24:5) (limitations 1[c] & [d]). In other words, the first interconnect must be in communication with *only* dies in the first group, and the second interconnect must be in communication with *only* dies in the

second group. For purposes of its decision, the Board accepted that the “first group” must include “**multiple** array dies.” Appx79 (emphasis added); *see* Appx84 n.9; Appx183(5:59) (first group has “at least **two** array dies” (emphasis added)).


The Board principally relied on Kim. There was no dispute that Kim does not teach the requirement of one die interconnect attached to multiple array dies.



Appx3627(Fig. 5). Petitioner mapped Kim’s “TSV1” (GIO\_Rank0) and “TSV2” (GIO\_Rank1) to the claimed “die interconnect[s],” and its memory dies (C1, C2) to the claimed “array dies.” Appx82. The Board did not dispute that Kim disclosed that each TSV communicates with only **one** memory die. Appx93.

The Board, however, accepted Petitioner’s contention that skilled artisans would have modified Kim so “**multiple** memory chips” would be connected to each of Kim’s “TSV[s].” Appx82 (emphasis added); *see* Appx84-85. Petitioner and the


Board thus further annotated Kim's Figure 5 so that what it depicts as a single die would instead be a group of dies:



Appx82 (annotated by Petitioner). Thus, using the image above, the Board posited a combination that replaced C1 (yellow) with *two* memory dies connected to TSV1, and replaced C2 (orange) with two memory dies connected to TSV2. Appx101-02.

Addressing motivation to combine, the Board pointed to Kim's disclosure that "any number of [memory] chips could be used." Appx93. The Board did not contend that Kim itself taught connecting multiple memory dies to a TSV. Instead, it accepted Petitioner's contention that skilled artisans "would have been motivated to connect additional memory dies," and would have done so by connecting them to "Kim's TSV1 and TSV2." Appx101. The Board determined that "using the existing TSVs of Kim would be beneficial to avoid having to create new ones, 'which would require space and additional circuitry.'" Appx101.

Netlist had argued that Kim “taught away from the proposed combination,” because modifying Kim to connect multiple memory dies to a TSV would reintroduce “unavoidable data collisions.” Appx11429-30; Appx602-03 (citing *DePuy Spine, Inc. v. Medtronic Sofamor Danek, Inc.*, 567 F.3d 1314, 1326-28 (Fed. Cir. 2009)). The Board’s combination would result in two memory dies sharing each of Kim’s data lines, resulting in the collision problem Kim identified. For example, two memory dies (Rank0 and Rank2) would connect to GIO\_Rank0 (*i.e.*, TSV1).



Appx601.

To solve the data-collision problem, Kim introduced a rank-selection switch to avoid conflicts between memory operations involving different ranks; it selectively switches read data to the shared read I/O path, and write data to the shared write I/O path. *See* pp. 19-20, *supra*. But Kim’s switch cannot perform that function when two dies *share* a *single* TSV. Appx601; *see* Appx3631(¶42); Appx6453-54(¶137); pp. 17-20, *supra*. Skilled artisans, Netlist further argued, would also lack

motivation to modify Kim as proposed, because that would render Kim “inoperative for its intended purpose.” Appx602-03.

The Board acknowledged Netlist’s argument that the Petition’s combination “would result in unavoidable data collisions.” Appx102. The Board did not deny that the combination in fact *would* result in data collisions. Appx102-03. And it did not deny that Kim “taught away from the proposed combination” because of those collisions, Appx11429; the Board did not even mention “teaching away.”

The Board instead “turn[ed]” to the separate “issue of whether the proposed combination would have been *within the ability* of” skilled artisans. Appx101 (emphasis added). The Board insisted that skilled artisans “knew how to deal with collisions.” Appx102. It pointed to Rajan, which the Board asserted disclosed how to avoid collisions. Appx105-06. The Board did not identify where the Petition invoked Rajan for that teaching (it did not). Nor did it explain how Rajan, which addresses *conventional* DRAM dies, would have taught skilled artisans how to solve data conflicts that arise from *Kim’s unique architecture*, involving memory dies that (unlike those in Rajan) share data lines and lack on-board I/O circuits.

The Board also read the ’060 patent’s specification to suggest “that there is no *prohibition* on signal collisions.” Appx103 (emphasis added). The “challenged claims,” it said, “do not recite that data collisions must be avoided.” Appx103. The Board cited no evidence to support its theory that skilled artisans would have found

it acceptable to have the data collisions Kim sought to avoid; what level of collisions might be acceptable; whether its combination would exceed that level; or why skilled artisans would ignore Kim’s warnings against collisions. The Board faulted the ’060 patent for not disclosing how to prevent data collisions. Appx103-04. It overlooked that the data-collision problem in Kim stems from Kim’s unique architecture (relocating the I/O circuits), which the patents at issue do not replicate.

For most other limitations of claim 1, the Board simply agreed with Petitioner that the prior art disclosed those limitations. Appx83 (limitations 1[a] and 1[b]); Appx111 (limitation 1[e]). The Board thus found claim 1 obvious over Kim/Rajan. Appx112. It deemed claims 2-5, 8-14, 16-19, and 29-35 unpatentable for the same reasons. Appx120.

2. *’060 Patent Claim 7: Considering Die-Interconnect Load and Reducing Load Differences (Kim, Rajan, and Riho)*

Claim 7 of the ’060 patent depends from claim 1. It recites that the “number of array dies” in the first and second groups—connected to the first and second “die interconnect[s]”—are “selected *in consideration of a load* of the first [and second] die interconnect[s],” “*so as to reduce the difference* between” the “load[s]” on the corresponding “data conduit[s].” Appx192(24:37-47) (emphasis added). The Board concluded that Riho taught claim 7’s additional limitation. Appx126. Skilled artisans, it also held, would have combined Riho with Kim/Rajan (for claim 1, discussed above). Appx120.

Riho, however, taught “using an *equal number of dies* in each group as part of its structure of using *equal length* TSVs.” Appx125 (emphasis added). Unlike claim 7, Netlist urged, Riho does not determine or vary the number of array dies in a group—using more or fewer dies—“‘to account for [differences in] TSV load . . . on different data conduits.’” Appx123. The Board rejected that argument. In its view, claim 7 does not “requir[e]” a “difference in loads” as a “starting point”—*i.e.*, before selecting array dies. Appx126. The Board also declared that claim 7 does not require adjusting the number of array dies—using more or fewer—to achieve any “actual[] . . . reduction in the difference in loads.” Appx126. Under claim 7, it ruled, no difference in load need exist, and no reduction need be possible.

According to the Board, Riho also taught that the number of dies is “selected in *consideration of a load* of the first [and second] die interconnect[s].” Appx124 (emphasis added). The Board pointed to Riho’s disclosure of dividing memory dies into “‘groups’” to “‘reduce . . . load.’” Appx124 (quoting Appx3681 (¶13)). That disclosure refers to Riho’s teaching that only half its memory dies (one of two “chip selection groups”) are enabled at any time; the data “ports” on the other half are “turned off,” reducing (but not eliminating) load from those disabled dies. Appx3688 (¶103); Appx3688-89 (¶¶100-114); Appx6510-12 (¶¶225-229). The Board nowhere explained how that taught *selecting* the number of array dies in a

group considering the *load of the corresponding die interconnects*, as the claim requires.

The Board made no findings as to motivation to combine Kim and Riho. It characterized the Petition as proposing that skilled artisans would replace Kim's *unequal-length* TSVs with Riho's *equal-length* TSVs to “reduce . . . load . . . and to compensate for phase variations.” Appx121; Appx125. The Board never stated it agreed with that argument, much less explained why. Moreover, Netlist had explained that, for cost reasons, skilled artisans would not have lengthened Kim's TSVs, which each extend up from the control die only as far as necessary to connect to the corresponding memory die, Appx3627(Fig. 5), to use Riho's equal-length TSVs that extend entirely through the stack, Appx11442; Appx11472. The Board never addressed that argument. The Board alternatively suggested that skilled artisans could use “Kim's *unequal length* TSVs.” Appx125 (emphasis added). But it never reconciled that assertion with its earlier characterization of the Petition as invoking Riho's equal-length TSVs, or explained what benefit Riho would provide in such a combination. Appx121.

3. *'060 Patent Claim 15: Differently Sized Drivers (Kim, Rajan, and Wyman)*

Claim 15 of the '060 patent recites the memory module of claim 11 wherein the first data conduit has a “first driver” with a “first driver size,” and the second data conduit has a “second driver” with a “second driver size.” Appx193(25:35-42).

Claim 1[e] of the '160 patent similarly recites that the “second driver size” is “different” from the “first driver size.” Appx169(24:2-3).

The Board determined that Wyman taught using different-sized drivers for TSVs of different lengths, as the claims require. Appx129-30; Appx44-45. The Board did not dispute that Wyman discloses using a single adjustable driver with multiple taps, rather than using drivers of different sizes, as the claims require. Appx3706(7:26-32); Appx6501-02(¶213). It likewise did not dispute that Wyman discloses using a *single* driver for *multiple* die interconnects, as opposed to multiple differently sized drivers, one for each interconnect, as the claims provide. Appx3704-05(4:46-51, 5:26-30); Appx6501-02(¶213). The Board acknowledged that Wyman’s choice of “using a single driver may be an acceptable option and may be even a better option than having separate drivers.” Appx132; Appx47. But it asserted that it was “[un]necessary to show that a combination is “the *best* option, only that it be a *suitable* option.”” Appx132. The Board claimed that a skilled artisan “would have recognized,” based on Wyman, “that different, physically separate drivers can be used.” Appx134; Appx50. The Board then relied on the same combination of Kim and Rajan and reasoning discussed above, pp. 23-28, *supra*, to find the remaining limitations and claims disclosed, Appx120-41; Appx51-61.

## **SUMMARY OF ARGUMENT**

I. The Board's ruling that the prior art teaches the claimed array die grouping—where a TSV (die interconnect) electrically connects with all dies in one group but none in another—cannot be sustained. The Board accepted that representative claim 1 requires at least one group of *multiple* array dies sharing a TSV. But the Board's primary reference, Kim, teaches a *one-to-one* correspondence between array dies and TSVs, because unacceptable data collisions will otherwise result. The Board's ruling that skilled artisans would modify Kim to share TSVs anyway, and combine it with Rajan, defies this Court's precedent many times over.

A. The Board failed to consider Kim's teaching away from connecting multiple array dies to each TSV; that itself is error. The Board, moreover, did not address the fact that its proposed modification would either reintroduce the fatal data collisions Kim was designed to avoid—rendering it inoperable for its intended purpose—or require abandoning Kim's principle of operation. And the Board asked only whether skilled artisans *could* have combined the references as claimed, instead of whether they would have been *motivated* to do so, when Kim teaches away. Finally, the Board's effort to dismiss data collisions as irrelevant given the claims cannot be reconciled with controlling precedent or the record.

B. The Board failed to explain why skilled artisans would reasonably have expected success in combining Kim and Rajan. The Board did not explain how

Rajan (which uses conventional memory dies with dedicated I/O circuits) could solve the collision problems in Kim (which proposes memory dies lacking dedicated I/O circuits). The Board’s suggestion that data collisions would be acceptable is unsupported and contrary to the very references the Board invokes.

II. Claim 7 requires grouping array dies on each TSV “in consideration of a load . . . so as to reduce a difference between” loads on different data conduits. The Board’s determination that Riho teaches that limitation rests on an erroneous claim construction that no “difference” in load or “reduc[tion]” of that difference is required. Riho teaches using equal-length die interconnects. As a result, there is no “difference” in load to begin with, contrary to claim 7. And Riho teaches using an equal number of dies in each group. That equal-numbered grouping cannot “reduce” load differences; it leaves them unchanged. The Board’s effort to rewrite claim 7 cannot be sustained. And its motivation-to-combine analysis is unreasoned.

III. The Board erred in finding claim 15 of the ’060 patent and claim 1 of the ’160 patent obvious over Wyman, Kim, and Rajan. The claims require multiple drivers of different sizes. Wyman teaches a *single* driver of a *single size*, disparaging multiple drivers.

IV. For the remaining limitations, the Board simply accepted Petitioner’s arguments without discussion. The APA and precedent require the Board to explain why it accepts even undisputed arguments.

## **STANDARD OF REVIEW**

Obviousness is a question of law reviewed de novo. *Virtek Vision Int’l ULC v. Assembly Guidance Sys., Inc.*, 97 F.4th 882, 886 (Fed. Cir. 2024). Underlying fact findings are reviewed for substantial evidence. *Id.* While motivation to combine is a question of fact, this Court applies an especially “‘thorough and searching’” review to that issue, demanding “‘specificity’” in the Board’s findings. *In re Nuvasive, Inc.*, 842 F.3d 1376, 1381-82 (Fed. Cir. 2016). Claim construction is reviewed de novo. *HTC Corp. v. Cellular Commc’ns Equip., LLC*, 877 F.3d 1361, 1367 (Fed. Cir. 2017). Board decisions must be vacated if they are “‘arbitrary, capricious, an abuse of discretion,’” “‘not in accordance with law,’” or “‘unsupported by substantial evidence.’” *Pers. Web Techs., LLC v. Apple, Inc.*, 848 F.3d 987, 992 (Fed. Cir. 2017).

## **ARGUMENT**

### **I. THE BOARD’S DETERMINATION THAT THE PRIOR ART TEACHES A FIRST AND SECOND GROUP OF ARRAY DIES, EACH GROUP SHARING SEPARATE INTERCONNECTS AS IN THE CLAIMS, CANNOT BE SUSTAINED**

Netlist’s invention reduces the load on each data conduit by separating the array dies into groups, where each die interconnect is connected to all the dies in *only* one group. Representative claim 1 of the ’060 patent thus requires: (1) a “first” and “second group” of array dies; and (2) a “first” and “second die interconnect,” wherein: (a) “the first die interconnect” is connected to every die in the “first group” but “not . . . the second group”; and (b) “the second die interconnect” is connected

to every die in the “second group” but “not . . . the first group.” Appx192(23:65-24:5) (limitations 1[c] & [d]).

The Board accepted for present purposes that the first group must have “**multiple** array dies.” Appx79 (emphasis added); *see* Appx183(5:59-60) (“at least two array dies”). Kim, the Board’s primary reference, undisputedly did not teach that limitation. Appx100. Kim taught precisely the opposite: It taught connecting each TSV (or “data line”) to only one die (or “rank”). Thus, Kim teaches a one-to-one correspondence—not connecting multiple dies in a group to a TSV. Indeed, Kim’s unique architecture, which requires dies to share a single I/O circuit, precludes connecting multiple dies to a single TSV because (according to Kim itself) unacceptable data collisions would result.

The Board nonetheless posited that skilled artisans would have modified Kim to connect each TSV to **two** dies. Appx100. That cannot be reconciled with this Court’s precedents or the Administrative Procedure Act’s (“APA”) requirements. It overlooks that Kim teaches away from that approach. It ignores that, according to Kim itself, that modification would render Kim’s invention inoperative because of “unavoidable” and unacceptable “data collisions.” Appx3631(¶42). Kim itself seeks to avoid those collisions by using separate data lines for each die and a rank-selecting unit.

The Board's effort to invoke Rajan as teaching how to avoid those data collisions is no answer. The motivation-to-combine analysis must begin with the reference and explain *why* skilled artisans would have been motivated to modify it. Here, the Board began with Kim, an invention for sharing I/O circuits that uses separate data lines and a rank-selecting unit to avoid the data collisions that would otherwise result. But its posited combination eliminated or rendered ineffective the shared I/O circuit, separate data lines, and rank-selecting circuit that are the core of Kim's invention, eviscerating its principle of operation. And the Board failed to explain why skilled artisans would think Rajan would somehow solve Kim's data-collision issue, when Rajan uses conventional memory dies that each have an I/O circuit—precisely the architecture Kim's innovation eliminated. The Board's analysis elides those issues.

The Board's finding that skilled artisans would have reasonably expected success in achieving the proposed combination fares even worse. It departs from how that analysis must be conducted. It purports to find data collisions acceptable despite Kim's clear statement they are not, along with the reference's strenuous efforts to avoid them. The decision is unreasoned, unreasonable, and unsupported. Reversal is warranted.

**A. The Board’s Motivation To Modify Kim—By Connecting Multiple Array Dies to Each TSV—Defies Precedent and the APA**

Obviousness requires showing that skilled artisans “would have been motivated to” modify or combine prior-art references to arrange the elements as in the claims. *Virtek Vision Int’l ULC v. Assembly Guidance Sys., Inc.*, 97 F.4th 882, 886-87 (Fed. Cir. 2024). The Board must provide a “full and reasoned explanation” why skilled artisans would have been motivated to pursue the posited combination. *In re Lee*, 277 F.3d 1338, 1342 (Fed. Cir. 2002). That obligation to provide “reasoned analysis” is “especially” important in motivation-to-combine cases because of the risk of impermissible “*ex post* reasoning.” *In re Van Os*, 844 F.3d 1359, 1361-62 (Fed. Cir. 2017).

Here, the Board insisted skilled artisans “‘would have been motivated to’” modify Kim to “‘connect [multiple] memory dies to Kim’s TSV1 and TSV2,’” so as to enhance memory capacity without multiplying the number of TSVs. Appx100-01. That finding, and the reasoning underlying it, defies this Court’s precedents in at least four ways. First, the Board invoked Kim as its principal reference, but failed to address the argument that Kim teaches away from the Board’s proposed modification. Second, the Board’s modification of Kim is not merely contrary to Kim’s principle of operation. It eviscerates the innovation Kim proposed and renders Kim inoperable for its intended purpose. Third, the Board’s cited motivation is entirely generic, divorced from the particulars of the prior art. The Board’s asser-


tion that motivation existed because the “proposed combination would have been within the *ability*” of skilled artisans is legal error as well. Appx101 (emphasis added). The Board was required to explain why skilled artisans would have been *motivated* to pursue the proposed combination, not only that (with the benefit of hindsight) it was *possible* to achieve. Fourth, the Board’s treatment of data collisions is contrary to this Court’s decisions and unmoored from the record.

1. *The Board Legally Erred in Overlooking that Kim Taught Away from Connecting Multiple Array Dies to Each TSV*

Skilled artisans generally lack motivation to modify a reference to achieve the claims where the reference “‘teaches away’ from making [those] modifications.” *Polaris Indus., Inc. v. Arctic Cat, Inc.*, 882 F.3d 1056, 1068 (Fed. Cir. 2018). A reference “‘teaches away’ from” a modification when skilled artisans, “‘upon reading [it], would be discouraged from’” making it or “‘would be led in a direction divergent from’” that path. *Id.* at 1069; see *DePuy Spine, Inc. v. Medtronic Sofamor Danek, Inc.*, 567 F.3d 1314, 1327 (Fed. Cir. 2009). Teaching away need not be “explicit”; a reference may teach away where “skilled artisan[s] would interpret” it to “implicit[ly] disparag[e]” a course of action. *AstraZeneca AB v. Mylan Pharms. Inc.*, 19 F.4th 1325, 1337 (Fed. Cir. 2021). Here, Kim teaches away from attaching multiple dies to a TSV. It expressly teaches the opposite. It requires a one-to-one correspondence between TSVs (data lines) and dies (ranks), disparaging the alternative. Appx84.

Kim’s “invention” is a memory package including a “plurality of data . . . lines *respectively connected* to [a] plurality of [memory] chips.” Appx3629(¶17) (emphasis added); Appx3629-30(¶28) (data lines “GIO\_Rank0” and “GIO\_Rank1 connect . . . to the respective first and second” ranks); Appx3630 (¶33) (similar). The term “respectively” means “each to each,” *Webster’s New International Dictionary* 2123 (2d ed. 1934), and “generally requires a *one-to-one relationship* between the words it modifies,” *Mars, Inc. v. Coin Acceptors, Inc.*, 514 F. Supp. 2d 624, 631 (D.N.J. 2007) (emphasis added). Kim’s statement that the data lines are “respectively connected” to the memory dies thus means that “each” data line is in a “one-to-one relationship” with a single corresponding memory die.


That is how Kim illustrates its invention: Each TSV (TSV1 and TSV2) connects to a single memory die (C1 and C2, respectively).



Appx3627.

Consistent with that, every embodiment in Kim discloses a *one-to-one correspondence* between data lines and memory dies (ranks). See Appx3629(¶14) (in the “present invention,” “a first data . . . line” is “connected to a first rank” and a “second data . . . line” is “connected to a second rank”); Appx3629(¶26) (data line “GIO\_Rank0” connected to “Rank0” and “GIO\_Rank1” connected to “Rank1”); Appx3630(¶33) (“first and second ranks” connected to “first and second data . . . lines”); Appx3631-32(¶¶42, 44, 45, 49).

Kim’s one-to-one correspondence between TSVs and memory dies is *essential* to the invention it discloses—having dies share an I/O circuit, rather than each having its own. As shown below, Kim’s memory dies (green, blue) share I/O circuitry (yellow, orange) in a control die. Memory dies (Rank0, green, and Rank1, blue) each have their own data line (GIO\_Rank0 and GIO\_Rank1, respectively) connecting them to that circuitry:



Appx3625(Fig. 3) (annotated). There is a reason for that: Because Kim’s dies lack I/O circuits of their own to control signal timing, two dies “cannot share” a data line (TSV) to a shared I/O circuit in an operable device; “unavoidable” “data collision[s]” would result. Appx3631(¶42); see Appx11424-29.<sup>5</sup> If that does not teach away from shared TSVs, it is hard to imagine what would.

Kim, moreover, connects the separate TSV from each die (rank) (GIO\_Rank0 and GIO\_Rank1) to a rank-selecting circuit (pink above). That rank-selecting circuit is critical to ensuring dies can share an I/O circuit. It functions like a switch or

<sup>5</sup> Indeed, Kim’s Figure 4A (reproduced p. 18, *supra*) illustrates the data collisions that would result from sharing data lines. There, “Rank0 and Rank1 share a data [I/O] circuit,” and their respective “data . . . lines operate in the same manner”—*i.e.*, also are effectively “share[d]”; the result is data collisions (red “X”). Appx3631(¶42); see pp. 17-19, *supra*.

multiplexer, directing data to read or write I/O pathways as appropriate. Appx3631(¶45). For example, during a write operation, the rank-selecting circuit directs data from the shared I/O write path (orange) through GIO\_Rank0 or GIO\_Rank1, depending on the memory die being written to. That design—where each memory die has a separate data line to the rank-selecting unit—ensures “a data collision *does not occur*.” Appx3631(¶45) (emphasis added). The rank-selection circuit has *two* lines to the shared I/O circuit, again to avoid data collisions. Appx3631(¶¶44-45).

Skilled artisans would have understood those disclosures to teach away from “sharing the same” TSV among memory dies, as in the Board’s combination. Appx6450-51(¶134). Consequently, skilled artisans following Kim would add memory dies by connecting them to additional, “*non-shared* TSVs.” *Id.* (emphasis added); see *AstraZeneca*, 19 F.4th at 1337; *DePuy*, 567 F.3d at 1327.

Confronted with that, the Board’s motivation-to-combine analysis said nothing. Netlist urged that Kim “taught away” from connecting two memory dies to a single TSV. Appx11424-29; Appx601-03 (citing *DePuy Spine*, 567 F.3d at 1326-28). But the Board never addressed teaching away in assessing motivation to combine. That was legal error. The APA required the Board to “fairly contemplat[e]” and “resolv[e]” Netlist’s key “arguments.” *Provisur Techs., Inc. v. Weber, Inc.*, 50 F.4th 117, 124 (Fed. Cir. 2022). The Board’s failure to do so here

“violate[d] the APA” and requires vacatur. *Id.* at 124; *see In re Nuvasive, Inc.*, 842 F.3d 1376, 1384 (Fed. Cir. 2016).<sup>6</sup>

2. *The Board’s Analysis Defies Precedent by Proposing Modifications That Eviscerate Kim’s Invention and Purpose*

Where “modifying” a reference in a particular manner would “undermine” the reference’s “purpose,” skilled artisans typically lack “motivation to make such modifications.” *Medtronic, Inc. v. Teleflex Innovations, S.à.r.l.*, 69 F.4th 1341, 1349 (Fed. Cir. 2023). Likewise, skilled artisans typically lack motivation to modify a reference in ways that “alter[ ] [its] inventive concept,” *Chemours Co. FC, LLC v. Daikin Indus., Ltd.*, 4 F.4th 1370, 1376 (Fed. Cir. 2021), change its “‘principle of operation,’” *Plas-Pak Indus., Inc. v. Sulzer Mixpac AG*, 600 F. App’x 755, 759 (Fed. Cir. 2015), or render it “inoperable for its intended purpose,” *In re Gordon*, 733 F.2d 900, 902 (Fed. Cir. 1984). Those authorities reflect the common-sense understanding that skilled artisans do not start with a prior-art invention only to modify it so extensively as to leave nothing of the original invention. The Board’s motivation-to-combine analysis defies those principles.

---

<sup>6</sup> The Board mentions data collisions *only* in connection with expectation of success. Appx102-07. But motivation and expectation of success are distinct inquiries. *See Endo Pharms. Inc. v. Actavis LLC*, 922 F.3d 1365, 1376 n.11 (Fed. Cir. 2019). And the Board’s expectation-of-success analysis is defective regardless. *See* pp. 51-54, *infra*.

Kim’s “invention”—*i.e.*, its inventive concept—is a memory device where a centralized “[I/O] circuit” is “share[d]” by all memory dies. Appx3628(¶12); *see* Appx3629-31 (¶¶ 14, 17, 25, 28-29, 34, 48). However, removal of the I/O circuit from the memory dies, and use of a shared I/O circuit, threatens fatal “data collision[s].” Appx3631(¶42). Removing the I/O circuit from the memory die eliminates the conventional delay between when a read is performed and when the data is loaded on the data lines. Thus, in Kim, a “read operation” results in data being “*immediately* loaded on the data . . . line.” Appx3631(¶43) (emphasis added). That change in timing—a consequence of eliminating individual I/O circuits on each die—“result[s] in a data collision,” with uncoordinated dies simultaneously loading data, in response to different commands, to a shared input on the shared I/O circuit. Appx3631(¶¶ 42-43); Appx3626(Fig. 4A). The remainder of Kim’s teachings—the one-to-one correspondence between TSVs and memory dies, and the “rank selecting unit”—represent efforts to *solve* that problem. In service of sharing I/O circuits among memory dies, they are directed to avoiding the data collisions that otherwise would result. Appx3630(¶35); *see* pp. 38-41, *supra*.

The Board’s modifications to Kim, however, would eviscerate Kim’s invention and undermine its purpose. Connecting two array dies to one TSV, as the Board proposed, would eliminate Kim’s one-to-one correspondence between TSVs and memory dies and its purpose of avoiding data collisions. *See* pp. 38-41, *supra*. The

resulting sharing of data lines between memory dies, moreover, would eliminate the switching/multiplexing function of the rank-selecting unit. The rank-selecting unit avoids data collisions by directing data between separate “data . . . lines” and separate read and write paths of the shared I/O circuit. Appx3630-31 (¶¶35-38, 44-45); *see pp. 40-41, supra*. Because the rank-selecting circuit selectively connects those separate *data lines* to separate pathways in the shared I/O circuit, it cannot mitigate data collisions that occur when two memory dies *sharing a data line* simultaneously attempt to load data on that shared line. Just as a railroad switch cannot solve collisions between trains on shared portions of track *before the switch*, Kim’s rank-selecting unit cannot solve collisions caused by sharing data lines before the lines meet the rank-selecting unit. The Board’s modifications reintroduce the very “unavoidable” “data collision[s]” Kim’s disclosure was directed at solving. Appx3631 (¶42).

Insofar as the Board suggested skilled artisans would have solved collisions by “chang[ing] [the] operation” of Kim’s memory dies to match those of Rajan, that eliminates Kim’s invention altogether. Appx107. Rajan discloses conventional “DRAM circuits,” where each memory die has an I/O circuit and read operations incur a “latency” before data is loaded on the data lines. Appx3580 (¶86). The Board nowhere found that Kim’s unique memory dies could operate like Rajan’s conventional memory dies without putting I/O circuits back into Kim’s memory

dies. But putting individual I/O circuits back into Kim’s dies abandons Kim’s “inventive concept” and “‘principle of operation.’” *Chemours*, 4 F.4th at 1376; see *Plas-Pak*, 600 F. App’x at 759; *Intel Corp. v. Koninklijke Philips, N.V.*, No. 22-2034, 2024 WL 725243, at \*1 (Fed. Cir. 2024) (skilled artisans would not have “modif[ied]” prior-art reference to “use . . . [a] multi-bit message,” where a “‘*single-bit* [message]’” was “‘essential’” to the prior-art invention).

This Court’s precedent forecloses such combinations. In *Chemours*, the reference’s “inventive concept” was a “polymer” with a “**narrow** molecular weight distribution.” 4 F.4th at 1376 (emphasis added). This Court ruled that the reference did not render obvious a claimed combination that “would have required **broadening** the molecular weight distribution of the polymer.” *Id.* (emphasis added). Skilled artisans, the Court explained, would not have been motivated to make modifications that “would necessarily involve altering” the reference’s “inventive concept” of a “narrow” molecular-weight distribution. *Id.* Likewise here, skilled artisans would not have been motivated to substitute Rajan’s conventional memory dies (with dedicated I/O circuits) into Kim to avoid data-collision problems, because that would require abandoning Kim’s inventive shared-I/O-die concept.

*Medtronic* underscores the Board’s error. There, the prior-art reference focused on using “sealing balloons to prevent embolic flow” in an artery. 69 F.4th at 1348. This Court rejected the argument that a “skilled artisan would be motivated

to remove the sealing balloons,” because that “would remove” the feature on which the reference’s “‘entire premise’ was founded.” *Id.* at 1347, 1350; *see id.* at 1350 (“destruction of [the reference’s] entire purpose”). Likewise here, dispensing with Kim’s shared I/O circuitry to achieve the same timing characteristics as Rajan’s conventional DRAM circuits would destroy Kim’s purpose. The Board’s departure from this Court’s motivation-by-combine precedents requires reversal. *See In re Rouffet*, 149 F.3d 1350, 1358 (Fed. Cir. 1998).

The Board addressed none of those issues. Netlist had urged skilled artisans would lack motivation to modify Kim as the Board proposed, because it would render Kim “inoperative for its intended purpose.” Appx601-03; Appx11434-35. The Board, however, nowhere addressed whether Kim would be operable for its original purpose of sharing the I/O circuit after the Board’s modifications. That, too, requires vacatur. *Provisur*, 50 F.4th at 124.

### 3. *The Board’s Generic Assertion of Motivation Was Insufficient*

Instead of analyzing motivation to modify Kim in view of Kim’s actual disclosures and teachings, the Board asserted that skilled artisans “‘would have been motivated to connect [multiple] memory dies to Kim’s TSV1 and TSV2,’” to avoid needing to “‘creat[e] new TSVs.’” Appx100-01. But that “generic” assertion of motivation is insufficient. *ActiveVideo Nets., Inc. v. Verizon Commc’ns, Inc.*, 694 F.3d 1312, 1328 (Fed. Cir. 2012). The motivation-to-modify analysis must reflect

the reference’s specific teachings and must “explain why” skilled artisans would have modified the prior art “*in the way the claimed invention does.*” *Id.* (emphasis added). Seeking to add additional memory dies is a broad goal; it is not a concrete motivation that explains why skilled artisans would have pursued *the particular modifications* required to achieve the claimed invention.<sup>7</sup>

In discussing expectation of success (not motivation), the Board insisted that modifying Kim to connect multiple memory dies to a TSV “would have been within the ability of a person of ordinary skill in the art.” Appx101. But whether skilled artisans would be *capable of* modifying the prior art in a particular way does not show that they would have been *motivated* to do so. Establishing motivation to combine requires proving not only what skilled artisans “*could have made,*” but what they “*would have been motivated to make*” in view of the prior art. *Belden Inc. v. Berk-Tek LLC*, 805 F.3d 1064, 1073 (Fed. Cir. 2015).

The Board’s short-circuiting of that requirement is textbook “hindsight” error. *In re Fritch*, 972 F.2d 1260, 1266 (Fed. Cir. 1992). Instead of analyzing whether Kim’s teachings would have persuaded skilled artisans to pursue a different path from the claims—*i.e.*, adding memory dies to separate TSVs to maintain Kim’s one-

---

<sup>7</sup> The Board’s assertion that sharing TSVs was one of “a finite number of known ways to connect additional dies,” Appx94, similarly fails to address whether sharing TSVs was a “known” solution in the context of Kim’s unique shared-I/O-die architecture.

to-one correspondence—the Board used the “claimed invention as [a] . . . ‘template’ to piece together the teachings of the prior art.” *Id.* The Board simply assumed that skilled artisans would have connected multiple array dies to each TSV—sharing it—and then offered *post hoc* rationales for how doing so would have been “within [their] ability.” Appx101-08. That defies patent law and the APA alike. *See* pp. 47-49, *supra*.

4. *The Board’s Effort To Downplay Data Collisions Is No Answer*

The Board did not dispute that connecting multiple of Kim’s memory dies to each TSV would reintroduce data collisions—Kim says as much and goes to great lengths to avoid them. Appx103; pp. 38-41, *supra*. In addressing likelihood of success, the Board suggested that perhaps collisions are not a big deal. Appx103. But likelihood of success differs from motivation to combine. And any such theory defies the law and the record regardless.

The Board asserted that data collisions were not a problem because the “*claims* do not recite that data collisions must be avoided.” Appx103 (emphasis added). That misapprehends the law. The motivation-to-combine analysis focuses on the “desirability of the modification” to skilled artisans. *Fritch*, 972 F.2d at 1265-66 & n.12. Thus, the Board must consider “[u]nclaimed factors” in evaluating motivation to combine if “skilled artisan[s] would reasonably consider” those factors. *Natera, Inc. v. NeoGenomics Labs., Inc.*, 106 F.4th 1369, 1378 (Fed. Cir.

2024); see *Auris Health, Inc. v. Intuitive Surgical Operations, Inc.*, 32 F.4th 1154, 1159 (Fed. Cir. 2022). For example, skilled artisans would be deterred from making a modification that would reduce a surgery’s “precision,” even if precision was unclaimed. *Auris*, 32 F.4th at 1159. Likewise, while benefits like “efficiency,” “reliability,” and “cost” typically are unclaimed, those factors are important to skilled artisans and must be considered by the Board. *Gen. Elec. Co. v. Raytheon Techs. Corp.*, 983 F.3d 1334, 1346-48 (Fed. Cir. 2020). The Board’s legal error—“discard[ing] evidence” bearing on motivation to combine under the mistaken view that it was unclaimed and thus not “relevant”—requires vacatur. *Polaris*, 882 F.3d at 1069.

Any suggestion that skilled artisans would have found data collisions acceptable, Appx103-04, is unsustainable. Findings must be supported by “substantial evidence,” *Virtek*, 97 F.4th at 886-87, and the Board must explain its reason for crediting some evidence over other sources, *Princeton Vanguard, LLC v. Frito-Lay N. Am., Inc.*, 786 F.3d 960, 970 (Fed. Cir. 2015). Kim is unambiguous that such collisions are *fatal*. Kim says that “Rank0 and Rank1 *cannot share*” an I/O circuit unless data collisions are eliminated. Appx3631(¶42) (emphasis added). Kim solves that problem by teaching a one-to-one correspondence between memory dies and TSVs, and a rank-selecting circuit, so “a data collision *does not occur*.”

Appx3631(¶42) (emphasis added). The Board never explained why skilled artisans would wholly disregard those disclosures.

The sole contrary evidence the Board invoked was a passing reference in the '060 patent supposedly about “reducing” (rather than eliminating) data collisions, from which the Board inferred that some level of collisions is acceptable. Appx103. But the '060 patent has nothing to say about whether skilled artisans would have considered the specific data-collision problem in *Kim* tolerable, when Kim itself said otherwise. Moreover, as explained below, the Board nowhere addressed whether—even if some collisions could be tolerated (through error-correction or otherwise)—the quantity of data collisions that would result from modifying Kim as proposed would have deterred skilled artisans from the Board’s proposed course.

#### **B. The Board’s Expectation-of-Success Analysis Was Deficient**

The Board was required to explain why skilled artisans would have had “‘a reasonable expectation of success’” in modifying Kim as the Board proposed. *TQ Delta, LLC v. CISCO Sys., Inc.*, 942 F.3d 1352, 1357 (Fed. Cir. 2019). Insofar as the Board was of the view skilled artisans would use Kim’s memory dies—which lack dedicated I/O circuits—the Board failed to provide the requisite explanation. As Netlist explained, connecting a group of Kim’s memory dies to a single TSV would have resulted in the same unacceptable data collisions that Kim itself described as rendering the device “‘inoperable.’” Appx102; pp. 43-47, *supra*.

Those collisions result from the fact that Kim's dies lack their own I/O circuits; the resulting lack of read latency causes unavoidable data collisions on shared lines leading to the shared I/O circuit. *See* pp. 17-19, *supra*. Indeed, Kim sought to avoid those collisions by using separate data lines for each die, along with a rank-selection switch. *Id.* The proposed combination, however, eliminates the separate, unshared data lines and denies the switch utility. *Id.* It is hard to see how skilled artisans would expect success using a combination that eliminates the very features Kim deemed critical to proper operation.

Insofar as the Board read the '060 patent to mean skilled artisans would have thought that the resulting data collisions are acceptable, or insisted that the "claims do not recite that data collisions must be avoided," Appx103, that fails for the reasons above: Those rationales defy precedent, the record, and Kim's disclosures. *See* pp. 49-51, *supra*. For expectation of success, the question is whether a combination would "render the *prior art*"—*i.e.*, **Kim**—"inoperable for its intended purpose" according to Kim itself. *Plas-Pak*, 600 F. App'x at 757-58.

The Board's insistence that skilled artisans would have understood they could avoid collisions through "timing solutions," Appx102, fares worse still. Even if "techniques for avoiding collisions on shared data lines were well-known" to skilled artisans, Appx102 (citing Appx11540-45), that says nothing about whether such solutions were known if one began with **Kim's** unique architecture. For one thing,

Kim discloses only its “changed architecture”—*i.e.*, its one-die-per-TSV structure with a rank-selecting unit—as “the known solution to [its data collision] problem.” Appx11598. It does not “suggest [other] solutions,” such as “modifying latency” or “otherwise adjusting the electrical or timing characteristics of its memory systems.” Appx6456(¶141). It was error for the Board to “ignore[] the express disclosure in [Kim]” (one die per TSV with a rank-selecting unit) and instead “rel[y] on teachings from other references” “that were not concerned with the particular problems”—collisions resulting from the absence of separate I/O circuits—that Kim “sought to solve.” *Chemours*, 4 F.4th at 1376.

The Board speculated Rajan’s “buffer chip” could solve Kim’s collisions by changing memory timing, Appx105-06, but the Petition did not rely on the buffer’s timing functions for claim 1. The AIA and APA precluded the Board from finding obviousness based on “portions of [the] prior art” the “Petition” did not rely on. *Corephotonics, Ltd. v. Apple Inc.*, 84 F.4th 990, 1002 (Fed. Cir. 2023). Regardless, Rajan relies on conventional memory chips, each with dedicated I/O circuits. Appx3628(¶¶7-11); Appx3668(15:3-12). Rajan’s buffer can mitigate different timing issues that arise where conventional chips are utilized. But Kim’s collision problem arises because **Kim’s** memory dies lack their own, individual I/O circuits. One can attempt to superimpose Rajan atop Kim only if one assumes away the core

of the invention Kim discloses—the use of shared rather than individual I/O circuits. That is impermissible. *Chemours*, 4 F.4th at 1376.

The Board also reasoned that, because the '060 patent does not provide “details on how to control timing to reduce or prevent collisions,” preventing collisions must be “within the knowledge and skill” of skilled artisans. Appx103-04. But the '060 patent relies on conventional array dies and nowhere mentions removing or sharing I/O circuits. Appx182(4:55-67). That skilled artisans would have understood how to manage collisions in the context of the '060 patent does not mean that, if they started with Kim, they could do so without the individual I/O circuits that Kim eliminates. Kim itself did not think that so easy: It proposed an entirely ***different solution*** because, without I/O circuits in each die, sharing data lines between memory dies would result in “unavoidable” collisions. Appx3631(¶42); pp. 38-41, *supra*. Moreover, the fact that the '060 patent’s “inventors were ultimately successful” is irrelevant to expectation of success at the time of the invention. *Life Techs, Inc. v. Clonetech Labs., Inc.*, 224 F.3d 1320, 1326 (Fed. Cir. 2000).

The Board faulted Netlist’s expert, Dr. Brogioli, for “fail[ing] to account for the combined teachings of Kim and Rajan and the knowledge and skill” of artisans. Appx107. The Board, however, never explained how those “combined teachings” or skilled artisans’ “knowledge and skill” would overcome Kim’s unique collision problem without Kim’s unique solution. The Board also cryptically suggested that

Rajan’s teachings regarding “‘rank multiplication’”—a feature that enables two memory dies to appear as one memory die to the memory controller—would solve Kim’s problem. Appx107-08. But, again, that solution assumes conventional memory dies, as in Rajan, not memory dies like Kim’s that lack dedicated I/O circuits, the starting point for the Board’s reasoning here.

## **II. THE BOARD’S DETERMINATION THAT RIHO TEACHES CLAIM 7 OF THE ’060 PATENT CANNOT BE SUSTAINED**

As explained above, the load on each “data conduit”—*i.e.*, die interconnect and attached array dies—includes the load from the die interconnect itself (which depends on its length) *plus* the load imposed by array dies in electrical communication with that interconnect. *See* p. 9, *supra*. Claim 7 of the ’060 patent, which depends from claim 1, requires that the “number of array dies” in the “first” and “second” groups of dies be “selected *in consideration of a load* of the first” and “second *die interconnect[s]*,” “*so as to reduce a difference between a first load* on the first data conduit and a *second load* on the second data conduit.” Appx192(24:37-47) (emphasis added). In other words, where data conduits have different loads because of different-length die interconnects, the differences in total load are reduced by adjusting the number of dies in each group of array dies. A data conduit with greater load due to a longer interconnect is connected to fewer array dies; one with a shorter interconnect and less resulting load can be connected to more array dies.

The Riho reference the Board invoked for that limitation teaches no such thing. Riho concededly teaches “*equal length*” TSVs with an “*equal number*” of dies connected to each. Appx125 (emphasis added). The Board could invoke Riho only by misconstruing claim 7: It construed claim 7 to encompass “using equal length TSVs,” so that there was no load “difference[.]” to start with, and “using an equal number of dies in each group,” which could never “reduce” load differences even if any existed. Appx123; Appx126. That construction defies the claim language and specification. And the Board’s analysis of motivation to combine and expectation of success falls far short of the APA’s demands.

**A. The Board Misconstrued Claim 7**

1. Despite the absence of an “express construction,” Appx80, the Board made specific determinations as to the “scope” of claim 7 that “fall within the ambit of claim construction.” *HTC Corp. v. Cellular Commc’ns Equip., LLC*, 877 F.3d 1361, 1367 (Fed. Cir. 2017). Those findings constitute “claim construction”—and are reviewed in this Court as such, de novo. *Id.*; see *Google LLC v. EcoFactor, Inc.*, 92 F.4th 1049, 1055-56 (Fed. Cir. 2024).

Here, claim 7 requires “select[ing]” “the number of array dies” for the “first” and “second” groups of dies “in consideration of” the load of their respective “die interconnect[s],” “*so as to reduce a difference between a first load*” and a “*second load*.” Appx192(24:37-47) (emphasis added). But Riho merely “discloses using . . .

equal length TSVs” and “an equal number of dies in each group.” Appx125. As the Board recognized, the sole dispute thus was whether that disclosure falls within the “scope [of] claim 7.” Appx123. As Netlist explained, it does not: Claim 7 requires “adjust[ing]” the “number of array dies in [each] group” to offset a difference in “TSV loads,” such as differences from TSVs having “different lengths.” Appx11445. That is what the claim means when it requires the number of dies in each group to be “selected” to “reduce *a difference*” between the loads of the first and second conduits. If there is no difference in loads to begin with because the TSVs are of equal length, as in Riho, there is no “difference” to “reduce” by adjusting the number of array dies.

2. To reach the opposite result, the Board asserted that claim 7 does not “require[ ]” any “difference in loads” as a “starting point.” Appx126. In its view, the claim is satisfied where the load from TSVs is identical, leaving no “difference in loads” to “reduc[e].” Appx126. But claim construction begins with the “‘plain and ordinary meaning’” of the words of the claim. *Cont’l Circuits LLC v. Intel Corp.*, 915 F.3d 788, 796 (Fed. Cir. 2019). Here, claim 7 plainly requires that the “number of array dies” in the first and second groups be “selected in consideration of a load of the” corresponding “die interconnect[s],” “so as *to reduce a difference* between a first load on the first data conduit and a second load on the second data conduit.” Appx192(24:37-47) (emphasis added). That language requires a

“difference between [the] . . . load[s]” of the data conduits—such as from die interconnects having different lengths—before the number of array dies is “selected.” “Difference” means the “measure of being different or unlike.” *Webster’s, supra*, at 726. Thus, the loads on the first and second data conduits must be “different” or “unlike” to start with, rather than the same or alike.

The claim also requires adjusting the number of array dies—using more or fewer dies in each conduit—“*so as to reduce a difference*” in loads. Appx192(24:37-47) (emphasis added). “Reduce” means to “to diminish, esp. in bulk, amount, or extent.” *Webster’s, supra*, at 2088. For the selection to reduce load differences, there must be a difference to reduce; one cannot reduce a difference that does not exist. And using the *same number* of dies for every data conduit, as Riho teaches, can never “reduce” or “diminish” a difference in load. Adding the same number of dies to each group adds an identical load to each data conduit, leaving any load difference unchanged. A difference between two numbers cannot be “reduced” by adding the same value to each.

The “specification” confirms the ordinary meaning. *Cont’l Circuits*, 915 F.3d at 796. It explains that the “*difference* between the load of any pair of [data] conduits” can be “*reduce[d]*” by appropriate selection of “the number of array dies” for each data conduit considering the “length of the [corresponding] die interconnect.” Appx186(11:40-50) (emphasis added). Table 1, below, shows how

the number of array dies can be varied (using one die or two dies), based on the length of the corresponding die interconnect (measured by the number of “Die Interconnect Segments”), to maintain the difference in loads between data conduits below a certain threshold.

TABLE 1

| Conduit | Number of<br>Array Dies | Number of Die<br>Interconnect<br>Segments | Capacitive<br>Load | Deviation from<br>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                              |

Appx187(13:8-17). Not one of those examples suggests that one can “reduce” a “difference” in loads where there is no difference to start with, or by selecting an equal number of dies for every group such that any load difference (if one exists) remains unchanged.

Simply put, the Board posited that the claims are satisfied even where there is *no* “difference in loads” to start with, each die interconnect having the same length, and the selection of an equal number of dies does not “reduce” load differences but leaves them constant. The claims plainly require otherwise. The Board’s erroneous construction requires reversal. *See Google*, 92 F.4th at 1059.

## **B. Reversal Is Warranted Because Riho Does Not Teach Claim 7**

1. Once claim 7 is properly construed, Riho cannot teach it as a matter of law. The Board acknowledged “Riho discloses using an equal number of dies in each group as part of its structure of using equal length TSVs.” Appx125. That does not satisfy claim 7 twice over.

*First*, Riho lacks the required “***difference*** between a first load on the first data conduit and a second load on the second data conduit” to start with (*i.e.*, load from the first and second “die interconnect[s],” before selection of the array dies). Appx192(24:37-47) (emphasis added). If each data conduit has TSVs of equal length, the resulting “load” from the die interconnects is “the same” for each. Appx6550-51 (¶287). Riho thus lacks any load “difference” to “reduce,” taking it outside claim 7.

*Second*, Riho does not teach that the “number of array dies” in each group is “***selected . . . so as to reduce*** a difference between” the loads. Appx192(24:37-47) (emphasis added). In Riho, all TSVs concededly are connected to an “equal number of dies.” Appx125. There is no possibility in Riho of connecting more or fewer array dies to a given TSV to reduce a difference in load. *See* Appx6551 (¶287).

The Board asserted that Riho taught reducing the “‘***skew*** between’” die interconnects, which the Board took to refer to reducing differences in “TSV load.” Appx124 (emphasis added). But skew—data signals sent at the same time on

different wires arriving at different times due to differences in the wires' lengths—is different from load. Appx4467. More importantly, Riho's solution to skew is fundamentally different from Netlist's invention. Riho teaches minimizing skew (timing differences) by “ma[king]” TSVs of “equal length.” Appx125. The invention, by contrast, compensates for differences in loads from TSVs having different lengths by adjusting the number of dies connected to each. *See* pp. 9-11, *supra*. Put differently, Riho **avoids** one problem resulting from TSVs of different lengths by mandating equal-length TSVs. The '060 patent, by contrast, **solves** a (different) problem by adjusting the number of dies connected to each TSV to reduce load differences. Because there is only one possible outcome under the proper construction, reversal rather than remand is required. *See In re Hodges*, 882 F.3d 1107, 1113 (Fed. Cir. 2018).

2. Nor does Riho teach that the “number of array dies” in a group “are selected *in consideration of a load of the first [and second] die interconnect[s]*,” as claim 7 requires. Appx192(24:37-47) (emphasis added). The Board pointed to a passage about Riho's chip-selection “‘groups’”—a concept where the memory dies are divided into two groups, and only one group is accessed at a time while data “ports” on the other group's dies are “turned off.” Appx124; Appx3688-89(¶¶ 100-114); Appx6510-12(¶¶ 225-229). According to Riho, that “‘reduce[s] . . . load’” on the die “‘interconnections.’” Appx124 (quoting Appx3681(¶ 13)). But claim 7

requires *considering* the load of each “die interconnect”—the load of the TSV itself—in *selecting the number of dies* to connect to each. The claim does not encompass a chip-selection feature that “reduce[s] [the] load” contributed by *attached array dies* by selectively enabling one group of dies at a time. Appx3688(¶103). The Board insisted Riho taught “chip grouping . . . based on TSV load considerations,” Appx124, but nowhere explained how Riho taught selecting *the number of dies in each group* considering the *load of the corresponding TSV*. Reversal is required for that reason as well. *See Hodges*, 882 F.3d at 1113.

**C. The Board’s Determination that Skilled Artisans Would Have Combined Riho with Kim Is Unreasoned and Unsupported**

Because claim 7 depends from claim 1, the Board was required to combine Riho (for claim 7’s added limitations) with Kim/Rajan (for claim 1’s limitations) to achieve the invention. Appx126. The Board, however, failed to explain why skilled artisans would have pursued that combination, or expected success in doing so. *See TQ Delta*, 942 F.3d at 1360; *Pers. Web Techs., LLC v. Apple, Inc.*, 848 F.3d 987, 994 (Fed. Cir. 2017). Indeed, the Board offered no independent analysis of motivation to combine.

The Petition argued that skilled artisans would have modified Kim—where the TSVs have *different* lengths—to use Riho’s *equal-length* TSVs to “reduce the load” and “‘compensate for phase variations.’” Appx121. The Board, however, never stated that it agreed with Petitioner’s contention, much less explained *why* it

agreed. Where the Board chooses to “credit[ ]” a party’s argument, it must explain “why” it accepts the argument “‘as its own analysis.’” *Nuvasive*, 842 F.3d at 1383-84. It cannot simply restate an argument and say it agrees without explanation—much less not even state its agreement. *Icon Health & Fitness, Inc. v. Strava, Inc.*, 849 F.3d 1034, 1047 (Fed. Cir. 2017).

The Board also failed to address Netlist’s argument that, for reasons of “cost,” skilled artisans would have lacked motivation to combine. Appx11442; Appx11472-73. As shown in Kim’s Figure 5, a TSV in Kim originates in the bottom layer and goes up through the stack of dies only as far as the single array die to which it is connected. Appx3627(Fig. 5). The remaining dies—those higher up in the stack—thus do not require a hole for that TSV, saving space. Modifying Kim to use Riho’s TSVs, which extend all the way through the stack, would increase the space used for holes, increasing cost. The Board nowhere addressed that argument. For that reason, too, its decision cannot be sustained. *See VirnetX Inc. v. Cisco Sys., Inc.*, 776 F. App’x 698, 702-03 (Fed. Cir. 2019). And the Board’s decision does not even mention expectation of success. It again quotes the Petition without further explanation. Appx121. That, too, requires vacatur. *Nuvasive*, 842 F.3d at 1383-84.

Lacking any reason why skilled artisans would modify Kim to use Riho’s *equal-length* TSVs, *see* Appx11442; Appx11472, the Board asserted, in the alternative, that skilled artisans could use “Kim’s *unequal length* TSVs.” Appx125

(emphasis added). But that contradicts the Board’s premise that skilled artisans would have modified Kim using Riho’s *equal-length* TSVs to “reduce the load” and ““compensate for phase variations.’” Appx121. Keeping Kim’s unequal-length TSVs would provide neither of the benefits that supposedly would have motivated skilled artisans to combine Kim and Riho. Such self-contradictory reasoning violates the APA and cannot support the Board’s determination. *See Del. Dep’t of Natural Res. & Env’t Control v. EPA*, 785 F.3d 1, 15-16 (D.C. Cir. 2015).

Moreover, Riho *does not teach anything* about reducing load differences where the TSVs have different lengths. It teaches only using equal-length TSVs. *See* p. 61, *supra*. Riho adds nothing to the combination unless equal-length TSVs are used. Indeed, that shows why Riho does not teach claim 7. Riho does not *solve* the problem the ’060 patent solves—reducing load differences where the TSVs are of *different* lengths. It instead *avoids* the problem, by mandating that TSVs have the *same* length.

### **III. THE BOARD’S DECISION ON CLAIM 15 OF THE ’060 PATENT CANNOT BE SUSTAINED**

Dependent claim 15 of the ’060 patent recites the memory module of claim 11 (which is materially the same as claim 1), where the first data conduit has a “first driver” with a “first driver size,” and the second data conduit has a “second driver” with a “second driver size.” Appx193(25:35-42). It was undisputed that claim 15 requires using multiple drivers of differing sizes, one for each data conduit.

Appx132; Appx134. Limitation 1[e] of the '160 patent likewise requires multiple drivers of “different” sizes. Appx169(23:59-24:3).

The Board found those limitations taught by Wyman, with the remaining limitations taught by Kim and Rajan. Appx126. To find obviousness, however, the Board must “explicit[ly]” identify the “reason” skilled artisans would have combined the references “in the fashion claimed by the patent.” *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398, 418 (2007). A prior-art reference’s “preferences are relevant to . . . whether a skilled artisan would be motivated to combine” them. *Polaris*, 882 F.3d at 1069. There is no motivation to combine where “doing so would necessarily involve altering the inventive concept” of the prior art. *Chemours*, 4 F.4th at 1376.

The Board’s proposed Kim/Rajan/Wyman combination runs afoul of that precedent. Wyman expressly teaches a single adjustable driver circuit—not multiple driver circuits of differing sizes. Appx3703(1:55-65). And it teaches using that single driver circuit to drive multiple TSVs—not assigning each die interconnect its own separate driver. Appx3703(1:55-65). The Board found that Wyman “teaches using different amounts of signal drive for different lengths of TSVs.” Appx45. But using different *amounts of drive power* is not the same as using multiple *drivers of different sizes*. Wyman teaches using a *single* driver with a “selectable” power level, Appx3703(1:36-37), not the *multiple, separate*, differently sized drivers recited by the claims. See Appx193(25:35-42); Appx169(23:59-24:3).

Each of Wyman’s embodiments involves “a single driver with sub-driver components that can be selectively combined to drive loads falling within given ranges.” Appx6501(¶212); *see* Appx3704(3:19-22) (single driver “made up of multiple cascaded and parallel transistor circuits”). Wyman’s adjustable single-driver circuit—as opposed to separate physical drivers—enables designers to, “depending on the necessary signal strength,” select only “the corresponding tap which supplies no less than the amount necessary for the particular element of device.” Appx3703-04(2:65-3:1). It therefore allows designers to use the same single-driver circuit for applications that require “high drive capability,” as well as applications that require “very low drive capability.” Appx3706(7:26-32).

Wyman explicitly contrasts its single-driver-circuit approach with the multi-driver approach proposed by the ’060 and ’160 patents, noting that its single-driver circuit “provides, as an advantage, a level of design flexibility because it enables a designer to adjust and minimize power consumption . . . by specifically tailoring driver outputs according to the particular element or device’s requirements, specification or application.” Appx3704(4:46-51). Netlist’s expert, Dr. Brogioli, explained that skilled artisans reading Wyman “would not have even used *different drivers*, let alone different driver *sizes*, to drive [the] TSV1 and TSV2 as asserted” in the proposed Kim/Rajan combination. Appx6501(¶212) (emphasis added). Wyman’s disclosure instead suggested use of Wyman’s actual, more “‘efficien[t]’”

invention—a single driver circuit with selectable power levels—not **multiple**, **separate** drivers for each TSV. Appx6501-02(¶213).

The Board’s response simply changed the subject. Even if “using a single driver may be an acceptable option and **may be even a better option** than having separate drivers,” the Board declared, the Petition’s proposed combination “need not be the best option.” Appx132 (emphasis added). But assessing motivation to combine requires considering a reference’s teachings regarding its relative benefits and drawbacks. The Board cannot consider prior-art references “in less than their entirety” or disregard “disclosures in the references that diverge from and teach away from the invention at hand.” *W.L. Gore & Assocs., Inc. v. Garlock, Inc.*, 721 F.2d 1540, 1550 (Fed. Cir. 1983). “‘The benefits, **both lost and gained**,’” of a proposed combination, must be weighed “‘against one another.’” *Henny Penny Corp. v. Frymaster LLC*, 938 F.3d 1324, 1332 (Fed. Cir. 2019). The Board simply dismissed the asserted benefits of Wyman’s single driver circuit, which the proposed combination would destroy. While the Board asserted that “common sense” would lead skilled artisans to adopt “physically separate drivers,” it did not square that observation with Wyman’s explicit teachings. Appx132-33. And while the Board invoked Wyman’s “alternative configuration in which taps are not used,” Appx133, Netlist explained—and the Board did not dispute—that those embodiments also involve only a “single driver.” Appx786-87. The Board’s failure to consider

Wyman’s explicit teachings violates the APA’s reasoned-decisionmaking requirement. *See Virtek*, 97 F.4th at 886-87; *Icon Health*, 849 F.3d at 1047.

#### **IV. THE BOARD’S UNREASONED ANALYSIS OF OTHER CLAIMS AND LIMITATIONS REQUIRES REMAND**

In an IPR, “the petitioner” has “the burden of proving . . . unpatentability.” 35 U.S.C. § 316(e). The patent owner has “no obligation to raise any objection” to the petition “at all.” *Fanduel, Inc. v. Interactive Games LLC*, 966 F.3d 1334, 1341-42 (Fed. Cir. 2020). The Board thus may not accept an argument simply because the patent owner “‘d[id] not challenge’” it. *Icon Health*, 849 F.3d at 1047. To overturn its own prior issuance of a patent, the agency must “‘examine the relevant data and articulate a satisfactory explanation for its action including a rational connection between the facts found and the choice made.’” *Id.* at 1043. To the extent the Board chooses to “‘[credit] a party’s argument,’” it must explain “‘why’” it accepts that argument “‘as its own analysis.’” *Nuvasive*, 842 F.3d at 1383.

The Board repeatedly violated those requirements. Time and again it found limitations and claims of the ’060 and ’160 patents obvious based in whole or in part on Petitioner’s mere “contention[s]” coupled with an assertion Netlist “does not dispute” them—specifically, as to limitations 1[a], 1[b], and 1[e] and claims 2-5, 8-14, and 16-24 of the ’060 patent; and limitations 1[a] and 1[b] and claims 2-14 and 17-20 of the ’160 patent. *E.g.*, Appx15-18, Appx27-28, Appx51-57; Appx83-86, Appx94-96, Appx111, Appx114-20, Appx122, Appx135. That is exactly the

shortcut this Court rejected in *Icon Health*: While the Board can credit a party's argument as part of a reasoned explanation, it "still must 'explain[] why [it] accepts the prevailing argument.'" 849 F.3d at 1047. Here, as in *Icon Health*, the Board "failed to do so." *Id.* The Board, moreover, repeatedly failed to address Netlist's arguments on myriad issues, as detailed above. That too requires remand. *VirnetX*, 776 F. App'x at 702-03.

### **CONCLUSION**

The final written decision should be reversed or vacated.

February 4, 2025

Elizabeth Kathleen Clarke  
Bonnie K. St. Charles  
MOLOLAMKEN LLP  
300 N. LaSalle Street, Suite 5350  
Chicago, IL 60654  
(312) 450-6700

Sara Margolis  
Catherine Martinez  
MOLOLAMKEN LLP  
430 Park Avenue  
New York, NY 10022  
(212) 607-8160

Respectfully submitted,

/s/ Jeffrey A. Lamken

Jeffrey A. Lamken  
Rayiner Hashem  
Jennifer Elizabeth Fischell  
Lidiya Mishchenko  
Kayvon Ghayoumi  
MOLOLAMKEN LLP  
The Watergate, Suite 500  
600 New Hampshire Avenue, N.W.  
Washington, D.C. 20037  
(202) 556-2000  
jlamken@mololamken.com

*Counsel for Appellant Netlist, Inc.*

## **ADDENDUM**

## **ADDENDUM TABLE OF CONTENTS**

|                                                                                                       | <u>Page</u> |
|-------------------------------------------------------------------------------------------------------|-------------|
| Final Written Decision in IPR2022-01427 (Apr. 1, 2024)<br>(U.S. Patent No. 9,318,160) (Paper 48)..... | Appx1       |
| Final Written Decision in IPR2022-01428 (Apr. 1, 2024)<br>(U.S. Patent No. 8,787,060) (Paper 48)..... | Appx69      |
| Decision Denying Rehearing (June 17, 2024)<br>(IPR2022-01427 Paper 52, IPR2022-01428 Paper 52) .....  | Appx144     |
| U.S. Patent No. 9,318,160.....                                                                        | Appx147     |
| U.S. Patent No. 8,787,060.....                                                                        | Appx171     |

UNITED STATES PATENT AND TRADEMARK OFFICE

---

BEFORE THE PATENT TRIAL AND APPEAL BOARD

---

SAMSUNG ELECTRONICS CO., LTD, MICRON TECHNOLOGY, INC.,  
MICRON SEMICONDUCTOR PRODUCTS, INC., and  
MICRON TECHNOLOGY TEXAS LLC,<sup>1</sup>  
Petitioner,

v.

NETLIST, INC.,  
Patent Owner.

---

IPR2022-01427  
Patent 9,318,160 B2

---

Before JON M. JURGOVAN, DANIEL J. GALLIGAN, and  
SHEILA F. McSHANE, *Administrative Patent Judges*.

GALLIGAN, *Administrative Patent Judge*.

JUDGMENT  
Final Written Decision  
Determining All Challenged Claims Unpatentable  
Dismissing Petitioner's Motion to Exclude  
*35 U.S.C. § 318(a)*

---

<sup>1</sup> Micron Technology, Inc., Micron Semiconductor Products, Inc., and Micron Technology Texas LLC filed a motion for joinder and a petition in IPR2023-00883 and have been joined as petitioners to this proceeding.

## I. INTRODUCTION

In this *inter partes* review, Samsung Electronics Co., Ltd. (“Samsung”), Micron Technology, Inc., Micron Semiconductor Products, Inc., and Micron Technology Texas LLC (collectively “Petitioner”) challenge the patentability of claims 1–20 of U.S. Patent No. 9,318,160 B2 (“the ’160 patent,” Ex. 1001), which is assigned to Netlist, Inc. (“Patent Owner”).

We have jurisdiction under 35 U.S.C. § 6. This Final Written Decision, issued pursuant to 35 U.S.C. § 318(a), addresses issues and arguments raised during the trial in this *inter partes* review. For the reasons discussed below, we determine that Petitioner has proven by a preponderance of the evidence that claims 1–20 of the ’160 patent are unpatentable. *See* 35 U.S.C. § 316(e) (2018) (“In an *inter partes* review instituted under this chapter, the petitioner shall have the burden of proving a proposition of unpatentability by a preponderance of the evidence.”).

### A. Procedural History

Samsung filed a Petition (Paper 1, “Pet.”) challenging claims 1–20 of the ’160 patent on the following grounds:

| Claims Challenged | 35 U.S.C. §         | References/Basis                                         |
|-------------------|---------------------|----------------------------------------------------------|
| 1–20              | 103(a) <sup>2</sup> | Kim, <sup>3</sup> Rajan, <sup>4</sup> Wyman <sup>5</sup> |

---

<sup>2</sup> The Leahy-Smith America Invents Act (“AIA”), Pub. L. No. 112-29, 125 Stat. 284, 287–88 (2011), amended 35 U.S.C. § 103 effective March 16, 2013. Petitioner applies the pre-AIA version of § 103 based on an effective filing date before March 16, 2013. Pet. 3–4.

<sup>3</sup> US 2011/0103156 A1, published May 5, 2011 (Ex. 1014).

<sup>4</sup> US 8,041,881 B2, issued Oct. 18, 2011 (Ex. 1015).

<sup>5</sup> US 7,969,192 B2, issued June 28, 2011 (Ex. 1017).

| Claims Challenged | 35 U.S.C. § | References/Basis                            |
|-------------------|-------------|---------------------------------------------|
| 1–20              | 103(a)      | Riho, <sup>6</sup> Rajan, Riho <sup>7</sup> |

Pet. 3. Patent Owner filed a Preliminary Response. Paper 6. With Board authorization (Ex. 3001), Samsung filed a preliminary reply (Paper 9) to the Preliminary Response, and Patent Owner filed a preliminary sur-reply (Paper 10). Trial was instituted on the asserted grounds of unpatentability. Paper 13 (“Inst. Dec.”) at 28.

During the trial, Patent Owner filed a Response (Paper 20, “PO Resp.”), and Samsung filed a Reply (Paper 24, “Pet. Reply”). After Samsung’s Reply, Micron Technology, Inc., Micron Semiconductor Products, Inc., and Micron Technology Texas LLC were joined as petitioners to this proceeding based on a petition and a motion for joinder in IPR2023-00883. Paper 25. After this joinder, Patent Owner filed a Sur-reply (Paper 28, “PO Sur-reply”).

Petitioner filed a motion to exclude certain evidence. Paper 34. Patent Owner opposed the motion (Paper 35), and Petitioner filed a reply in support of the motion (Paper 36).

An oral hearing was held on January 11, 2024, a transcript of which appears in the record. Paper 46 (“Tr.”).

Petitioner relies on testimony from Andrew Wolfe, Ph.D. Ex. 1003. Patent Owner relies on testimony from Michael C. Brogioli, Ph.D. Ex. 2023. The parties have entered in the record transcripts of the depositions of these declarants. Exs. 2025 (Wolfe Deposition), 1052 (Brogioli Deposition).

---

<sup>6</sup> US 2011/0026293 A1, published Feb. 3, 2011 (Ex. 1016).

<sup>7</sup> US 2010/0195364 A1, published Aug. 5, 2010 (Ex. 1018).

### *B. Real Parties in Interest*

The identified real parties in interest on the petitioner side are the following: Samsung, Samsung Semiconductor, Inc., Micron Technology, Inc., Micron Semiconductor Products, Inc., and Micron Technology Texas LLC. Pet. 1; IPR2023-00883, Paper 2 at 1.

Patent Owner identifies itself as the real party in interest. Paper 3 at 1.

### *C. Related Matters*

As required by 37 C.F.R. § 42.8(b)(2), the parties identify various related matters. Pet. 1; Paper 3 at 1; IPR2023-00883, Paper 2 at 1. We are issuing concurrently a final decision in IPR2022-01428 involving related U.S. Patent 8,787,060 No. B2 (“the ’060 patent”).

### *D. The ’160 Patent and Illustrative Claim*

The ’160 patent relates to computer memory devices and, more specifically, to reducing the load of drivers in memory. Ex. 1001, 1:20–24. As background, the ’160 patent describes an existing memory package with reference to Figure 1A, reproduced at right, which shows memory package 100 with control die 130 and three array dies 110. Ex. 1001, 1:32–46. Control die 130 has driver 134, which drives data signals from control die 130 to each array die via interconnect 142, and driver 140, which drives command and address signals to each array die. Ex. 1001, 1:37–44.

The ’160 patent explains that “a load exists on each of the drivers 134,




FIG. 1A

140 . . . by virtue of the drivers being in electrical communication with the corresponding die interconnects and the corresponding circuitry of the array dies.” Ex. 1001, 2:10–13. The ’160 patent discloses that, “to drive a signal along a die interconnect, a driver typically must be large enough to overcome the load on the driver” and that “a larger driver not only consumes more space on the control die, but also consumes more power.” Ex. 1001, 2:13–17.

The ’160 patent states that driver size and power consumption can be reduced “by increasing the number of die interconnects and reducing the number of array dies that are in electrical communication with each die interconnect.” Ex. 1001, 4:25–29. The ’160 patent illustrates an exemplary configuration in Figure 2, reproduced at right, which shows memory package 200 with control die 230 and four array dies 210a–210d. Ex. 1001, 5:15–16. Control die 230 is connected to array dies 210a and 210b by die interconnect 220a, as shown by the darkened circles.

Ex. 1001, 5:66–6:4. Die interconnect 220b connects control die 230 to array dies 210c and 210d, as shown by the darkened circles, but die interconnect 220b is not electrically connected to array dies 210a and 210b, as shown by




FIG. 2

the unfilled circles. Ex. 1001, 6:12–26. The ’160 patent also states that “[e]xamples of die interconnects include, but are not limited to, through-silicon vias (TSV), conducting rods, wire bonds, and pins.” Ex. 1001, 5:54–56.

Claim 1 is illustrative and is reproduced below with Petitioner’s claim element identifiers in brackets.

1. [1.a] A memory package, comprising:
  - [1.b] data terminals and control terminals;
  - [1.c] stacked array dies including a first group of array dies and a second group of at least one array die;
  - [1.d.1] 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, [1.d.2] 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
  - [1.e.1] a control die comprising [1.e.2] first data conduits between the first die interconnects and the data terminals, and [1.e.3] second data conduits between the second die interconnects and the data terminals, [1.e.4] 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, [1.e.5] 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.

## II. ANALYSIS

### A. *Principles of Law*

A patent claim is unpatentable under 35 U.S.C. § 103(a) if the differences between the claimed subject matter and the prior art are such that the subject matter, as a whole, would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. *KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 406 (2007). The question of obviousness is resolved on the basis of underlying factual determinations including (1) the scope and content of the prior art; (2) any differences between the claimed subject matter and the prior art; (3) the level of ordinary skill in the art; and (4) any secondary considerations, if in evidence. *Graham v. John Deere Co. of Kan. City*, 383 U.S. 1, 17–18 (1966).

### B. *Level of Ordinary Skill in the Art*

Citing the testimony of its declarant, Dr. Wolfe, Petitioner contends that a person of ordinary skill in the art “would have had an advanced degree in electrical or computer engineering, or a related field, and two years working or studying in the field of design or development of memory systems, or a bachelor’s degree in such engineering disciplines and at least three years working in the field.” Pet. 5 (citing Ex. 1003 ¶ 60). Petitioner also asserts that “[a]dditional training can substitute for educational or research experience, and vice versa,” and Petitioner identifies particular technology with which a person of ordinary skill in the art would have been familiar, including JEDEC (Joint Electron Devices Engineering Council) industry standards, DRAM (dynamic random access memory) and SDRAM

(synchronous DRAM) memory modules, and “the structure and operation of circuitry used in stacked memory devices.” Pet. 5 (citing Ex. 1003 ¶ 60).

Patent Owner states that, “[f]or purposes of this proceeding, Patent Owner applies the education and work-experience level specified on Petition, page 5.” PO Resp. 6. Patent Owner then contends that “Petitioner does not allege that a [person of ordinary skill in the art] would have been familiar with designs and operations of non-DRAM devices or non-JEDEC specifications.” PO Resp. 6. Patent Owner contends, therefore, that “under Petitioner’s unpatentability theories” involving non-DRAM devices, persons of ordinary skill in the art “would be engaged in the project of making changes to memory devices that they have no experience with.” PO Resp. 6 (citing Ex. 2023 ¶ 30).

We do not agree with Patent Owner’s assertions regarding non-DRAM devices. The obviousness inquiry under 35 U.S.C. § 103(a) requires us to determine

if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains.

As Patent Owner confirmed during oral argument, the term “array dies” in the claims encompasses non-DRAMs. Tr. 53:4–7. Because the subject matter pertains to configurations involving non-DRAM devices, the “person having ordinary skill in the art,” by definition under the statute, would have had knowledge of non-DRAM devices. To the extent Petitioner’s proposed skill level, on which the parties agree, does not make that clear, we do.

Thus, we accept the uncontested assessment offered by Petitioner with the exception of the qualifier “at least,” which introduces vagueness as to the

amount of experience. We also determine that a person of ordinary skill in the art would have been familiar with designs and operations of non-DRAM devices for the reasons discussed above.

### *C. Claim Construction*

We interpret claim terms using “the same claim construction standard that would be used to construe the claim in a civil action under 35 U.S.C. 282(b).” 37 C.F.R. § 42.100(b).

#### *1. Array Die*

The challenged claims recite the term “array dies.” In related litigation, a district court construed this term based on the following argument made by the applicant during prosecution of the ’060 patent regarding U.S. Patent Application Publication No. 2008/0025137 A1 (Ex. 1011 (“Rajan 137”)): “Rajan[137] does not disclose ‘a plurality of stacked array dies.’ Rajan[137] merely stacks DRAM circuits 206A–D, which are different from array dies.” Ex. 2004, 31–32 (quoting Ex. 1002 (’060 patent prosecution history), 465 (January 13, 2014 Amendment at 10)). The district court found that Patent Owner “did not explain the structural difference between array dies and DRAM circuits, or even how they might be stacked differently” but, nonetheless, found that Patent Owner “*structurally* distinguished ‘stacked DRAM circuits’ from ‘stacked array dies’ to obtain the patent.” Ex. 2004, 32. Based on this prosecution history, the district court adopted the defendants’ proposed construction that “array die” means “array die that is different from a DRAM circuit.” Ex. 2004, 32. Samsung is a defendant in that litigation.

Patent Owner argues that it “agreed to be bound by this construction and proceeded with the jury trial on that construction.” PO Resp. 17 (citing

Exs. 2010, 2011). Yet, during the trial in this proceeding, Patent Owner qualifies the district court's construction to suggest that the term "array dies" means dies that are different from a particular kind of DRAM circuit. *See, e.g.*, PO Resp. 47 ("[T]he Petition contains no analysis on how or why Kim's memory chips are different from Rajan137's DRAM circuits."); *see also* Inst. Dec. 10–11 (addressing similar pre-institution arguments); Tr. 58:11–15 (Patent Owner's counsel's assertion that Rajan137 "describes, for instance, monolithic memory circuits may take the form of, and then it lists an example. And it lists DRAM as an example of a monolithic memory circuit. So, we take DRAM circuit to be a monolithic memory circuit that holds state while powered."). The district court's construction, which Patent Owner asserts it adopted, does not contain any such qualification.

In our analysis below, we apply the district court's construction, and we find that the combination of Kim, Rajan, and Wyman teaches array dies that are not DRAMs and, therefore, are "different from a DRAM circuit." Therefore, we need not address any further qualifications of the district court's construction or construe the term "DRAM circuit" to determine what Patent Owner actually disclaimed, if anything. *See Omega Eng'g, Inc. v. Raytek Corp.*, 334 F.3d 1314, 1324 (Fed. Cir. 2003) ("[P]rosecution disclaimer promotes the public notice function of the intrinsic evidence and protects the public's reliance on definitive statements made during prosecution."); *see also Genuine Enabling Tech. LLC v. Nintendo Co.*, 29 F.4th 1365, 1374 (Fed. Cir. 2022) ("If the challenged statements are ambiguous or amenable to multiple reasonable interpretations, prosecution disclaimer is not established."). Based on our findings as to non-DRAM

dies, we also need not determine whether the term “array die” should be construed to also encompass DRAM circuits.

## 2. *Chip Select*

Claims 15 and 16 recite “first chip select signals” and “second chip select signals.” Patent Owner argues that the district court construed the term “chip select signals” to “exclude[] situations in which a chip select signal could enable multiple array dies at once” and contends that we should adopt this construction. PO Resp. 18 (citing Ex. 2004, 33–34). The district court’s construction is actually for the “chip select signal” term recited in various claims of the ’060 patent. Ex. 2004, 33–34, 36. As explained in the Institution Decision, we agree with this interpretation as applied to the “second chip select signals” term of claims 15 and 16 of the ’160 patent because these claims recite “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*” (emphasis added). *See* Inst. Dec. 11–12. Claims 15 and 16 recite that there are fewer “first chip selection signals” than “second chip selection signals,” and, therefore, each of the “first chip selection signals” appears to apply to multiple array dies, which makes sense in the context of the “rank multiplication” of claims 15 and 16. *See* Inst. Dec. 12.

We apply this interpretation of the claim language below in addressing the patentability of claims 15 and 16.

## 3. *Group of Array Dies*

Petitioner argues that the phrase “a first group of array dies,” which is recited in the independent claims, can be just one die. Pet. Reply 8–9; *see also* Pet. 40 (arguing that multiple chips would have been obvious “[i]nsofar

as one might argue that ‘first group of array dies’ requires multiple slave chips coupled to” one through-silicon via (TSV)).<sup>8</sup> Patent Owner disagrees and asserts that the first group must have plural dies. PO Resp. 20. We need not resolve this dispute because, as explained below in § II.D.2.d, we find persuasive Petitioner’s contentions that the combination of Kim, Rajan, and Wyman teaches a first group that has multiple array dies.

#### 4. *Driver Size*

Various claims recite the term “driver size.” Neither party argues for an express construction of this term, but Patent Owner notes that a district court in related litigation construed “driver size” as “driver physical size.” PO Sur-reply 18; *see* Ex. 1066 (district court claim construction order), 13, 43. As explained below in § II.D.2.e, we find that the combination of Kim, Rajan, and Wyman teaches the subject matter under this construction, and we need not address the construction of this term further.

#### 5. *Remaining Terms*

We determine that no other terms or phrases require express construction. *See Realtime Data, LLC v. Iancu*, 912 F.3d 1368, 1375 (Fed. Cir. 2019) (“The Board is required to construe ‘only those terms . . . that are in controversy, and only to the extent necessary to resolve the controversy.’” (quoting *Vivid Techs., Inc. v. Am. Sci. & Eng’g, Inc.*, 200 F.3d 795, 803 (Fed. Cir. 1999))).

---


<sup>8</sup> We omit the emphasis added by Petitioner for reference names and claim language when quoting Petitioner’s arguments.

*D. Obviousness over Kim, Rajan, and Wyman  
(Claims 1–20)*

Petitioner asserts that claims 1–20 are unpatentable as obvious over the combined teachings of Kim, Rajan, and Wyman. Pet. 3, 23–80. Patent Owner opposes. PO Resp. 20–53; PO Sur-reply 1–20.

*1. Overview of the Prior Art*

Kim discloses a semiconductor memory apparatus including stacked chips, as depicted in Figure 5 below.



Kim's Figure 5 above shows a memory including a plurality of stacked chips, including main chip C0 and first and second slave chips C1 and C2, connected by through-silicon vias (TSVs). Ex. 1014 ¶¶ 47–48.

Rajan discloses using an interface circuit to emulate characteristics of particular memory even though the physical memory may not have those



interconnects”; TSV2 (teal) from C0 to C2 as “second die interconnects”; box labeled “DQ” (gray) as “data terminals”; line (pink) from data input/output section 1000 to TSV1 as “first data conduits”; and line (rose) from data input/output section 1000 to TSV2 as “second data conduits.” Pet. 34–53.

Petitioner argues that a person of ordinary skill in the art would have been motivated to make Kim’s memory package compatible with JEDEC standards based on Rajan’s teachings. Pet. 26–29. Petitioner also argues that a person of ordinary skill in the art would have been motivated to have multiple memory chips share a TSV because this was a known option and “would not require creating new TSVs (which would add space and circuitry) and would allow emulating the JEDEC standard required by external devices.” Pet. 30–31. Petitioner relies on Wyman to teach using different driver sizes depending on load. Pet. 32–33.

*b) Preamble (1.a)*

Petitioner argues that Kim’s Figure 5 depicts a “memory package,” as recited in the preamble. Pet. 35 (citing Ex. 1014 ¶¶ 46–47; Ex. 1003 ¶¶ 174–179). Patent Owner does not dispute this contention. We agree with Petitioner, and we find, that Kim’s disclosure of semiconductor memory apparatus 3 in Figure 5 teaches a “memory package.” In view of this finding, we need not decide whether the preamble is limiting.

*c) Data and control terminals (1.b)*

For claim 1’s “data terminals and control terminals,” Petitioner argues that Kim’s Figure 5 shows data terminals as “DQ,” that Kim’s Figure 2 shows control terminals as CS0 and CS1, and that Kim’s memory package would have control terminals to accept control signals in a

JEDEC-compliant implementation. Pet. 35–38. Patent Owner does not dispute that the combination of Kim and Rajan teaches data terminals and control terminals, but Patent Owner does dispute the JEDEC aspect of Petitioner’s combination. We agree with Petitioner, and we find, that Kim teaches a memory package having data terminals and control terminals, and we further address the parties’ dispute about JEDEC below.

*d) Stacked array dies (1.c) and die interconnects (1.d.1, 1.d.2)*

Claim 1 recites “stacked array dies including a first group of array dies and a second group of at least one array die” (1.c). Claim 1 further recites that the memory package has “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,” (1.d.1) and “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” (1.d.2).

As noted above in reference to Petitioner’s annotated version of Kim’s Figure 5, Petitioner identifies Kim’s chip C1 as “a first group of array dies” and Kim’s chip C2 as “a second group of at least one array die.” Pet. 39 (citing Ex. 1014 ¶¶ 26, 48, Fig. 5; Ex. 1003 ¶¶ 197–214). Petitioner argues that having multiple dies in the first group all sharing the same die interconnect would have been obvious. Pet. 40–42 (citing Ex. 1014 ¶¶ 48, 50; Ex. 1024, 9:14–18, Fig. 5; Ex. 1015, Fig. 4; Ex. 1003 ¶¶ 204–212).<sup>9</sup>

---

<sup>9</sup> Petitioner also argues that “a first group of array dies” can be just one die. Pet. Reply 8–9; *see also* Pet. 40. We do not rely on this argument and, therefore, need not resolve the parties’ dispute regarding whether this phrase encompasses a single die.

Petitioner argues that the annotated version of Rajan's Figure 4 below shows "multiple stacked memory chips connected to the same data bus."

Pet. 41–42.




FIG. 4

Rajan's Figure 4 "shows a buffered stack of DRAM circuits having one address, control, and clock bus and two data busses." Ex. 1015, 2:6–7. In the annotated figure above, Petitioner maps depicted components to elements of claim 1 as follows: DRAMs 417C and 417D (yellow) as "first group of array dies"; DRAMs 417A and 417B (orange) as "second group of at least one array die"; buffer chip 413 (green) as "control die"; data bus (light green) from buffer chip 413 to DRAMs 417C and 417D as "first die interconnects"; and data bus from buffer chip 413 to DRAMs 417A and 417B as "second die interconnects." Pet. 42.

Petitioner argues that Kim and Rajan are analogous art. Pet. 26 (citing Ex. 1003 ¶¶ 150–151). Patent Owner does not dispute that Kim and Rajan are analogous art to the '160 patent. See PO Resp. We agree with Petitioner that Kim and Rajan are analogous art because they are both in the same field

of endeavor as the '160 patent—memory packages and in particular those involving stacked memory. Ex. 1001, 1:20–21 (“The present disclosure relates to memory devices and memory modules.”), claim 1 (“memory package” with “stacked array dies”); Ex. 1014 ¶ 3 (“Various embodiments of the present disclosure generally relate to a semiconductor memory apparatus . . . .”), Fig. 5 (stacked memory chips); Ex. 1015, 1:14–16 (“This invention relates generally to digital memory such as used in computers[], and more specifically to organization and design of memory modules such as DIMMs [(dual in-line memory modules)].”), Figs. 2–6 (stacked memories).

Petitioner argues that a person of ordinary skill in the art “would be motivated to create a package with an interface that complied with the well-known JEDEC standards, as taught by Rajan,” such that Kim’s memory package would be compatible with JEDEC. Pet. 26–28 (citing Ex. 1014 ¶ 38, Fig. 5; Ex. 1015, 3:52–54, 4:20–24, 5:36–43, 8:8–11, Figs. 4, 18; Ex. 1003 ¶¶ 147–149, 151–153; Ex. 1019, 12, Fig. 3).

Petitioner also argues that, based on Kim’s disclosure of using “any number of” chips (Ex. 1014 ¶ 48), a person of ordinary skill in the art “would have been motivated to look at Rajan for the details about adding more memory chips in the stack.” Pet. 28 (citing Ex. 1014 ¶¶ 48, 50; Ex. 1003 ¶¶ 154–156). Petitioner argues that a person of ordinary skill in the art “would have understood that there were a finite number of known ways to connect additional dies, including having a subset of dies sharing a through-silicon via (TSV),” and that “[s]haring a TSV among two or more dies was a known option, as confirmed by other references at the time.” Pet. 30 (citing Ex. 1003 ¶¶ 157–159; Ex. 1015, 5:36–43, Figs. 2–6;

Ex. 1025, Fig. 5). Petitioner also argues that a person of ordinary skill in the art “would have been motivated to connect additional memory dies to Kim’s TSV1 and TSV2 because it would not require creating new TSVs (which would add space and circuitry) and would allow emulating the JEDEC standard required by external devices.” Pet. 30 (citing Ex. 1003 ¶ 160; Ex. 1015, 3:27–30, 3:52–61). Petitioner further argues that “it was common at the time to have multiple memory dies sharing a single data bus” and “to use a ‘fork-in-the-road’ arrangement with two data buses for two ranks of memory devices . . . with the option of adding two additional ranks of memory devices (resulting in four ranks) to the existing data busses,” such that two dies share a bus. Pet. 31 (citing Ex. 1003 ¶¶ 161–162; Ex. 1026, Figs. 12, 13).

As to claim 1’s “electrical communication” recitations (1.d.1 and 1.d.2), Petitioner argues that Kim’s Figure 5 shows die interconnects (TSV1) connected from chip C0 to chip C1 and not extending to chip C2 and die interconnects (TSV2) connected from chip C0 to chip C2 and passing through chip C1 without being connected. Pet. 43 (citing Ex. 1014, Fig. 5; Ex. 1003 ¶ 220). Petitioner argues that each array die receives multiple data signals and that each signal would be on its own TSV, such that there would be multiple die interconnects to each array die. Pet. 42–43 (citing Ex. 1014 ¶ 28; Ex. 1016 ¶ 30; Ex. 1003 ¶ 220). And, as discussed above, in Petitioner’s proposed combination, multiple dies would share a TSV (Pet. 30–31) such that each group would be in electrical communication with the TSVs shared by that group and not other TSVs.

Below, we address Patent Owner’s arguments pertinent to recitations 1.c, 1.d.1, and 1.d.2.

*(1) Array dies*

As noted above in § II.C.1, a district court construed “array die” to mean “array die that is different from a DRAM circuit.” Ex. 2004, 31–32. Patent Owner argues that “the Petition contains no analysis on how or why Kim’s memory chips are different from Rajan137’s DRAM circuits.” PO Resp. 47. Patent Owner’s argument, therefore, reflects an attempt to limit the district court’s construction to exclude only certain types of DRAM circuits, similar to arguments Patent Owner made before institution. *See* Inst. Dec. 11 (“Patent Owner appears to argue that the district court’s construction is that the claimed array dies are different from Rajan 137’s DRAM circuits 206A–D, not all DRAM circuits.”).

As noted at institution, Kim does not even mention DRAM and thus is not limited to DRAM, and Rajan states that “any type of memory whatsoever” may be used in its disclosure and lists non-DRAM memories. Inst. Dec. 24 (quoting Ex. 1015, 15:3–9). Therefore, we find that the combination of Kim and Rajan teaches “array dies” under the district court’s construction and under Patent Owner’s interpretation of that construction.

Patent Owner further argues that,

if, in Petitioner’s combination, the stacked dies are non-DRAMs, neither the Petition nor the declaration explains why a [person of ordinary skill in the art] would be motivated to have the package emulate a DRAM interface or how the resulting structure would work given Petitioner’s theory that the interface circuit should provide a JEDEC-compliant interface to the host, citing to SDRAM standards.

PO Resp. 48 (citing Pet. 23–32). According to Patent Owner, “Petitioner has not presented any evidence that with non-DRAM circuits such as NANDs or SRAMs [(static RAM)], a [person of ordinary skill in the art]

would have or could have emulated the memory package as JEDEC-compliant DRAM devices.” PO Resp. 48–49 (citing Pet. 8–11, 23–34; Ex. 2023 ¶¶ 173–177); *see also* PO Resp. 50 (“Petitioner has not shown that a [person of ordinary skill in the art] would have been motivated to use non-DRAM dies in combination with an interface circuit that emulates a JEDEC-compliant interface to the host.”).

We disagree with Patent Owner’s arguments because the evidence in support of Petitioner’s position is the express disclosure of Rajan. As noted above, Petitioner argues that a person of ordinary skill in the art “would be motivated to create a package with an *interface* that complied with the well-known JEDEC standards, as taught by Rajan.” Pet. 26 (emphasis added). Rajan discloses an “interface circuit [that] may take the form of or incorporate, or be incorporated into, a register, an AMB [(advanced memory buffer)], a buffer, or the like, and may comply with Joint Electron Device Engineering Council (JEDEC) standards, and may have forwarding, storing, and/or buffering capabilities.” Ex. 1015, 4:20–24. Rajan explains that “the interface circuit presents to the system device an interface to emulated memory devices which differ in some aspect from the physical memory circuits which are actually present.” Ex. 1015, 3:24–27. With reference to Figure 1, Rajan discloses that the “physical memory circuits may be any type of memory circuits.” Ex. 1015, 2:59–60. Rajan further discloses that, “[a]lthough the embodiments described here show the stack consisting of multiple DRAM circuits, a stack may refer to any collection of memory circuits (e.g. DRAM circuits, flash memory circuits, or combinations of memory circuit technologies, etc.).” Ex. 1015, 4:51–55. Rajan still further discloses that the “physical memory circuits employed in practicing this

invention may be any type of memory whatsoever, such as: DRAM, DDR DRAM, DDR2 DRAM, DDR3 DRAM, SDRAM, QDR DRAM, DRDRAM, FPM DRAM, VDRAM, EDO DRAM, BEDO DRAM, MDRAM, SGRAM, MRAM, IRAM, NAND flash, NOR flash, PSRAM, wetware memory, etc.” Ex. 1015, 15:3–9. Thus, Rajan expressly discloses using non-DRAM dies as the memory circuits.

Patent Owner argues that “Rajan never suggests that its interface die would be emulating a different type of memory circuits from the backend memory.” PO Sur-reply 21 (citing Ex. 1015, 15:3–9). We disagree with this argument. As noted above, Rajan’s “interface circuit presents to the system device an interface to emulated memory devices which differ in some aspect from the physical memory circuits which are actually present” (Ex. 1015, 3:24–27), and we see nothing in Rajan that limits its teachings as Patent Owner’s argument suggests. Furthermore, Rajan discloses that the

memory subsystem includes a buffer chip 202 which presents the host system with emulated interface to emulated memory, and a plurality of physical memory circuits which, in the example shown, are DRAM chips 206A-D. In one embodiment, the DRAM chips are stacked, and the buffer chip is placed electrically between them and the host system. Although the embodiments described here show the stack consisting of multiple DRAM circuits, a stack may refer to any collection of memory circuits (e.g. DRAM circuits, flash memory circuits, or *combinations of memory circuit technologies*, etc.).

Ex. 1015, 4:45–55 (emphasis added). The fact that the memory stack may include “combinations of memory circuit technologies,” such as DRAM and flash, shows that the interface would emulate a memory different from at least one of the memory technologies in the combination when it presents as one type of memory to the host. Rajan also discloses that “the memory

circuits may be symmetrical, meaning each has the same capacity, *type*, speed, etc., while in other embodiments they may be asymmetrical.”

Ex. 1015, 2:62–67 (emphasis added). Thus, Rajan discloses that symmetrical circuits are of the same type, suggesting that asymmetrical circuits can be of different types behind the interface.

Patent Owner also asserts in a parenthetical that Rajan’s disclosure at column 15, lines 3–9 “has no link to 4:20-24 which is part of the description for Fig. 1, related to ‘DRAM circuits.’” PO Sur-reply 21. We disagree with this assertion. Rajan at column 4, lines 20–24 discloses that the interface circuit may comply with JEDEC, and this is the “interface circuit” of the disclosed invention. *See* Ex. 1015, 1:65–67 (“FIG. 1 shows a system coupled to multiple memory circuits and an interface circuit according to one embodiment of this invention.”). Thus, Rajan’s disclosure that the “physical memory circuits employed in practicing *this invention* may be any type of memory whatsoever,” including non-DRAMs, (Ex. 1015, 15:3–9 (emphasis added)) is linked to the interface circuit of the invention, which is described throughout Rajan and is described as being JEDEC-compliant in column 4, lines 20–24.

For all of the reasons discussed above, we find that Rajan teaches using non-DRAM memories in a JEDEC-compliant memory package and, therefore, that Rajan provides an express motivation to “to use non-DRAM dies in combination with an interface circuit that emulates a JEDEC-compliant interface to the host.” *See* PO Resp. 50.

We also disagree with Patent Owner’s argument that “Petitioner has not presented any evidence that with non-DRAM circuits such as NANDs or SRAMs, a [person of ordinary skill in the art] . . . *could have* emulated the

memory package as JEDEC-compliant DRAM devices.” *See* PO Resp. 48–49 (emphasis added) (citing Pet. 8–11, 23–34; Ex. 2023 ¶¶ 173–177). We disagree because the express disclosure of Rajan, discussed above, is evidence that a person of ordinary skill in the art could have used non-DRAM memories in a JEDEC-compliant DRAM package. Petitioner cites Rajan’s Figure 18 (Pet. 27), which shows an embodiment of an interface circuit, and Rajan’s accompanying disclosures provide more details as to how the interface circuit functions. Ex. 1015, 14:6–62.

Furthermore, we have considered the cited testimony from Dr. Brogioli, which focuses on how NAND flash and SRAM both differ from DRAM. Ex. 2023 ¶¶ 173–177. When considered in view of the disclosure of Rajan, the listed differences do not indicate or suggest difficulties beyond the skill of an ordinarily skilled artisan to implementing JEDEC-compliant interfaces with non-DRAM memories. For example, Dr. Brogioli testifies that NAND “signaling, electrical and physical interface are drastically different than that of DRAM memory devices, or DRAM memory modules that are JEDEC or JEDEC adjacent.” Ex. 2023 ¶ 175. But Rajan discloses that memory characteristics emulated by the interface circuit “may be electrical in nature, physical in nature, logical in nature, pertaining to a protocol, etc.” Ex. 1015, 3:36–38. As to signaling in particular, Rajan discloses that “[t]he interface circuit may emulate the number of signals, type of signals, duration of signal assertion, and so forth” and “may combine multiple signals to emulate another signal.” Ex. 1015, 3:48–51. Dr. Brogioli also testifies about “vastly different timing” between NAND and DRAM (Ex. 2023 ¶ 175), but Rajan discloses that “[a]n example of an emulated protocol characteristic might be a timing” (Ex. 1015, 3:42–43) and

that the interface circuit’s “emulation logic may, in various embodiments, alter a *timing*, value, latency, etc. of any of the address, control, clock, and/or data signals it sends to or receives from the system and/or the physical memory” (Ex. 1015, 14:55–59 (emphasis added)). Therefore, Dr. Brogioli’s testimony fails to account for Rajan’s disclosure. Dr. Brogioli’s testimony as to the “different signaling mechanism” that SRAM uses (Ex. 2023 ¶ 177) is similarly unpersuasive.

Furthermore, Dr. Brogioli’s testimony addressing NAND flash and SRAM differences fails to appreciate that Rajan is not so limited and lists other non-DRAM memories that can be used, including NOR flash. *See* Ex. 1015, 15:3–9.

For the foregoing reasons, we find that the combination of Kim and Rajan teaches “array dies.”

*(2) Electrical communication*

Patent Owner argues that Petitioner has not shown that the combination of Kim and Rajan teaches “first die interconnects in electrical communication with” multiple array dies. PO Resp. 20–44.

Patent Owner argues that Kim discloses that “each TSV is in communication with only a single chip.” PO Resp. 20 (citing Ex. 1014, Fig. 5; Ex. 2023 ¶ 84). Although Kim’s Figure 5 depicts TSV1 in communication with chip C1 and TSV2 in communication with chip C2, Kim discloses that “any number of . . . chips may be used” (Ex. 1014 ¶ 48), as Petitioner points out. *See* Pet. 28 (citing Ex. 1003 ¶¶ 154–156; Ex. 1014 ¶¶ 48, 50). Petitioner argues that this disclosure would have motivated a person of ordinary skill in the art “to look at Rajan for the details about adding more memory chips in the stack (resulting, e.g., in four chips . . . in


two groups . . .).” Pet. 28. We agree with Petitioner, and we find, that Kim’s disclosure suggests having more than the two chips depicted in Figure 5. *See* Ex. 1014 ¶¶ 48 (“Although only one main chip and two slave chips are shown, it is to be understood that any number of main and slave chips may be used.”), 50 (“While the semiconductor memory apparatuses having two ranks are explained with reference to FIGS. 2 and 5, a person having ordinary skill in the art will appreciate that the technical concept of the present invention can be applied to a semiconductor memory apparatus which are divided into three or more ranks.”); *see also* Ex. 1003 ¶ 155 (“A Skilled Artisan would have understood from this disclosure that Kim’s invention can include more slave chips (and more ranks) and that some of those additional slave chips proposed by Kim can be in the same group as chips C1 or C2 and connected to TSV1 or TSV2, respectively.”).

Petitioner contends that a person of ordinary skill in the art would have been motivated to implement a shared data bus for multiple memory chips as taught by Rajan (e.g., 415A and 415B, above) using, e.g., Kim’s TSV interconnects . . . , in part because a [person of ordinary skill in the art] would have understood that there were a finite number of known ways to connect additional dies, including having a subset of dies sharing a through-silicon via (TSV).

Pet. 30 (citing Ex. 1003 ¶¶ 157–159; Ex. 1015, 5:36–43, Figs. 2–6). Thus, Petitioner’s contention is that, given Kim’s express disclosure of having more than two memory chips, a person of ordinary skill would have looked to Rajan for details about adding more memory chips. Petitioner argues that Rajan teaches “a shared data bus for multiple memory chips” as depicted in Rajan’s Figure 4 (Pet. 30), and Patent Owner agrees. *See* PO Resp. 21 (“In

Rajan, Figure 4, Rajan’s DRAM circuits 417A and 417B share one data bus and DRAM circuits 417C and 417D share another data bus.”).

Patent Owner does not dispute Petitioner’s contention (Pet. 30) that there are a finite number of known ways to connect additional dies, one of which being to have dies share a TSV. *See* PO Resp. In support of its contention that “[s]haring a TSV among two or more dies was a known option,” as confirmed by other references, such as Foster, Petitioner provides the annotated figure below.



Pet. 30 (citing Ex. 1025 (“Foster”),<sup>10</sup> Fig. 5). Foster’s Figure 5 shows a cross-sectional view of die stack 20 including first substack 78 of dies 21 and 22 and second substack 80 of dies 23 and 24. Ex. 1025, 7:28–45. Petitioner identifies first substack 78 as a “first group of array dies” in yellow and second substack 80 as a “second group of at least one array die” in orange. Pet. 30. In Figure 5, TSVs 30 (identified by Petitioner in light

<sup>10</sup> US 8,258,619 B2, filed Nov. 12, 2009, issued Sept. 4, 2012.

green) connect to the first substack but do not extend to the second substack, and pass-through vias (PTVs) 29 (identified by Petitioner in teal) go through the first stack and connect to TSVs in the second substack. Ex. 1025, 7:45–54. Foster explains that PTVs are “conductive pathways through each die with no connections to any circuitry on the die” and that TSVs are “conductive pathways through the dies that also connect to electronic circuitry(36) on the die.” Ex. 1025, 2:53–56.

Patent Owner and Dr. Brogioli do not disagree with Petitioner’s assertion that Foster shows that “[s]haring a TSV among two or more dies was a known option.” *See* Pet. 30. Rather, they address different issues. Patent Owner argues that Petitioner cites Foster “for die grouping.” PO Resp. 21 n.3. And Dr. Brogioli testifies that “Foster does not specify the relationship of the dies forming the substack that shares a TSV or PTV” or suggest a certain die grouping. Ex. 2023 ¶ 164. Regardless of whether Foster teaches a certain die grouping, we agree with Petitioner that Foster shows that “[s]haring a TSV among two or more dies was a known option.” *See* Pet. 30.

Petitioner’s contention, as noted above, is that a person of ordinary skill in the art “would have been motivated to implement a shared data bus for multiple memory chips as taught by Rajan . . . using, e.g., Kim’s TSV interconnects.” Pet. 30. Rajan’s Figure 4 is shown below.




**FIG. 4**

Rajan’s Figure 4 above “shows a buffered stack of DRAM circuits having one address, control, and clock bus and two data busses.” Ex. 1015, 2:6–7; *see also* Ex. 1015, 4:51–55 (“Although the embodiments described here show the stack consisting of multiple DRAM circuits, a stack may refer to any collection of memory circuits (e.g. DRAM circuits, flash memory circuits, or combinations of memory circuit technologies, etc.).”). As shown in Figure 4, memory chips 417A and 417B share data bus 415A, and memory chips 417C and 417D share data bus 415B. Ex. 1015, 5:39–42.

Dr. Brogioli testifies that “Rajan uses wire bonds for connecting its DRAM circuits to the data busses, and wire bonds ordinarily cannot be shared in the way that TSVs are shared.” Ex. 2023 ¶ 143 (citing Ex. 1011 ¶ 50). In support of this statement, Dr. Brogioli does not cite the Rajan reference on which Petitioner’s challenge is based; rather, he cites Rajan137, which states that “data signals may be wired as one common bus, several busses or as an individual bus to each DRAM circuit.” Ex. 1011 ¶ 50, *quoted in* Ex. 2023 ¶ 143. Rajan, however, does not include this disclosure and does not state that its dies are connected with wire bonds. Thus, we do not credit Dr. Brogioli’s testimony that Rajan uses wire bonds.

Dr. Brogioli also testifies that

[t]he ‘060 patent discloses that an example of a die interconnect is a wire bond. But that wire bond is different from a wire bond dangling from the side of a DRAM circuit such as those in Rajan. Instead, the type of wire bonds referenced by the ‘060 patent are wires internal to a die, acting, for example, like a redistribution line.

Ex. 2023 ¶ 39 n.1 (citing “e.g., the red and blue wires in DDX4-48” (Ex. 2030)<sup>11</sup>). Even if Rajan’s Figure 4 shows wire bonds, an opinion that we do not credit as explained above, Dr. Brogioli does not explain why a wire bond in Rajan would be different from the wire bonds that can be die interconnects in the ’160 patent. *See* Ex. 1001, 5:54–56 (“Examples of die interconnects include, but are not limited to, through-silicon vias (TSV), conducting rods, wire bonds, and pins.”). In addition, it is not clear what “the red and blue wires in DDX4-48” (Ex. 2023 ¶ 39 n.1) show because the cited page is about “Accused HBM Products” and does not mention wire

---

<sup>11</sup> Dr. Brogioli appears to be referring to Ex. 2030. *See* Ex. 2023 ¶ 57 (citing “EX2030, DDX4-53”).

bonds. *See* Ex. 2030, DDX4-48. Thus, this testimony is not helpful in explaining any alleged differences in wire bonds. *See* 37 C.F.R. § 42.65(a) (“Expert testimony that does not disclose the underlying facts or data on which the opinion is based is entitled to little or no weight.”). Furthermore, we do not agree with Dr. Brogioli’s testimony that wire bond die interconnects in the ’160 patent “are wires internal to a die” (Ex. 2023 ¶ 39 n.1) because wires that are only internal to one die would not interconnect with another die, and, thus, would not be a “die interconnect.”

In view of the foregoing, we find that Rajan does not disclose the particular structure in which the shared data busses in Figure 4 are implemented. In other words, Rajan’s disclosure is agnostic as to the particular interconnect technology by which the memories that share data busses are connected to the buffer chip.

Dr. Brogioli also testifies that a “data bus can be shared without sharing interconnects, especially in Rajan,” “because Rajan uses wire bonds for connecting its DRAM circuits to the data busses.” Ex. 2023 ¶ 143 (citing Ex. 1011 ¶ 50). As discussed above, we do not credit Dr. Brogioli’s opinion that Rajan uses wire bonds because it is based on a different reference. Furthermore, the fact that a data bus can be shared without sharing interconnects is irrelevant because Petitioner’s proposed combination is to share interconnects (Kim’s TSVs). *See* Pet. 30. Rajan’s Figure 4 discloses a stack of memory circuits, as discussed above (*see* Ex. 1015, 2:6–7, 4:51–55), and Kim discloses using TSVs in stacks of memory chips, as also discussed above (*see* Ex. 1014 ¶¶ 48, 50, Fig. 5). Thus, Petitioner’s combination is a straightforward one – “implement a shared data bus for multiple memory

chips as taught by Rajan . . . using, e.g., Kim’s TSV interconnects.” Pet. 30 (discussed above).

At oral argument, Patent Owner’s counsel argued that the “TSV is the individual data wire going down to the bus that’s shared.” Tr. 65:12–13.

According to Patent Owner’s counsel,

Rajan talks about [two chips] sharing a data bus. So, I don’t see how, from Rajan, you would get the idea of a shared TSV when Rajan speaks solely about sharing data buses. And the data bus, under JEDEC standard, involves individual data wires coming from each of the chips.

Tr. 65:26–66:4. By this argument, Patent Owner’s counsel appears to be suggesting that, even if TSVs are used, they would just connect each chip to a shared data bus. But this is not the proposed combination, as stated above, which is to implement a shared bus per Rajan using Kim’s TSVs. *See* Pet. 30. Implementing a bus using TSVs is consistent with the evidence of record. *See* Ex. 1025, 1:39–42 (“A more recent approach for wafer stacking is to connect the signals together with vias, effectively sending a bus of signal lines vertically through a stack of dies.”), 2:56–64 (“[A]ggregations of PTVs and TSVs will often be used to effect connection of bus signals from a substrate up through the die stack to circuitry somewhere in the die stack.”).

Based on the foregoing, we find persuasive Petitioner’s contention that a person of ordinary skill in the art

would have been motivated to implement a shared data bus for multiple memory chips as taught by Rajan (e.g., 415A and 415B, above) using, e.g., Kim’s TSV interconnects . . . , in part because a [person of ordinary skill in the art] would have understood that there were a finite number of known ways to connect additional dies, including having a subset of dies sharing a through-silicon via (TSV).

*See* Pet. 30. Furthermore, we find persuasive Petitioner’s reasoning that a person of ordinary skill in the art “would have been motivated to connect additional memory dies to Kim’s TSV1 and TSV2 because it would not require creating new TSVs (which would add space and circuitry).” Pet. 30 (citing Ex. 1003 ¶ 160). As discussed above, Kim suggests having more than two dies, Rajan discloses how to share a data bus among multiples dies, and the evidence of record shows that TSVs were used to implement data busses. Ex. 1014 ¶¶ 48, 50; Ex. 1025, 1:39–42, 2:56–64. We find that using the existing TSVs of Kim would be beneficial to avoid having to create new ones, “which would require space and additional circuitry.” Ex. 1003 ¶ 160.

Having found persuasive the above contentions, we turn to the issue of whether the proposed combination would have been within the ability of a person of ordinary skill in the art. *See KSR*, 550 U.S. at 421 (noting that “a person of ordinary skill has good reason to pursue the known options *within his or her technical grasp*” (emphasis added)); *see also id.* at 417 (“[I]f a technique has been used to improve one device, and a person of ordinary skill in the art would recognize that it would improve similar devices in the same way, using the technique is obvious unless its actual application is beyond his or her skill.”).

Patent Owner argues that the asserted combination could result in four possible configurations and that each configuration either would be inoperable or would not meet the claims. PO Resp. 22 (citing Ex. 2023 ¶ 86). Patent Owner asserts that the

four possibilities are: (1) four chips belong to two ranks and two chips from the same rank share a common TSV; (2) four chips belong to two ranks and a chip from one rank shares a TSV with a chip from another rank; (3) four chips belong to four ranks and

each of the four ranks has its own TSV; and (4) four chips belong to four ranks and every two ranks share a TSV.

PO Resp. 22 (citing Ex. 2023 ¶ 86).

We need not address the first three configurations because, for the reasons explained below, we find Petitioner’s combination persuasive in the fourth configuration. Patent Owner argues that a person of ordinary skill in the art would not combine the teachings of Kim and Rajan in the fourth configuration because it would result in unavoidable data collisions.

PO Resp. 32–44. Patent Owner asserts that, “[a]s Dr. Brogioli testified, handling data collisions involves ‘a lot of complexities’ that [are] ‘beyond the scope of a person of skill.’” PO Sur-reply 9 (citing Ex. 1052, 283:17–284:12). Patent Owner argues, therefore, that “the requirement that at least one die interconnect be in connection with a group of array dies (plural) would be incompatible with Kim’s system and the proposed modifications are inoperable for Kim’s memory system for its intended purposes.” PO Resp. 43 (citing Ex. 2023 ¶¶ 149–150); *see* Ex. 2023 ¶ 150 (“A person of ordinary skill in the art would not have had a reason or a reasonable expectation of success to modify Kim so that multiple dies share the same TSV signaling path.”).

Petitioner argues that techniques for avoiding collisions on shared data lines were well-known to persons of ordinary skill in the art, citing as one example Rajan’s teachings of timing solutions to avoid collisions. Pet. Reply 2–8 (citing, *inter alia*, Ex. 1015, 8:59–12:13, Figs. 9–15). For the reasons explained below, we find persuasive Petitioner’s contentions that persons of ordinary skill in the art knew how to deal with collisions.

The specification of the ’160 patent mentions collisions only once:

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.

Ex. 1001, 21:24–41 (emphasis added). The fact that the ’160 patent mentions *reducing or preventing* signal collisions suggests that there is no prohibition on signal collisions. This is in contrast with Dr. Brogioli’s more limited view that if there are any data collisions the resulting structure would be inoperable. Ex. 2023 ¶ 150. Moreover, the challenged claims do not recite that data collisions must be avoided, and “[t]he reasonable expectation of success requirement refers to the likelihood of success in combining references *to meet the limitations of the claimed invention.*” *Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd.*, 821 F.3d 1359, 1367 (Fed. Cir. 2016) (emphasis added). Thus, Dr. Brogioli’s testimony about the lack of “a reasonable expectation of success” (Ex. 2023 ¶ 150) is not commensurate in scope with the claims. Furthermore, the above-quoted passage from the ’160 patent does not provide any particular details on how to control timing to reduce or prevent collisions, suggesting that this was within the

knowledge and skill of a person of ordinary skill in the art, as Petitioner asserts. *See* Pet. Reply 2–8. Thus, “the [’160] patent itself does not disclose the level of detail that [Patent Owner] would have us require of the prior art.” *See Lockwood v. Am. Airlines, Inc.*, 107 F.3d 1565, 1570 (Fed. Cir. 1997).

At oral argument, Patent Owner’s counsel argued that the ’160 patent “avoids the collisions” by “dramatically remov[ing] the latency on the TSVs,” which “dramatically increases the speed at which the load and reads can be cleared.” Tr. 40:4–11. According to Patent Owner’s counsel, the ’160 patent’s collision solution is at column 8, lines 1–21 and is “a physical solution to a timing problem, not a timing solution to a physical problem.” Tr. 39:15–40:11. The cited disclosure is below.

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 232a is less than the maximum load and the load of the data conduit 232b is less than the maximum load. Further, in many implementations, the combined load of the data conduit 232a and the data conduit 232b is less than the maximum load of a single data conduit. In other words, it is possible to design the data conduit 232a and the data conduit 232b 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.


Ex. 1001, 7:66–8:21. This passage says nothing about avoiding collisions. As discussed above, the '160 patent discloses that “[c]ontrolling the command/address time slots and the data bus time slots enables the control dies 822 to reduce or prevent signal collisions.” Ex. 1001, 21:36–39. Thus, the '160 patent suggests that there is a timing solution to address collisions, contrary to Patent Owner’s counsel’s argument. *See* Tr. 39:15–40:11.

Petitioner asserts that Rajan “provides much more detail than the 160 Patent with respect to timing solutions for avoiding collisions.” Pet. Reply 6 (citing Ex. 1015, 8:59–12:13, Figs. 9–15). We agree because, as discussed above, the '160 patent simply mentions controlling timing to reduce or prevent collisions but does not provide any particular details on how to do this. Ex. 1001, 21:36–39. Rajan, however, actually explains how to avoid collisions through the timing of various operations. *See, e.g.*, Ex. 1015, 11:12–12:13, Figs. 14, 15. For example, Rajan discloses that “FIGS. 14 and 15 are a timing diagram 1400 and a timing diagram 1500 illustrating methods of avoiding such collisions.” Ex. 1015, 11:12–14; *see also* Ex. 1052, 200:15–19 (Dr. Brogioli testified as follows: “It looks like Figure 14 says it’s a timing diagram illustrating methods of avoiding such collisions. These look like command collisions of things like activate commands, read and write commands.”).

Patent Owner argues that Rajan does not teach how to avoid collisions for operations that “involved cross-rank read-after-write operations and certainly not for Kim’s memory systems” and that “[d]elaying write commands will just further delay the write data, and does not solve Kim’s data collision problem.” PO Sur-reply 10 (citing Ex. 1014 ¶ 42, Fig. 4A).

Rajan's disclosure is more robust than simply delaying write operations. For example, Rajan discloses that the buffer chip can delay activate operations by one to three clock cycles and that "[t]he actual delay selected may depend on the presence or absence of other DRAM operations that may conflict with the activate operation, and may optionally change from one activate operation to another. In other words, the delay may be dynamic." Ex. 1015, 11:40–47.

In support of its arguments regarding the fourth configuration, in which four chips belong to four ranks and every two ranks share a TSV, Patent Owner refers to Dr. Brogioli's modification of Kim's Figure 2 below.



PO Resp. 33; Ex. 2023 ¶ 120. Kim's Figure 2 illustrates a memory configuration including shared data input/output (I/O) section 1000 connected via first and second global I/O lines GIO\_Rank0 and GIO\_Rank1 to first and second I/O driving sections 100 and 200, which are connected, respectively, to Rank0 and Rank1. Ex. 1014 ¶¶ 25–26. The modified figure above includes a third I/O driving section connected between GIO\_Rank0 and Rank 2 and a fourth I/O driving section connected between GIO\_Rank1 and Rank 3. Ex. 2023 ¶ 120. Dr. Brogioli testifies that this illustrates

Dr. Wolfe’s proposal that “Kim could be modified by having four ranks sharing two data busses, for instance, in a rank multiplication scenario.”

Ex. 2023 ¶¶ 119–120. At deposition, Dr. Brogioli confirmed that he “assume[d] that Kim’s circuits would operate the same way as in Kim’s original Figure 2.” Ex. 1052, 320:5–11.

Dr. Brogioli’s assumption of no change in operation with the addition of two ranks fails to account for the combined teachings of Kim and Rajan and the knowledge and skill of a person of ordinary skill in the art. As the Federal Circuit instructs,

*KSR* does not require that a combination only unite old elements without changing their respective functions. Instead, *KSR* teaches that “[a] person of ordinary skill is also a person of ordinary creativity, not an automaton.” And it explains that the ordinary artisan recognizes “that familiar items may have obvious uses beyond their primary purposes, and in many cases a person of ordinary skill will be able to fit the teachings of multiple patents together like pieces of a puzzle.”

*ClassCo, Inc. v. Apple, Inc.*, 838 F.3d 1214, 1219 (Fed. Cir. 2016) (citations omitted; alteration in original). The Federal Circuit further cautioned in *ClassCo* that the “rationale of *KSR* does not support [the] theory that a person of ordinary skill can only perform combinations of a puzzle element A with a perfectly fitting puzzle element B.” *Id.* Here, Dr. Brogioli alleges that the modified figure reflects “a rank multiplication scenario” (Ex. 2023 ¶ 119), but when asked about the fact that there are four ranks but only two chip select signals shown in the modified figure, Dr. Brogioli testified that “Dr. Wolfe didn’t articulate anything about additional chip select signals that would relate to what this figure is visualizing.” Ex. 1052, 320:12–24.

Dr. Wolfe, however, testifies that Rajan teaches “emulation techniques, including rank multiplication techniques” (Ex. 1003 ¶ 161), and that Rajan

teaches rank multiplication by generating chip select signals for each chip in the stack. Ex. 1003 ¶¶ 140, 447–448 (quoting Ex. 1015, 6:34–38). In particular, Rajan discloses that “extra address bits may be decoded by the buffer chip to individually select the DRAM chips, utilizing separate chip select signals (not shown) to each of the DRAM chips in the stack.”

Ex. 1015, 6:34–38. Thus, Dr. Brogioli’s modified version of Kim’s Figure 2 does not reflect “a rank multiplication scenario” (Ex. 2023 ¶ 119) in view of the teachings of Rajan.

Patent Owner also argues that “Petitioner has not shown that its relied-on non-DRAM chips would utilize chip-select signals.” PO Resp. 51–52 (citing Ex. 2023 ¶¶ 169–171, 193; Ex. 1015, 6:34–38; Exs. 2043, 2044). Dr. Brogioli identifies NAND and SRAM datasheets that he says do not show “pins for chip-select signals.” Ex. 2023 ¶¶ 169–170 (citing Ex. 2043, 8; Ex. 2044, 5–6). But at deposition, Dr. Brogioli acknowledged that these devices use chip enable signals. Ex. 1052, 236:23–239:3 (discussing Ex. 2043), 239:10–241:6 (discussing Ex. 2044). We agree with Petitioner that the chip enable signals are used “to select the chip that performs the read/write operation.” See Pet. Reply 19; Ex. 2043, 16 (“CE# is used to enable the device.”); Ex. 2044, 5 (pin for “Chip enable input”). The use of chip enable signals is consistent with the district court construction, which Patent Owner urges “should be applied by the Board,” that ““excludes situations in which a chip select signal could *enable* multiple array dies at once.”” See PO Resp. 18–19 (quoting Ex. 2004, 33–34 (emphasis added)).

Patent Owner asserts that “[t]here is also no evidence that the concept of ‘rank’/‘rank multiplication,’ would apply to non-DRAM memory

systems.” PO Resp. 52 n.8 (quoting Ex. 1022, 413 (“[T]he word *rank* is now used to denote a set of DRAM devices that operate in lockstep to respond to a given command in a memory system.”)). We disagree for the reasons discussed above in section II.D.2.d.1, which explains that Kim does not even mention DRAM and thus is not limited to DRAM and that Rajan states that “any type of memory whatsoever” may be used in its disclosure and lists non-DRAM memories. Ex. 1015, 15:3–9.<sup>12</sup> Thus, we find that the combined teachings of Kim and Rajan teach ranks and rank multiplication with non-DRAM array dies.

*(3) Summary of findings for 1.c, 1.d.1, 1.d.2*

For the reasons discussed above, we find that the combination of Kim and Rajan teaches the subject matter of limitations 1.c, 1.d.1, and 1.d.2 and that a person of ordinary skill in the art would have combined the teachings of Kim and Rajan in the manner asserted with a reasonable expectation of success.


*e) Control die (1.e.1), data conduits (1.e.2, 1.e.3), and drivers (1.e.4, 1.e.5)*

Claim 1 recites a “control die” (1.e.1) comprising “first data conduits between the first die interconnects and the data terminals” (1.e.2) and “second data conduits between the second die interconnects and the data terminals” (1.e.3).

Petitioner argues that Kim’s main chip C0 is a control die and that Kim’s Figure 5 (annotated version below) shows first and second data conduits between the TSVs and the data terminals DQ. Pet. 44–49.

---

<sup>12</sup> We further note that Patent Owner’s argument is conclusory and appears only in a footnote. *Cf. CommScope Techs. LLC v. Dali Wireless Inc.*, 10 F.4th 1289, 1296 (Fed. Cir. 2021) (“[A]n argument that is only made in a footnote of an appellant’s brief is forfeited.”).




Pet. 44. In the annotated figure above, Petitioner identifies the line from TSV1 to shared data I/O section 1000 as "first data conduits" in pink and identifies the line from TSV2 to shared data I/O section 1000 as "second data conduits" in rose. Pet. 45–46. As shown and as Petitioner asserts, the conduits are between the interconnects (TSVs) and the data terminals (DQ). See Pet. 46–47.

Claim 1 further recites "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" (1.e.4) and "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" (1.e.5).

Petitioner argues that it would have been obvious for the data conduits to have drivers of different strengths for the TSVs of different lengths in Kim based on Wyman's teachings of using different amounts of drive for

different path lengths. Pet. 51–53 (citing Ex. 1017, 1:22–24, 2:61–65, 6:15–30, 6:44–50, 7:14–18, Figs. 5, 7, 8, 10, 11; Ex. 1030, 135–38; Ex. 1003 ¶¶ 257–265); *see also* Pet. 32–34 (discussing combination with Wyman).

Below is Petitioner’s annotated version of a portion of Kim’s Figure 3.




Pet. 34, 50. The portion of Kim’s Figure 3 above shows rank selecting unit 1100 as part of shared data I/O section 1000 including write selecting unit 1110 and read selecting unit 1120. Ex. 1014 ¶¶ 34–35. Petitioner modifies the figure above to include a navy blue triangle representing “smaller drivers” between write selecting unit 1110 GIO\_Rank0 (first die interconnects) and a larger dark green triangle representing “larger drivers” between write selecting unit 1110 and GIO\_Rank1 (second die interconnects). Pet. 33–34, 49–50.

Patent Owner presents two main arguments in response to Petitioner’s reliance on Wyman: (1) that Wyman teaches using a single driver such that

a person of ordinary skill in the art would not have used different drivers as recited in the claim; and (2) that Wyman and Kim do not suggest that TSVs of different lengths in Kim would warrant using different amounts of drive. PO Resp. 44–47.

We disagree with Patent Owner’s second argument because it ignores the express disclosure of Wyman. Figure 8 of Wyman is below.



Wyman’s Figure 8 above shows “different through-chip via connections” for stacked chips 700-4 (the “mother chip”) and 700-1, 700-2, and 700-3 (the “daughter chips”). Ex. 1017, 2:21–23, 6:9–14. Figure 8A shows through-chip via connection 802 from mother chip 700-4 to the closest daughter chip labeled 700-3, and Figure 8B shows through-chip via connection 804 from mother chip 700-4 to the farthest daughter chip labeled 700-1. Ex. 1017, 6:15–23. Wyman discloses that “[t]he increased resistance, capacitance and impedance of such a connection (804) might

require additional drive than that referred to in connection with FIG. 8A” and that the drive used for connection 802 “would be inadequate.”

Ex. 1017, 6:24–28. Based on this disclosure of Wyman, we agree with Petitioner that “Wyman discloses that shorter paths (e.g., 802 . . . ) have less load and thus require less drive, while longer paths (e.g., 804 . . . ) have more load and thus require a larger drive.” *See* Pet. 32.

We also have considered the following testimony from Dr. Brogioli:

Wyman and Kim do not suggest that the two TSVs with different lengths would result in sufficiently different load to even warrant using different taps in a Wyman-style driver. For instance, load contribution from TSV “may be negligible compared to load contribution from the array dies.” *See* EX1001, 4:36-38. In such situations, same driver strengths can be used and a driver circuitry like that shown in Wyman, Figure 2 will suffice.


Ex. 2023 ¶ 216. The quoted passage from the ’160 patent states in full: “*In some embodiments*, the load contribution from a die interconnect may be negligible compared to the load contribution from the array dies.” Ex. 1001, 4:39–41 (emphasis added).<sup>13</sup> Dr. Brogioli may be correct that, “[i]n such situations, same driver strengths can be used.” Ex. 2023 ¶ 216. Wyman, however, expressly discloses a different situation in which the signal drive for a shorter connection “would be inadequate” for a “longer through-chip via connection” due to “increased resistance, capacitance and impedance of” the longer connection. Ex. 1017, 6:21–30.

Based on the foregoing, we find that Wyman teaches using different amounts of signal drive for different lengths of TSVs, and we turn to Patent Owner’s argument that Wyman teaches using a single driver. More

---

<sup>13</sup> The same passage appears at column 4, lines 36–38 of the ’060 patent, which is what Dr. Brogioli cites.

particularly, Patent Owner argues that Wyman teaches a single driver circuit with different taps at which different drive levels can be selected. PO Resp. 44–45 (citing Ex. 1017, 1:36–37, 2:66–3:1, 6:15–50; Ex. 2023 ¶ 211). Patent Owner’s contentions refer to a configuration such as the one below in Figure 5 of Wyman.



Wyman’s Figure 5 above shows current drive portion 500 including driver circuits 502-1, 502-2, 502-3, and 502-4, each of which, via respective enabling leads 504-1, 504-2, 504-3, and 504-4, can be enabled, disabled, or placed in a tri state where minimal power is drawn. Ex. 1017, 3:37–41, 4:66–5:11. Wyman explains that “a designer may now utilize and tap-off at five locations (506-1 [T0], 506-2 [T1], 506-3 [T2], 506-4 [T3], 508) depending on the drive requirements for a specific element or device.” Ex. 1017, 5:11–14.

According to Patent Owner, “Wyman’s driver circuitry, once fabricated, would be regarded as a single driver of a given size regardless of the number of driver components used.” PO Resp. 45–46 (citing Ex. 2024, 141:18–145:14, 145:24–146:7, 149:24–150:8; Ex. 2023 ¶ 215). Patent

Owner argues, therefore, that a person of ordinary skill in the art would not have used different drivers but, instead, would have used a single driver with selectable drive levels. PO Resp. 45 (citing Ex. 1017, 6:44–50; Ex. 2023 ¶¶ 213–214).

Patent Owner’s proposed combination of using a single driver may be an acceptable option and may be even a better option than having separate drivers. The proposed combination, however, need not be the best option. *See Intel Corp. v. Qualcomm Inc.*, 21 F.4th 784, 800 (Fed. Cir. 2021) (“Our caselaw is clear. It’s not necessary to show that a combination is ‘the *best* option, only that it be a *suitable* option.” (quoting *PAR Pharm., Inc. v. TWI Pharms., Inc.*, 773 F.3d 1186, 1197–98 (Fed. Cir. 2014))). Petitioner proposes using separate drivers in its annotation of Kim’s Figure 3, which is reproduced above in this section. *See* Pet. 33–34 (asserting that a person of ordinary skill in the art would have been motivated to use “smaller drivers (navy blue) for the shorter TSV1 (light green) and larger drivers (dark green) for the longer TSV2 (teal)”), 49–50. Dr. Wolfe testifies that “you can provide a separate copy of [Wyman’s] circuit 500 for each thing that needs to be driven after you analyze the current requirements.” Ex. 2025, 129:2–7, *cited in* Pet. Reply 15. As noted above, Wyman teaches using different amounts of signal drive for different lengths of TSVs. We see no patentable distinction between whether the different amounts of drive come from one set of transistors that can be tapped for different drive strengths or from two different sets of transistors. *See In re Harza*, 274 F.2d 669, 671 (CCPA 1960) (“It is well settled that the mere duplication of parts has no patentable significance unless a new and unexpected result is produced . . .”).

According to Patent Owner, Petitioner’s proposal to use one driver with one number of transistors and another driver with a different number of transistors “relies on nothing more than appeal to common sense, which is improper under *Arendi*.” PO Sur-reply 19 (citing *Arendi SARL v. Apple Inc.*, 832 F.3d 1355 (Fed. Cir. 2016)). We disagree that Petitioner’s position relies on nothing more than common sense because Wyman teaches using different amounts of signal drive for different TSVs, as discussed above. That physically separate drivers could be used to produce different amounts of drive in different paths is an unremarkable observation. *See KSR*, 550 U.S. at 421 (“Rigid preventative rules that deny factfinders recourse to common sense . . . are neither necessary under our case law nor consistent with it.”).

Patent Owner and Dr. Brogioli also rely on testimony from Dr. Harold Stone, a technical expert for the Micron entities in district court litigation. *See* Ex. 2024, 141:18–145:14, 145:24–146:7, 149:24–150:8, *cited in* PO Resp. 45–46 and Ex. 2023 ¶ 215. As Patent Owner notes, this testimony concerns “an example of ten-transistor driver where a subset or all ten of the transistors can be used to drive a load.” PO Resp. 45–46; *see also* PO Sur-reply 19 (parenthetical noting “when discussing the size of a 10-transistor driver that uses two transistors for some operations and ten transistors for others, Dr. Stone considers the physical dimension of the ten transistors as the driver size without ever referring to those as two drivers”). In this line of questioning, Dr. Stone confirmed that his testimony was based on a hypothetical in which you could “turn on all ten transistors,” and he explained that “[i]f you can turn on all ten, the driver size is related to the load that you can drive. And if you’re going to drive ten transistors, you

have to have enough area to dissipate the power they generate.” Ex. 2024, 147:3–11. This testimony, therefore, does not concern the situation in which different and separate drivers drive separate loads, as in Petitioner’s proposed combination discussed above.

Furthermore, although Wyman discloses that different taps can be provided to provide different drive levels, Wyman also discloses an alternative configuration in which taps are not used, as Petitioner correctly points out. *See* Pet. Reply 15 (citing Ex. 1017, 5:34–44). More particularly, Wyman discloses that its “approach could be used with a conventional drive circuit made up of multiple stages by using a via approach.” Ex. 1017, 5:35–44. In this configuration, Wyman discloses forming an “electrically conductive connection to an intermediate point between stages where the current drive is adequate, as needed, and to disable any remaining downstream stage(s) by using a via to break one or more connections and leaving the via unfilled or filling the via with an insulator.” Ex. 1017, 5:35–44. Thus, Wyman expressly discloses that a suitable option involves disabling unneeded transistors and not using taps to provide selectable drive levels. *See Intel*, 21 F.4th at 800. Furthermore, we find that Wyman’s disclosure of providing different drive strengths with different numbers of transistors would result in drivers of different physical size, consistent with the disclosure of the ’160 patent that “[t]he size of the driver may be adjusted by the selection of the transistor size and/or number of transistors included in the driver.” Ex. 1001, 17:32–34; *see* Pet. Reply 15 (discussing this disclosure); *see also* Ex. 2024, 147:7–9 (Dr. Stone’s testimony that “the driver size is related to the load that you can drive”).

Based on the foregoing, we find that Wyman teaches using different amounts of signal drive for different lengths of TSVs and that a person of ordinary skill would have recognized that different, physically separate drivers can be used to provide those different amounts of signal drive.

*f) Objective evidence of non-obviousness*

We must consider any evidence of objective indicia of non-obviousness before reaching our conclusion on obviousness. *WBIP, LLC v. Kohler Co.*, 829 F.3d 1317, 1328 (Fed. Cir. 2016). This is one of the factual considerations underlying an obviousness determination. *Graham*, 383 U.S. at 17.

Patent Owner argues that the “history of commercial 3DS [(3D-Stacked)] and HBM [(high bandwidth memory)] in particular is objective evidence that the patented design is not obvious.” PO Resp. 1. Patent Owner, however, does not make any arguments to establish the nexus required for consideration of secondary considerations. *See ClassCo*, 838 F.3d at 1220 (“For objective evidence of secondary considerations to be accorded substantial weight, its proponent must establish a nexus between the evidence and the merits of the claimed invention.” (quotations and citation omitted)); *see also* Tr. 42:9–25 (Patent Owner confirming that it did not present an argument for secondary considerations).

Thus, the record does not contain any supportable arguments for objective indicia of non-obviousness.

*g) Conclusion for Claim 1*

We have considered the full trial record, and, for the reasons discussed above and based on Petitioner’s contentions and evidence, we conclude that the subject matter of claim 1 would have been obvious to a

person of ordinary skill in the art based on the combined teachings of Kim, Rajan, and Wyman.

### *3. Dependent Claims 2–5*

Claim 2 depends from claim 1 and recites “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.” Petitioner relies on Kim’s teaching that TSV2 is longer than TSV1 and Wyman’s teachings of using greater drive strength for longer TSVs. Pet. 53–54 (citing Ex. 1014, Fig. 5; Ex. 1003 ¶¶ 267–272). Referring to its arguments for independent claim 1, Patent Owner asserts that “Petitioner has not shown that a [person of ordinary skill in the art] would have used two different drivers or drivers of different sizes.” PO Resp. 81. We disagree with this argument for the reasons discussed above in § II.D.2.e.

Claim 3 depends from claim 1 and recites “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.” Petitioner argues that “it would be an obvious design choice to have an ‘asymmetrical’ stack” with fewer array dies in the second group than in the first group based on Kim’s teaching that any number of chips may be used and Rajan’s teaching that memory circuits in a stack can be symmetrical or asymmetrical. Pet. 55–56 (citing Ex. 1014 ¶¶ 27, 48, 50; Ex. 1015, 2:62–65, 7:63–67; Ex. 1003 ¶¶ 282–285). Petitioner contends that a person of ordinary skill in the art would have been motivated to arrange such an asymmetrical “stack so that a longer TSV (e.g., TSV2 in Kim) has fewer dies connected to it to reduce the power requirements of driving the longer TSV, as taught by

Wyman.” Pet. 56–57 (citing Ex. 1017, 2:61–65, 6:15–30, 7:14–26, Figs. 7, 8; Ex. 1003 ¶¶ 286–288). Petitioner also argues that U.S. Patent 7,796,446 B2 (Ex. 1024 (“Ruckerbauer”)) discloses an asymmetric arrangement of dies, thus “providing a further motivation for there to be fewer ‘array dies in the second group’ than ‘in the first group of array dies.’” Pet. 56 (citing Ex. 1024, 8:58–66, 9:14–18, Fig. 5; Ex. 1003 ¶¶ 284–285).

Patent Owner disputes Petitioner’s reliance on “Ruckerbauer as evidence that an asymmetric number of chips may share respective TSVs,” arguing that Ruckerbauer discloses that “each TSV is connected to data port on every die” and, thus, does not show the claimed arrangement. PO Resp. 81 (citing Ex. 1024, Figs. 5A–5B; Ex. 2025, 105:11–106:7; Ex. 2023 ¶¶ 157–159, 295). We need not resolve whether Ruckerbauer’s arrangement shows electrical communication with all dies because Petitioner relies on Ruckerbauer as “providing a further motivation” to create an asymmetrical stack, and we find persuasive Petitioner’s other evidence for this assertion. *See* Pet. 55–56. For example, Kim discloses that “any number of main and slave chips may be used.” Ex. 1014 ¶ 48; *see also* Ex. 1014 ¶ 50 (“While the semiconductor memory apparatuses having two ranks are explained with reference to FIGS. 2 and 5, a person having ordinary skill in the art will appreciate that the technical concept of the present invention can be applied to a semiconductor memory apparatus which are divided into three or more ranks.”). Rajan also discloses that stacked memory circuits may be “asymmetrical.” Ex. 1015, 2:62–65. Dr. Brogioli’s testimony that “Rajan shows the same number of dies per group” (Ex. 2023 ¶ 295) does not account for Rajan’s “asymmetrical” disclosure.

Claim 4 depends from claim 1 and recites “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.” Referring to its contentions for the driver limitations of claim 1 (1.e.4 and 1.e.5), Petitioner argues that it would have been obvious to a person of ordinary skill in the art “to follow Wyman’s teaching of using different driver sizes corresponding to the different loads created by the TSVs of different lengths in Kim’s stack.” Pet. 57 (citing Ex. 1017, 7:14–26; Ex. 1003 ¶¶ 291–297).

Claim 5 depends from claim 1 and recites “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.” Petitioner argues that Kim teaches a control circuit in rank selecting unit 1100 that control the states of the TSVs to be active or not active based on signals received from an external device. Pet. 57–58 (citing Ex. 1014 ¶¶ 32, 35–36, 38, Fig. 3; Ex. 1003 ¶¶ 298–306; Ex. 1019, 6–14, 18, 33; Ex. 1023, 9, Fig. 16; Ex. 1022, 318–20, 332–35).

Other than its arguments for independent claim 1, which we address above in § II.D.2, and the arguments we address above in this section, Patent Owner does not raise additional arguments for these claims. We have reviewed Petitioner’s arguments and evidence, and we find them persuasive. Therefore, having considered the full record developed during the trial, we conclude that claims 2–5 are unpatentable as obvious over the combined teachings of Kim, Rajan, and Wyman.

#### *4. Claims 6–9*

Independent claim 6 is directed to a “memory package” and recites limitations substantially similar to subject matter recited in claims 1 and 3.

Claims 7–9 recite subject matter that is substantially similar to subject matter recited in claims 1, 2, and 5. For these claims, Petitioner refers to its contentions for similarly-recited subject matter in claims 1–3 and 5. Pet. 59. Patent Owner does not raise arguments in addition to those already addressed above.

We have reviewed Petitioner’s arguments and evidence, and we find them persuasive for the reasons discussed above for claims 1–3 and 5. Therefore, having considered the full record developed during the trial, we conclude that claims 6–9 are unpatentable as obvious over the combined teachings of Kim, Rajan, and Wyman.

*5. Claims 10–14 and 17–20*

Below we address Petitioner’s contentions for claims 10–14 and 17–20, which we find persuasive. Patent Owner does not raise arguments for these claims in addition to those already addressed above.

Independent claim 10 is directed to 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.” Claim 10 recites that each DRAM package comprises “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.” Claim 10 also recites that each DRAM package comprises subject matter that is substantially similar to limitations 1.c–1.e.

Petitioner argues that the combination of Kim, Rajan, and Wyman teaches the subject matter of claim 10 that is not recited in claim 1, and

Petitioner refers to its claim 1 contentions for the remaining subject matter. Pet. 60–65. In particular, Petitioner argues that the combination teaches a JEDEC memory module that has a register device (Rajan’s register 804). Pet. 60–62 (citing Ex. 1015, 1:28–32, 3:5–7, 4:20–24, 8:52–58, 15:3–12, Fig. 8; Ex. 1014 ¶ 50; Ex. 1019, 13, 33; Ex. 1022, 316–20, 418–19, Fig. 7.2; Ex. 1023, 9, Fig. 16; Ex. 1037, 1:34–37, Fig. 1; Ex. 1003 ¶¶ 353–368). As discussed above with respect to claim 1, we find that Rajan teaches using non-DRAM memories in a JEDEC-compliant memory package. We also find that the combination of Kim and Rajan teaches “a plurality of DRAM packages” and “a register device configured to receive input command/address signals from the system memory controller and to output control signals” because Rajan discloses “high capacity DIMM 800 using a plurality of buffered stacks of DRAM circuits 802 and a register device 804,” which “performs the addressing and control of the buffered stacks.” Ex. 1015, 8:52–56; *see* Pet. 62–63 (discussing this disclosure and Fig. 8 of Rajan). We also agree with Petitioner’s contention, and we find, that Rajan’s Figure 8 disclosure teaches data and control terminals, as recited in claim 10. *See* Pet. 63–64 (citing Ex. 1015, Fig. 8; Ex. 1019, 6–13; Ex. 1022, 418–419; Ex. 1003 ¶¶ 374–379). We find Petitioner’s contentions persuasive for the remaining subject matter of claim 10 for the reasons discussed for similar limitations in claim 1.

Claim 11 recites, “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.” Petitioner refers to its contentions for claim 5, which recites similar subject matter. Pet. 65.

Claim 12 recites, “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.” Petitioner argues that chip select and read/write command signals would be used per JEDEC to control the direction of data transfer. Pet. 65–67 (citing Ex. 1014 ¶¶ 31–38, Fig. 3; Ex. 1019, 6–14, 18, 33; Ex. 1023, 9, Fig. 16; Ex. 1022, 318–20, 332–35; Ex. 1003 ¶¶ 401–407).

Claim 13 recites, “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.” As with claim 12, Petitioner cites chip select and read/write signals as controlling the direction of data transfer. Pet. 68 (citing Ex. 1014 ¶¶ 37–38; Ex. 1003 ¶¶ 401–416).

Claim 14 recites, “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.” Petitioner argues that Rajan’s “Address & Control” signals (Fig. 8) are received at the memory module and that address signals are sent to the dies to identify where to store or retrieve data. Pet. 68–72 (citing Ex. 1014 ¶¶ 29–30, 32, 38; Ex. 1015, 14:11–18, 14:55–60, Figs. 4, 18; Ex. 1019, 6–13, 18, 33; Ex. 1022, 318–20, 332–35; Ex. 1023, 9, Fig. 6; Ex. 1003 ¶¶ 418–434).

Claim 17 recites, “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.” Petitioner refers to its contentions for claim 4, which recites similar subject matter. Pet. 78.

Claim 18 recites, “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.” Petitioner argues that Rajan teaches delaying signals by two clock cycles and that a person of ordinary skill in the art would have implemented such a delay in Kim to emulate JEDEC. Pet. 78–79 (citing Ex. 1015, 9:46–10:27, Fig. 11; Ex. 1019, 23–24; Ex. 1003 ¶¶ 483–490).

Claim 19 recites, “The memory module of claim 10, wherein the control die includes data buffers to control the timing of the data signals.” Petitioner cites Rajan’s disclosure of controlling the CAS latency timing. Pet. 79–80 (citing Ex. 1015, 8:59–9:45, Fig. 9; Ex. 1019, 23–24; Ex. 1003 ¶¶ 491–496).

Claim 20 recites, “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.” Petitioner refers to its contentions for claim 3, which recites similar subject matter. Pet. 80.

As noted above, Patent Owner does not raise additional arguments for these claims. We have reviewed Petitioner’s arguments and evidence, and we find them persuasive. Therefore, having considered the full record developed during the trial, we conclude that claims 10–14 and 17–20 are unpatentable as obvious over the combined teachings of Kim, Rajan, and Wyman.

6. *Claims 15 and 16*

Claim 15 recites,

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.

Petitioner argues that chip select signals are received as part of JEDEC, and Petitioner relies on Rajan's disclosure of using separate chip select signals for each chip to show rank multiplication. Pet. 72–73 (citing Ex. 1015, 3:27–30, 6:30–7:67, 8:6–13, 8:56–58; Ex. 1014 ¶¶ 32, 38; Ex. 1003 ¶¶ 436–450); *see also* Pet. 8–11 (explaining rank multiplication as known in the art). Petitioner further argues that it would have been obvious to use chip select die interconnects to conduct chip select signals based on, among other things, Rajan's disclosure of using separate chip signals for each chip. Pet. 73–76 (citing Ex. 1014 ¶¶ 29–30, 49, Figs. 2, 5; Ex. 1015, 6:34–48, Fig. 4; Ex. 1019, 12–13, Fig. 2; Ex. 1022, 319, Ex. 1023, 2–4, 9; Ex. 1016 ¶ 38; Ex. 1003 ¶¶ 451–461).

Claim 16 recites,

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.

Petitioner's contentions here are similar to its contentions for claim 15. *See* Pet. 76–78.

Patent Owner argues that “Petitioner has not shown that its relied-on non-DRAM chips would utilize chip-select signals.” PO Resp. 51–52 (citing Ex. 2023 ¶¶ 169–171, 193; Ex. 1015, 6:34–38; Exs. 2043, 2044). Patent Owner asserts that “[t]here is also no evidence that the concept of ‘rank’/‘rank multiplication,’ would apply to non-DRAM memory systems.” PO Resp. 52 n.8 (citing Ex. 1022, 413). We address these arguments above in § II.D.2.d.2 and disagree with them for the reasons explained above.

Patent Owner also argues that “TSVs that penetrate through multiple chips were difficult and expensive to form at the time of the invention” and that “one failed TSV would cause the entire die stack to be discarded.” PO Resp. 52 (citing Ex. 2023 ¶¶ 195–196; Ex. 2007; Ex. 2046). Patent Owner includes a parenthetical regarding “a JEDEC presentation that indicated commercially viable speed bin yield would not be viable until 2016.” PO Resp. 52 (citing Ex. 2046). According to Patent Owner, a person of ordinary skill in the art “would not have incurred the cost or the risk to build TSVs that are not needed for the memory chips’ operation.” PO Resp. 52 (citing Ex. 2023 ¶¶ 195–196). The evidence of record shows that a known configuration was to have multiple dies sharing TSVs, as discussed above in § II.D.2.d.2. *See* Ex. 1025 (Foster); *see also* Ex. 2023 ¶ 164 (Dr. Brogioli testifying with respect to Foster about “the dies forming the substack that shares a TSV”). The record evidence shows that using TSVs was within the

skill of a person of ordinary skill in the art. *See* Exs. 1014 (Kim), 1025 (Foster). We find that this evidence shows that the use of TSVs was an obvious option for a person of ordinary skill in the art, notwithstanding Patent Owner’s commercial viability arguments. *See Uber Technologies, Inc. v. X One, Inc.*, 957 F.3d 1334, 1340 (Fed. Cir. 2020) (“Because a person of ordinary skill ‘has good reasons to pursue the known options within his or her technical grasp,’ § 103 bars the patentability of such obvious variations.” (quoting *KSR*, 550 U.S. at 421)); *cf. CFMT, Inc. v. Yieldup Int’l Corp.*, 349 F.3d 1333, 1338 (Fed. Cir. 2003) (“Title 35 does not require that a patent disclosure enable one of ordinary skill in the art to make and use a perfected, commercially viable embodiment absent a claim limitation to that effect.”).

We find Petitioner’s contentions for claims 15 and 16 persuasive.

Rajan discloses the following:

If the buffer chip is emulating a memory device which has a larger capacity than each of the physical DRAM chips in the stack, the buffer chip may receive from the host system’s memory controller more address bits than are required to address any given one of the DRAM chips. In this instance, the extra address bits may be decoded by the buffer chip to individually select the DRAM chips, utilizing separate chip select signals (not shown) to each of the DRAM chips in the stack.

Ex. 1015, 6:30–38. We agree with Petitioner that this disclosure teaches rank multiplication that results in a number of “second chip select signals” that equals the number of array dies, as recited in claims 15 and 16. *See* Pet. 72–73; *see also* Ex. 1003 ¶¶ 445–447 (explaining that Rajan’s disclosure in column 6, lines 34–38 shows rank multiplication). This meets the claim interpretation discussed above in § II.C.2 because each chip has its own chip select signal such that none of the second chip select signals enables multiple dies at once.

Based on Petitioner's persuasive contentions and evidence, summarized above, and for the reasons discussed above, we conclude that claims 15 and 16 are unpatentable as obvious over the combined teachings of Kim, Rajan, and Wyman.

*E. Patent Owner's Procedural Arguments*

Patent Owner makes various procedural arguments with which we disagree as discussed below.

*1. Alleged Improper Incorporation by Reference*

Patent Owner asserts that, "[t]hroughout the Petition, Petitioner tries to evade the word limit by including no substantive analysis in the Petition itself, but attempting to incorporate paragraphs and paragraphs of its expert's declaration." PO Resp. 19. As example, Patent Owner notes that "Petitioner devotes two paragraphs in Ground 1" for limitations 1.d.1. and 1.d.2 but cites eleven paragraphs of its expert declaration. PO Resp. 19. Patent Owner asserts that "[t]his is improper, and such arguments should be deemed as absent from the Petition itself." PO Resp. 19.

We disagree with Patent Owner that there is "no substantive analysis" for certain limitations, and we refer to our findings above as to the Petition's analysis. Furthermore, there is no prohibition on expert declarations that provide more discussion of claim limitations than a petition. Indeed, our Rules provide that "[e]xpert testimony that does not disclose the underlying facts or data on which the opinion is based is entitled to little or no weight." 37 C.F.R. § 42.65(a). Thus, an expert declaration should provide the necessary factual bases underlying the declarant's opinion, and the expert is not limited to only what a petitioner asserts in a petition.

## 2. *Alleged Improper New Arguments*

Patent Owner argues that having each die in the stack in a separate rank with its own chip select signal is a “new argument” that we should not consider. PO Sur-reply 22. At oral argument, Patent Owner’s counsel stated, “I heard so many times, I couldn’t count, oh, our combination is a chip select for each die. So, in Riho, that would be 16 chip selects, and in Kim plus Rajan, that would be a chip select for each separate chip. That is nowhere in their petition.” Tr. 36:25–37:2.

We disagree with Patent Owner because the Petition cites Rajan’s disclosure of providing a separate chip select signal for each chip in the stack for claims 15 and 16. *See* Pet. 72, 78 (citing Ex. 1015, 6:34–38). Thus, Patent Owner’s contention that this is “nowhere in their petition” is incorrect. Petitioner’s reliance on this disclosure for claim 1 in the Reply was an appropriate response to Patent Owner’s arguments about collisions. *See Apple Inc. v. Andrea Elecs. Corp.*, 949 F.3d 697, 706 (Fed. Cir. 2020) (“[A]ny ambiguity as to whether Apple raised a new argument on reply is eliminated when we consider whether Apple’s reply arguments are responsive to arguments raised in Andrea’s Patent Owner Response.”).

Patent Owner is also incorrect in asserting that having each die in the stack in a separate rank is a new argument. The Petition states the following:

[W]hen implementing “rank multiplication” (discussed above, pp.8-11), it was common to use a “fork-in-the-road” arrangement with two data buses for two ranks of memory devices (shown below left, yellow and orange), with the option of adding two additional ranks of memory devices (*resulting in four ranks*) to the existing data busses (shown below right, two yellow, two orange), meaning two memory dies (e.g., two orange) would share a given data bus.

Pet. 31 (emphasis added; citing Ex. 1003 ¶ 161; Ex. 1026, Figs. 12–13).

This passage refers to annotated versions of Figures 12 and 13 of U.S. Patent Application Publication No. 2006/0277355 A1 (Ex. 1026 (“Ellsberry”)).

Thus, the Petition expressly proposes four separate ranks in the combination.

*See also* Pet. 8 (discussing rank multiplication and citing Ex. 1015, 6:30–7:67 (Rajan’s disclosure of using separate chip select signals for each chip)).

Indeed, Patent Owner’s declarant Dr. Brogioli testifies that, “in Ellsberry, the 4 ranks are in four separate physical ranks.” Ex. 2023 ¶ 111. Petitioner, therefore, put Patent Owner on notice of its contention of using separate ranks. The passage quoted above also expressly states that “it was common to use a ‘fork-in-the-road’ arrangement.” Pet. 31. Thus, we also disagree with Patent Owner’s assertions that “fork in the road” is not in the Petition. *See, e.g.*, Tr. 37:4–7 (“The concept of a fork in the road is the magical elixir that makes clear that Rajan has a shared data interconnect, data terminal, and that the prior art references would share that data terminal. That is nowhere in the petition or the declaration.”); 70:16–17 (“Where does the fork in the road . . . appear in their petition? It does not.”).

Patent Owner also asserts that Petitioner’s arguments about known solutions to collisions are “new arguments that should be disregarded.” PO Sur-reply 3; *see also* PO Sur-reply 17–18 (asserting that Petitioner’s “suggest[ion] that Kim’s timing for read and write operations would change in the Kim-Rajan combination . . . is a new theory”). We disagree. Patent Owner asserts in its Response that having dies share a TSV would result in collisions. PO Resp. 32–44. In this case, Petitioner’s arguments that there were known solutions to data collision issues are directly responsive to Patent Owner’s collision arguments. *See Apple*, 949 F.3d at 706; *see also*

*Rembrandt Diagnostics, LP v. Alere, Inc.*, 76 F.4th 1376, 1385 (Fed. Cir. 2023) (finding proper a “reply argument discussing cost and time savings [that] has a nexus to [a] prior argument and is responsive”).

Furthermore, Petitioner’s discussion of the skilled artisan’s knowledge of collision avoidance solutions is proper because it “it is used ‘to document the knowledge that skilled artisans would bring to bear in reading the prior art identified as producing obviousness.’” *Anacor Pharms., Inc v. Iancu*, 889 F.3d 1372, 1380–81 (Fed. Cir. 2018) (quoting *Genzyme Therapeutic Prods. Ltd. P’ship v. Biomarin Pharm. Inc.*, 825 F.3d 1360, 1369 (Fed. Cir. 2016)).

We also find Petitioner’s reliance on *Lockwood* to be appropriate in pointing out that the prior art contains more detail on avoiding collisions than the ’160 patent, which is directly responsive to Patent Owner’s criticisms of the asserted combination. *See* Pet. Reply 13 (citing *Lockwood*, 107 F.3d at 1570); *but see* Tr. 39:3–9 (Patent Owner’s counsel’s assertion that “[t]his Lockwood argument, an argument just like essentially every argument that was made by Samsung in this oral argument was an argument in reply. And the reason why it was put in reply is so their expert can’t have been tested on it, because if he would have been tested upon it, he would have told the truth, which is that there’s a basic physics behind the speed at which a given memory die can clear a load or clear a read.”).

Patent Owner also asserts that Petitioner’s argument that a person of ordinary skill in the art “could fabricate drivers of different sizes for different TSVs” rather than using “a single driver circuitry” is “a new argument.” PO Sur-reply 19. Petitioner’s reliance on different drivers of different sizes is not a new theory because the Petition expressly asserts this

and illustrates it, as discussed above in § II.D.2.e. In particular, Petitioner asserts that a person of ordinary skill in the art “would have been motivated to implement Wyman’s teachings in Kim to improve power efficiency by using smaller drivers (navy blue) for the shorter TSV1 (light green) and larger drivers (dark green) for the longer TSV2 (teal), as shown below,” referring to an annotated version of a portion of Kim’s Figure 3. Pet. 33–34; *see also* Pet. 50 (illustrating separate drivers). Patent Owner acknowledges as much by arguing that a person of ordinary skill in the art “would not have even used different drivers to drive TSV1 and TSV2 *as asserted*.” PO Resp. 45 (emphasis added). Thus, our Decision relies on the Petition’s express assertions of separate drivers.

Patent Owner further asserts that reliance on non-DRAM devices such as NAND flash and SRAM is improper and “a complete violation of due process” because “the petition makes no discussion of NAND Flash and no discussion of SRAM as being a basis for the combination.” Tr. 55:12–17. We disagree. Before institution, we authorized Petitioner to file a preliminary reply addressing the district court’s construction of “array dies,” and we authorized Patent Owner to file a preliminary sur-reply. Ex. 3001. In the preliminary reply, Petitioner cited disclosure in Rajan that allows for “any type of memory whatsoever” to be used, including NAND flash and SRAM. Paper 9 at 3 (quoting Ex. 1015, 15:3–9). The Petition also cites this disclosure. *See* Pet. 60 (citing Ex. 1015, 15:3–12). The Decision on Institution discussed this disclosure in view of Patent Owner’s argument that array dies must be different from Rajan137’s DRAM circuits. Inst. Dec. 24–25. Thus, Patent Owner was provided with notice and an opportunity to be heard. *See Apple Inc. v. Corephotonics, Ltd.*, 81 F.4th

1353, 1361 (Fed. Cir. 2023) (stating that the Board’s “decisions must be reached only after the parties have been provided fair notice and an opportunity to be heard”).

*F. Obviousness over Riho, Rajan, and Riho2  
(Claims 1–20)*

Because we determine that all challenged claims are unpatentable as discussed above, we need not separately assess the ground of unpatentability based on the combination of Riho, Rajan, and Riho2. 35 U.S.C. § 318(a) (“If an inter partes review is instituted and not dismissed under this chapter, the Patent Trial and Appeal Board shall issue a final written decision with respect to the patentability of any patent claim challenged by the petitioner and any new claim added under section 316(d).”); *Bos. Sci. Scimed, Inc. v. Cook Grp. Inc.*, 809 F. App’x 984, 990 (Fed. Cir. 2020) (nonprecedential) (“We agree that the Board need not address issues that are not necessary to the resolution of the proceeding.”).

### III. PETITIONER’S MOTION TO EXCLUDE

Petitioner filed a Motion to Exclude Exhibits 2024 and 2026–2031. Paper 34. Because we do not rely on this evidence in a manner adverse to Petitioner, we dismiss the Motion to Exclude as moot.

#### IV. CONCLUSION<sup>14</sup>

For the reasons discussed above, we determine that Petitioner has proven, by a preponderance of the evidence, that claims 1–20 of the '160 patent are unpatentable, as summarized in the following table:

| <b>Claim(s)</b>            | <b>35<br/>U.S.C. §</b> | <b>Reference(s)/Basis</b>          | <b>Claims<br/>Shown<br/>Unpatentable</b> | <b>Claims<br/>Not Shown<br/>Unpatentable</b> |
|----------------------------|------------------------|------------------------------------|------------------------------------------|----------------------------------------------|
| 1–20                       | 103(a)                 | Kim, Rajan,<br>Wyman               | 1–20                                     |                                              |
| 1–20                       | 103(a)                 | Riho, Rajan,<br>Riho <sup>15</sup> |                                          |                                              |
| <b>Overall<br/>Outcome</b> |                        |                                    | 1–20                                     |                                              |

---

<sup>14</sup> Should Patent Owner wish to pursue amendment of the challenged claims in a reissue or reexamination proceeding subsequent to the issuance of this decision, we draw Patent Owner's attention to the April 2019 *Notice Regarding Options for Amendments by Patent Owner Through Reissue or Reexamination During a Pending AIA Trial Proceeding*. See 84 Fed. Reg. 16,654 (Apr. 22, 2019). If Patent Owner chooses to file a reissue application or a request for reexamination of the challenged patent, we remind Patent Owner of its continuing obligation to notify the Board of any such related matters in updated mandatory notices. See 37 C.F.R. § 42.8(a)(3), (b)(2).

<sup>15</sup> As explained above, because we determine that the challenged claims are unpatentable based on the combination of Kim, Rajan, and Wyman, we decline to address this ground.

V. ORDER

Accordingly, it is

ORDERED that claims 1–20 of the '160 patent have been shown to be unpatentable;

FURTHER ORDERED that Petitioner's Motion to Exclude (Paper 34) is *dismissed*; and

FURTHER ORDERED that, because this is a Final Written Decision, parties to the proceeding seeking judicial review of the Decision must comply with the notice and service requirements of 37 C.F.R. § 90.2.

For PETITIONER:

Eliot D. Williams  
Theodore W. Chandler  
Ferenc Pazmandi  
Michael E. Knierim  
Brianna L. Potter  
BAKER BOTTS L.L.P.  
eliot.williams@bakerbotts.com  
ted.chandler@bakerbotts.com  
ferenc.pazmandi@bakerbotts.com  
michael.knierim@bakerbotts.com  
brianna.potter@bakerbotts.com

For PATENT OWNER:

Hong Annita Zhong  
IRELL & MANELLA  
hzhong@irell.com

UNITED STATES PATENT AND TRADEMARK OFFICE

---

BEFORE THE PATENT TRIAL AND APPEAL BOARD

---

SAMSUNG ELECTRONICS CO., LTD, MICRON TECHNOLOGY, INC.,  
MICRON SEMICONDUCTOR PRODUCTS, INC., and  
MICRON TECHNOLOGY TEXAS LLC,<sup>1</sup>  
Petitioner,

v.

NETLIST, INC.,  
Patent Owner.

---

IPR2022-01428  
Patent 8,787,060 B2

---

Before JON M. JURGOVAN, DANIEL J. GALLIGAN, and  
SHEILA F. McSHANE, *Administrative Patent Judges*.

GALLIGAN, *Administrative Patent Judge*.

JUDGMENT  
Final Written Decision  
Determining All Challenged Claims Unpatentable  
Dismissing Petitioner's Motion to Exclude  
*35 U.S.C. § 318(a)*

---

<sup>1</sup> Micron Technology, Inc., Micron Semiconductor Products, Inc., and Micron Technology Texas LLC filed a motion for joinder and a petition in IPR2023-00882 and have been joined as petitioners to this proceeding.

## I. INTRODUCTION

In this *inter partes* review, Samsung Electronics Co., Ltd. (“Samsung”), Micron Technology, Inc., Micron Semiconductor Products, Inc., and Micron Technology Texas LLC (collectively “Petitioner”) challenge the patentability of claims 1–34 of U.S. Patent No. 8,787,060 B2 (“the ’060 patent,” Ex. 1001), which is assigned to Netlist, Inc. (“Patent Owner”).

We have jurisdiction under 35 U.S.C. § 6. This Final Written Decision, issued pursuant to 35 U.S.C. § 318(a), addresses issues and arguments raised during the trial in this *inter partes* review. For the reasons discussed below, we determine that Petitioner has proven by a preponderance of the evidence that claims 1–34 of the ’060 patent are unpatentable. *See* 35 U.S.C. § 316(e) (2018) (“In an *inter partes* review instituted under this chapter, the petitioner shall have the burden of proving a proposition of unpatentability by a preponderance of the evidence.”).

### A. Procedural History

Samsung filed a Petition (Paper 1, “Pet.”) challenging claims 1–34 of the ’060 patent on the following grounds:

| Claims Challenged       | 35 U.S.C. §         | References/Basis                     |
|-------------------------|---------------------|--------------------------------------|
| 1–6, 8–14, 16–19, 29–34 | 103(a) <sup>2</sup> | Kim, <sup>3</sup> Rajan <sup>4</sup> |
| 1–14, 16–19, 29–34      | 103(a)              | Kim, Rajan, Riho <sup>5</sup>        |

---

<sup>2</sup> The Leahy-Smith America Invents Act (“AIA”), Pub. L. No. 112-29, 125 Stat. 284, 287–88 (2011), amended 35 U.S.C. § 103 effective March 16, 2013. Petitioner applies the pre-AIA version of § 103 based on an effective filing date before March 16, 2013. Pet. 3–4.

<sup>3</sup> US 2011/0103156 A1, published May 5, 2011 (Ex. 1014).

<sup>4</sup> US 8,041,881 B2, issued Oct. 18, 2011 (Ex. 1015).

<sup>5</sup> US 2011/0026293 A1, published Feb. 3, 2011 (Ex. 1016).

| <b>Claims Challenged</b> | <b>35 U.S.C. §</b> | <b>References/Basis</b>         |
|--------------------------|--------------------|---------------------------------|
| 1–6, 8–34                | 103(a)             | Kim, Rajan, Wyman <sup>6</sup>  |
| 1–14, 16–19, 29–34       | 103(a)             | Riho, Rajan                     |
| 1–34                     | 103(a)             | Riho, Rajan, Riho2 <sup>7</sup> |

Pet. 3. Patent Owner filed a Preliminary Response. Paper 6. With Board authorization (Ex. 3001), Samsung filed a preliminary reply (Paper 9) to the Preliminary Response, and Patent Owner filed a preliminary sur-reply (Paper 10). Trial was instituted on the asserted grounds of unpatentability. Paper 13 (“Inst. Dec.”) at 28.

During the trial, Patent Owner filed a Response (Paper 20, “PO Resp.”), and Samsung filed a Reply (Paper 24, “Pet. Reply”). After Samsung’s Reply, Micron Technology, Inc., Micron Semiconductor Products, Inc., and Micron Technology Texas LLC were joined as petitioners to this proceeding based on a petition and a motion for joinder in IPR2023-00882. Paper 25. After this joinder, Patent Owner filed a Sur-reply (Paper 28, “PO Sur-reply”).

Petitioner filed a motion to exclude certain evidence. Paper 34. Patent Owner opposed the motion (Paper 35), and Petitioner filed a reply in support of the motion (Paper 36).

An oral hearing was held on January 11, 2024, a transcript of which appears in the record. Paper 46 (“Tr.”).

Petitioner relies on testimony from Andrew Wolfe, Ph.D. Ex. 1003. Patent Owner relies on testimony from Michael C. Brogioli, Ph.D. Ex. 2023. The parties have entered in the record transcripts of the

---

<sup>6</sup> US 7,969,192 B2, issued June 28, 2011 (Ex. 1017).

<sup>7</sup> US 2010/0195364 A1, published Aug. 5, 2010 (Ex. 1018).

depositions of these declarants. Exs. 2025 (Wolfe Deposition), 1052 (Brogioli Deposition).

*B. Real Parties in Interest*

The identified real parties in interest on the petitioner side are the following: Samsung, Samsung Semiconductor, Inc., Micron Technology, Inc., Micron Semiconductor Products, Inc., and Micron Technology Texas LLC. Pet. 1; IPR2023-00882, Paper 2 at 1.

Patent Owner identifies itself as the real party in interest. Paper 3 at 1.

*C. Related Matters*

As required by 37 C.F.R. § 42.8(b)(2), the parties identify various related matters. Pet. 1; Paper 3 at 1; IPR2023-00882, Paper 2 at 1. We are issuing concurrently a final decision in IPR2022-01427 involving related U.S. Patent No. 9,318,160 B2 (“the ’160 patent”).

*D. The ’060 Patent and Illustrative Claim*

The ’060 patent relates to computer memory devices and, more specifically, to reducing the load of drivers in memory. Ex. 1001, 1:18–21. As background, the ’060 patent describes an existing memory package with reference to Figure 1A, reproduced at right, which shows memory package 100 with control die 130 and three array dies 110. Ex. 1001, 1:30–44. Control die 130 has driver 134, which drives data signals from control die 130 to each array die via interconnect 142, and driver 140, which drives command and address




FIG. 1A

signals to each array die. Ex. 1001, 1:35–42. The '060 patent explains that “a load exists on each of the drivers 134, 140 . . . by virtue of the drivers being in electrical communication with the corresponding die interconnects and the corresponding circuitry of the array dies.” Ex. 1001, 2:8–11. The '060 patent discloses that, “to drive a signal along a die interconnect, a driver typically must be large enough to overcome the load on the driver” and that “a larger driver not only consumes more space on the control die, but also consumes more power.” Ex. 1001, 2:11–15.

The '060 patent states that driver size and power consumption can be reduced “by increasing the number of die interconnects and reducing the number of array dies that are in electrical communication with each die interconnect.” Ex. 1001, 4:22–26. The '060 patent illustrates an exemplary configuration in Figure 2, reproduced at right, which shows memory package 200 with control die 230 and four array dies 210a–210d. Ex. 1001, 5:12–13. Control die 230 is connected to array dies 210a and 210b by die interconnect 220a, as shown by the darkened circles. Ex. 1001, 5:23–6:1. Die interconnect 220b connects control die 230 to array dies 210c and 210d, as shown by the darkened circles, but die interconnect




FIG. 2

220b is not electrically connected to array dies 210a and 210b, as shown by the unfilled circles. Ex. 1001, 6:9–23. The '060 patent also states that “[e]xamples of die interconnects include, but are not limited to, through-silicon vias (TSV), conducting rods, wire bonds, and pins.” Ex. 1001, 5:51–53.

Claim 1 is illustrative and is reproduced below with Petitioner’s claim element identifiers in brackets.

1. [1.a] A memory package, comprising:

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

[1.c] 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;

[1.d.1] 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, [1.d.2] 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

[1.e.1] a control die comprising [1.e.2] at least a first data conduit between the first die interconnect and a first terminal of the plurality of input/output terminals, [1.e.3] and at least a second data conduit between the second die interconnect and the first terminal, the first terminal being a data terminal, [1.e.4] 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.

## II. ANALYSIS

### A. *Principles of Law*

A patent claim is unpatentable under 35 U.S.C. § 103(a) if the differences between the claimed subject matter and the prior art are such that the subject matter, as a whole, would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. *KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 406 (2007). The question of obviousness is resolved on the basis of underlying factual determinations including (1) the scope and content of the prior art; (2) any differences between the claimed subject matter and the prior art; (3) the level of ordinary skill in the art; and (4) any secondary considerations, if in evidence. *Graham v. John Deere Co. of Kan. City*, 383 U.S. 1, 17–18 (1966).

### B. *Level of Ordinary Skill in the Art*

Citing the testimony of its declarant, Dr. Wolfe, Petitioner contends that a person of ordinary skill in the art “would have had an advanced degree in electrical or computer engineering, or a related field, and two years working or studying in the field of design or development of memory systems, or a bachelor’s degree in such engineering disciplines and at least three years working in the field.” Pet. 5 (citing Ex. 1003 ¶ 60). Petitioner also asserts that “[a]dditional training can substitute for educational or research experience, and vice versa,” and Petitioner identifies particular technology with which a person of ordinary skill in the art would have been familiar, including JEDEC (Joint Electron Devices Engineering Council) industry standards, DRAM (dynamic random access memory) and SDRAM

(synchronous DRAM) memory modules, and “the structure and operation of circuitry used in stacked memory devices.” Pet. 5 (citing Ex. 1003 ¶ 60).

Patent Owner states that, “[f]or purposes of this proceeding, Patent Owner applies the education and work-experience level specified on Petition, page 5.” PO Resp. 6. Patent Owner then contends that “Petitioner does not allege that a [person of ordinary skill in the art] would have been familiar with designs and operations of non-DRAM devices or non-JEDEC specifications.” PO Resp. 6. Patent Owner contends, therefore, that “under Petitioner’s unpatentability theories” involving non-DRAM devices, persons of ordinary skill in the art “would be engaged in the project of making changes to memory devices that they have no experience with.” PO Resp. 6 (citing Ex. 2023 ¶ 30).

We do not agree with Patent Owner’s assertions regarding non-DRAM devices. The obviousness inquiry under 35 U.S.C. § 103(a) requires us to determine

if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains.

As Patent Owner confirmed during oral argument, the term “array dies” in the claims encompasses non-DRAMs. Tr. 53:4–7. Because the subject matter pertains to configurations involving non-DRAM devices, the “person having ordinary skill in the art,” by definition under the statute, would have had knowledge of non-DRAM devices. To the extent Petitioner’s proposed skill level, on which the parties agree, does not make that clear, we do.

Thus, we accept the uncontested assessment offered by Petitioner with the exception of the qualifier “at least,” which introduces vagueness as to the

amount of experience. We also determine that a person of ordinary skill in the art would have been familiar with designs and operations of non-DRAM devices for the reasons discussed above.

### *C. Claim Construction*

We interpret claim terms using “the same claim construction standard that would be used to construe the claim in a civil action under 35 U.S.C. 282(b).” 37 C.F.R. § 42.100(b).

#### *1. Array Die*

The challenged claims recite the term “array dies.” In related litigation, a district court construed this term based on the following argument made by the applicant during prosecution of the ’060 patent regarding U.S. Patent Application Publication No. 2008/0025137 A1 (Ex. 1011 (“Rajan 137”)): “Rajan[137] does not disclose ‘a plurality of stacked array dies.’ Rajan[137] merely stacks DRAM circuits 206A–D, which are different from array dies.” Ex. 2004, 31–32 (quoting Ex. 1002 (’060 patent prosecution history), 465 (January 13, 2014 Amendment at 10)). The district court found that Patent Owner “did not explain the structural difference between array dies and DRAM circuits, or even how they might be stacked differently” but, nonetheless, found that Patent Owner “*structurally* distinguished ‘stacked DRAM circuits’ from ‘stacked array dies’ to obtain the patent.” Ex. 2004, 32. Based on this prosecution history, the district court adopted the defendants’ proposed construction that “array die” means “array die that is different from a DRAM circuit.” Ex. 2004, 32. Samsung is a defendant in that litigation.

Patent Owner argues that it “agreed to be bound by this construction and proceeded with the jury trial on that construction.” PO Resp. 17 (citing

Exs. 2010, 2011). Yet, during the trial in this proceeding, Patent Owner qualifies the district court's construction to suggest that the term "array dies" means dies that are different from a particular kind of DRAM die. *See, e.g.*, PO Resp. 44 ("[T]he Petition contains no analysis on how or why Kim's memory chips are different from Rajan 137's DRAM circuits."); *see also* Inst. Dec. 10–11 (addressing similar pre-institution arguments); Tr. 58:11–15 (Patent Owner's counsel asserting that Rajan 137 "describes, for instance, monolithic memory circuits may take the form of, and then it lists an example. And it lists DRAM as an example of a monolithic memory circuit. So, we take DRAM circuit to be a monolithic memory circuit that holds state while powered."). The district court's construction, which Patent Owner asserts it adopted, does not contain any such qualification.

In our analysis below, we apply the district court's construction, and we find that the combination of Kim and Rajan teaches array dies that are not DRAMs and, therefore, are "different from a DRAM circuit." Therefore, we need not address any further qualifications of the district court's construction or construe the term "DRAM circuit" to determine what Patent Owner actually disclaimed, if anything. *See Omega Eng'g, Inc. v. Raytek Corp.*, 334 F.3d 1314, 1324 (Fed. Cir. 2003) ("[P]rosecution disclaimer promotes the public notice function of the intrinsic evidence and protects the public's reliance on definitive statements made during prosecution."); *see also Genuine Enabling Tech. LLC v. Nintendo Co.*, 29 F.4th 1365, 1374 (Fed. Cir. 2022) ("If the challenged statements are ambiguous or amenable to multiple reasonable interpretations, prosecution disclaimer is not established."). Based on our findings as to non-DRAM

dies, we also need not determine whether the term “array die” should be construed to also encompass DRAM circuits.

## 2. *Chip Select*

Claims 6, 11–14, 16–19, and 29–34 recite “chip select signals” or “chip select conduits.” Patent Owner argues that the district court construed the terms “chip select signals” and “chip select conduits” to “exclude[] situations in which a chip select signal could enable multiple array dies at once” and contends that we should adopt this construction. PO Resp. 18–19 (citing Ex. 2004, 33–34). For the reasons explained below in section II.D, we determine that Petitioner has demonstrated unpatentability under the district court’s construction, and, therefore, we need not further address this claim construction to render a final decision on patentability.

## 3. *Group of Array Dies*

Petitioner argues that the phrase “a first group of array dies,” which is recited in the independent claims, can be just one die. Pet. Reply 7–8; *see also* Pet. 37 (arguing that multiple chips would have been obvious “[i]nsofar as one might argue that ‘first group of array dies’ requires multiple slave chips coupled to” one through-silicon via (TSV)).<sup>8</sup> Patent Owner disagrees and asserts that the first group must have plural dies. PO Resp. 20. We need not resolve this dispute because, as explained below in § II.D.2.d, we find persuasive Petitioner’s contentions that the combination of Kim and Rajan teaches a first group that has multiple array dies.

---

<sup>8</sup> We omit the emphasis added by Petitioner for reference names and claim language when quoting Petitioner’s arguments.

#### 4. *Driver Size*

Various claims recite the term “driver size.” Neither party argues for an express construction of this term, but Patent Owner notes that a district court in related litigation construed “driver size” as “driver physical size.” PO Sur-reply 23; *see* Ex. 1066 (district court claim construction order), 13, 43. As explained below in § II.F, we find that the combination of Kim, Rajan, and Wyman teaches the subject matter under this construction, and we need not address the construction of this term further.

#### 5. *Remaining Terms*

We determine that no other terms or phrases require express construction. *See Realtime Data, LLC v. Iancu*, 912 F.3d 1368, 1375 (Fed. Cir. 2019) (“The Board is required to construe ‘only those terms . . . that are in controversy, and only to the extent necessary to resolve the controversy.’” (quoting *Vivid Techs., Inc. v. Am. Sci. & Eng’g, Inc.*, 200 F.3d 795, 803 (Fed. Cir. 1999))).

#### *D. Obviousness over Kim and Rajan (Claims 1–6, 8–14, 16–19, 29–34)*


Petitioner asserts that claims 1–6, 8–14, 16–19, and 29–34 are unpatentable as obvious over the combined teachings of Kim and Rajan. Pet. 3, 21–81. Patent Owner opposes. PO Resp. 20–51; PO Sur-reply 1–20.

##### *1. Overview of the Prior Art*

Kim discloses a semiconductor memory apparatus including stacked chips, as depicted in Figure 5 below.

FIG.5

3



Kim's Figure 5 above shows a memory including a plurality of stacked chips, including main chip C0 and first and second slave chips C1 and C2, connected by through-silicon vias (TSVs). Ex. 1014 ¶¶ 47–48.


Rajan discloses using an interface circuit to emulate characteristics of particular memory even though the physical memory may not have those characteristics, which allows the “use of low cost memory chips in manufacturing high capacity memory modules.” Ex. 1015, 1:51–54, 3:24–43. Rajan explains that the interface circuit may comply with JEDEC standards. Ex. 1015, 4:20–24.

We discuss additional pertinent details of the references below.

## 2. Claim 1

### a) Overview of Petitioner's contentions for claim 1

To illustrate its contentions based on Kim, Petitioner provides the following annotated version of Kim's Figure 5:



Pet. 27. In the annotated figure above, Petitioner maps depicted components to elements of claim 1 as follows: chip C1 (yellow) as "first group of array dies"; chip C2 (orange) as "second ground of at least one array die"; chip C0 (green) as "control die"; TSV1 (light green) from C0 to C1 as "first die interconnect"; TSV2 (teal) from C0 to C2 as "second die interconnect."

Pet. 27.

Petitioner argues that a person of ordinary skill in the art would have been motivated to make Kim's memory package compatible with JEDEC standards based on Rajan's teachings. Pet. 24–27. Petitioner also argues that a person of ordinary skill in the art would have been motivated to have multiple memory chips share a TSV because this was a known option and "would not require creating new TSVs (which would add space and

circuitry) and would allow emulating the JEDEC standard required by external devices.” Pet. 28–29.

*b) Preamble (1.a)*

Petitioner argues that Kim’s Figure 5 depicts a “memory package,” as recited in the preamble. Pet. 30 (citing Ex. 1014 ¶¶ 46–47; Ex. 1003 ¶¶ 184–188). Patent Owner does not dispute this contention. We agree with Petitioner, and we find, that Kim’s disclosure of semiconductor memory apparatus 3 in Figure 5 teaches a “memory package.” In view of this finding, we need not decide whether the preamble is limiting.

*c) Input/output terminals (1.b)*

For claim 1’s “plurality of input/output terminals via which the memory package communicates data and control/address signals with one or more external devices” (1.b), Petitioner argues that Kim’s Figure 5 shows input/output data terminals as “DQ,” that Kim’s Figure 2 shows control terminals as CS0 and CS1, and that Kim’s memory package would have control and address terminals to accept control and address signals in a JEDEC-compliant implementation. Pet. 31–34. Patent Owner does not dispute that the combination of Kim and Rajan teaches control, address, and data terminals, but Patent Owner does dispute the JEDEC aspect of Petitioner’s combination. We agree with Petitioner, and we find, that Kim teaches a memory package having data, address, and control terminals, and we further address the parties’ dispute about JEDEC below.

*d) Stacked array dies (1.c) and die interconnects (1.d.1, 1.d.2)*

Claim 1 recites “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” (1.c). Claim 1 further recites that the memory

package has “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,” (1.d.1) and “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” (1.d.2).

As noted above in reference to Petitioner’s annotated version of Kim’s Figure 5, Petitioner identifies Kim’s chip C1 as “a first group of array dies” and Kim’s chip C2 as “a second group of at least one array die.” Pet. 34–35 (citing Ex. 1014 ¶¶ 26, 48, Fig. 5; Ex. 1003 ¶¶ 211–219). Petitioner argues that having multiple dies in the first group all sharing the same die interconnect would have been obvious. Pet. 37–39 (citing Ex. 1014 ¶¶ 48, 50; Ex. 1024, 9:14–18, Fig. 5; Ex. 1015, Fig. 4; Ex. 1003 ¶¶ 219–227).<sup>9</sup> Petitioner argues that the annotated version of Rajan’s Figure 4 below shows “multiple stacked memory chips connected to the same data bus.” Pet. 38–39.

---

<sup>9</sup> Petitioner also argues that “a first group of array dies” can be just one die. Pet. Reply 7–8; *see also* Pet. 37. We do not rely on this argument and, therefore, need not resolve the parties’ dispute regarding whether this phrase encompasses a single die.




FIG. 4

Rajan's Figure 4 "shows a buffered stack of DRAM circuits having one address, control, and clock bus and two data busses." Ex. 1015, 2:6–7. In the annotated figure above, Petitioner maps depicted components to elements of claim 1 as follows: DRAMs 417C and 417D (yellow) as "first group of array dies"; DRAMs 417A and 417B (orange) as "second ground of at least one array die"; buffer chip 413 (green) as "control die"; data bus (light green) from buffer chip 413 to DRAMs 417C and 417D as "first die interconnect"; and data bus from buffer chip 413 to DRAMs 417A and 417B as "second die interconnect." Pet. 42.

Petitioner argues that Kim and Rajan are analogous art. Pet. 24 (citing Ex. 1003 ¶¶ 154–155). Patent Owner does not dispute that Kim and Rajan are analogous art to the '060 patent. *See* PO Resp. We agree with Petitioner that Kim and Rajan are analogous art because they are both in the same field of endeavor as the '060 patent—memory packages and in particular those involving stacked memory. Ex. 1001, 1:18–19 ("The present disclosure relates to memory devices and memory modules."), claim 1 ("memory

package” with “stacked array dies”); Ex. 1014 ¶ 3 (“Various embodiments of the present disclosure generally relate to a semiconductor memory apparatus . . . .”), Fig. 5 (stacked memory chips); Ex. 1015, 1:14–16 (“This invention relates generally to digital memory such as used in computers[], and more specifically to organization and design of memory modules such as DIMMs [(dual in-line memory modules)].”), Figs. 2–6 (stacked memories).

Petitioner argues that a person of ordinary skill in the art “would be motivated to create a package with an interface that complied with the well-known JEDEC standards, as taught by Rajan,” such that Kim’s memory package would be compatible with JEDEC. Pet. 24–26 (citing Ex. 1014 ¶ 38, Fig. 5; Ex. 1015, 3:52–54, 4:20–24, 5:36–43, 8:8–11, Figs. 4, 18; Ex. 1003 ¶¶ 151–157; Ex. 1019, 12, Fig. 3).

Petitioner also argues that, based on Kim’s disclosure of using “any number of” chips (Ex. 1014 ¶ 48), a person of ordinary skill in the art “would have been motivated to look at Rajan for the details about adding more memory chips in the stack.” Pet. 26 (citing Ex. 1014 ¶¶ 48, 50; Ex. 1003 ¶¶ 158–160). Petitioner argues that a person of ordinary skill in the art “would have understood that there were a finite number of known ways to connect additional dies, including having a subset of dies sharing a through-silicon via (TSV),” and that “[s]haring a TSV among two or more dies was a known option, as confirmed by other references at the time.” Pet. 28 (citing Ex. 1003 ¶¶ 161–163; Ex. 1015, 5:36–43, Figs. 2–6; Ex. 1025, Fig. 5). Petitioner also argues that a person of ordinary skill in the art “would have been motivated to connect additional memory dies to Kim’s TSV1 and TSV2 because it would not require creating new TSVs (which

would add space and circuitry) and would allow emulating the JEDEC standard required by external devices.” Pet. 28–29 (citing Ex. 1003 ¶ 164; Ex. 1015, 3:27–30, 3:52–61). Petitioner further argues that “it was common at the time to have multiple memory dies sharing a single data bus” and “to use a ‘fork-in-the-road’ arrangement with two data buses for two ranks of memory devices . . . with the option of adding two additional ranks of memory devices (resulting in four ranks) to the existing data busses,” such that two dies share a bus. Pet. 29 (citing Ex. 1003 ¶¶ 165–166; Ex. 1026, Figs. 12, 13).

As to claim 1’s “electrical communication” recitations (1.d.1 and 1.d.2), Petitioner argues that Kim’s Figure 5 shows die interconnect (TSV1) connected from chip C0 to chip C1 and not extending to chip C2 and die interconnect (TSV2) connected from chip C0 to chip C2 and passing through chip C1 without being connected. Pet. 39–40 (citing Ex. 1014, code (57), ¶¶ 45, 49, Fig. 5; Ex. 1003 ¶¶ 230–239). And, as discussed above, in Petitioner’s proposed combination, multiple dies would share a TSV (Pet. 28–29) such that each group would be in electrical communication with the TSVs shared by that group and not other TSVs.

Below, we address Patent Owner’s arguments pertinent to recitations 1.c, 1.d.1, and 1.d.2.

*(1) Array dies*

As noted above in § II.C.1, a district court construed “array die” to mean “array die that is different from a DRAM circuit.” Ex. 2004, 31–32. Patent Owner argues that “the Petition contains no analysis on how or why Kim’s memory chips are different from Rajan137’s DRAM circuits.” PO Resp. 44. Patent Owner’s argument, therefore, reflects an attempt to limit

the district court's construction to exclude only certain types of DRAM circuits, similar to arguments Patent Owner made before institution. *See* Inst. Dec. 11 ("Patent Owner appears to argue that the district court's construction is that the claimed array dies are different from Rajan 137's DRAM circuits 206A–D, not all DRAM circuits.").

As noted at institution, Kim does not even mention DRAM and thus is not limited to DRAM, and Rajan states that "any type of memory whatsoever" may be used in its disclosure and lists non-DRAM memories. Inst. Dec. 23 (quoting Ex. 1015, 15:3–9). Therefore, we find that the combination of Kim and Rajan teaches "array dies" under the district court's construction and under Patent Owner's interpretation of that construction.

Patent Owner further argues that,

if, in Petitioner's combination, the stacked dies are non-DRAM dies, neither the Petition nor the declaration explains why a [person of ordinary skill in the art] would be motivated to have the package emulate a DRAM interface or how the resulting structure would work given Petitioner's theory that the interface circuit should provide a JEDEC-compliant interface to the host, citing to SDRAM standards.

PO Resp. 45 (citing Pet. 31–33). According to Patent Owner, "Petitioner has not presented any evidence that with non-DRAM circuits such as NANDs or SRAMs [(static RAM)], a [person of ordinary skill in the art] would have or could have emulated the memory package as JEDEC-compliant DRAM devices." PO Resp. 46 (citing Pet. 8–11, 24–33; Ex. 2023 ¶¶ 173–177); *see also* PO Resp. 47–48 ("Petitioner has not shown that a [person of ordinary skill in the art] would have been motivated to use non-DRAM dies in combination with an interface circuit that emulates a JEDEC-compliant interface to the host.").

We disagree with Patent Owner's arguments because the evidence in support of Petitioner's position is the express disclosure of Rajan. As noted above, Petitioner argues that a person of ordinary skill in the art "would be motivated to create a package with an *interface* that complied with the well-known JEDEC standards, as taught by Rajan." Pet. 24 (emphasis added). Rajan discloses an "interface circuit [that] may take the form of or incorporate, or be incorporated into, a register, an AMB [(advanced memory buffer)], a buffer, or the like, and may comply with Joint Electron Device Engineering Council (JEDEC) standards, and may have forwarding, storing, and/or buffering capabilities." Ex. 1015, 4:20–24. Rajan explains that "the interface circuit presents to the system device an interface to emulated memory devices which differ in some aspect from the physical memory circuits which are actually present." Ex. 1015, 3:24–27. With reference to Figure 1, Rajan discloses that the "physical memory circuits may be any type of memory circuits." Ex. 1015, 2:59–60. Rajan further discloses that, "[a]lthough the embodiments described here show the stack consisting of multiple DRAM circuits, a stack may refer to any collection of memory circuits (e.g. DRAM circuits, flash memory circuits, or combinations of memory circuit technologies, etc.)." Ex. 1015. 4:51–55. Rajan still further discloses that the "physical memory circuits employed in practicing this invention may be any type of memory whatsoever, such as: DRAM, DDR DRAM, DDR2 DRAM, DDR3 DRAM, SDRAM, QDR DRAM, DRDRAM, FPM DRAM, VDRAM, EDO DRAM, BEDO DRAM, MDRAM, SGRAM, MRAM, IRAM, NAND flash, NOR flash, PSRAM, wetware memory, etc." Ex. 1015, 15:3–9. Thus, Rajan expressly discloses using non-DRAM dies as the memory circuits.

Patent Owner argues that “Rajan never suggests that its interface die would be emulating a different type of memory circuits from the backend memory.” PO Sur-reply 18 (citing Ex. 1015, 15:3–9). We disagree with this argument. As noted above, Rajan’s “interface circuit presents to the system device an interface to emulated memory devices which differ in some aspect from the physical memory circuits which are actually present” (Ex. 1015, 3:24–27), and we see nothing in Rajan that limits its teachings as Patent Owner’s argument suggests. Furthermore, Rajan discloses that the

memory subsystem includes a buffer chip 202 which presents the host system with emulated interface to emulated memory, and a plurality of physical memory circuits which, in the example shown, are DRAM chips 206A-D. In one embodiment, the DRAM chips are stacked, and the buffer chip is placed electrically between them and the host system. Although the embodiments described here show the stack consisting of multiple DRAM circuits, a stack may refer to any collection of memory circuits (e.g. DRAM circuits, flash memory circuits, or *combinations of memory circuit technologies, etc.*).

Ex. 1015, 4:45–55 (emphasis added). The fact that the memory stack may include “combinations of memory circuit technologies,” such as DRAM and flash, shows that the interface would emulate a memory different from at least one of the memory technologies in the combination when it presents as one type of memory to the host. Rajan also discloses that “the memory circuits may be symmetrical, meaning each has the same capacity, *type*, speed, etc., while in other embodiments they may be asymmetrical.”

Ex. 1015, 2:62–67 (emphasis added). Thus, Rajan discloses that symmetrical circuits are of the same type, suggesting that asymmetrical circuits can be of different types behind the interface.

Patent Owner also asserts in a parenthetical that Rajan’s disclosure at column 15, lines 3–9 “has no link to 4:20-24 which is part of the description for Fig. 1, related to ‘DRAM circuits.’” PO Sur-reply 18. We disagree with this assertion. Rajan at column 4, lines 20–24 discloses that the interface circuit may comply with JEDEC, and this is the “interface circuit” of the disclosed invention. *See* Ex. 1015, 1:65–67 (“FIG. 1 shows a system coupled to multiple memory circuits and an interface circuit according to one embodiment of this invention.”). Thus, Rajan’s disclosure that the “physical memory circuits employed in practicing *this invention* may be any type of memory whatsoever,” including non-DRAMs, (Ex. 1015, 15:3–9 (emphasis added)) is linked to the interface circuit of the invention, which is described throughout Rajan and is described as being JEDEC-compliant in column 4, lines 20–24.

For all of the reasons discussed above, we find that Rajan teaches using non-DRAM memories in a JEDEC-compliant memory package and, therefore, that Rajan provides an express motivation to “to use non-DRAM dies in combination with an interface circuit that emulates a JEDEC-compliant interface to the host.” *See* PO Resp. 47–48.

We also disagree with Patent Owner’s argument that “Petitioner has not presented any evidence that with non-DRAM circuits such as NANDs or SRAMs, a [person of ordinary skill in the art] . . . *could have* emulated the memory package as JEDEC-compliant DRAM devices.” *See* PO Resp. 46 (emphasis added) (citing Pet. 8–11, 24–33; Ex. 2023 ¶¶ 173–177). We disagree because the express disclosure of Rajan, discussed above, is evidence that a person of ordinary skill in the art could have used non-DRAM memories in a JEDEC-compliant DRAM package. Petitioner

cites Rajan’s Figure 18 (Pet. 24–25), which shows an embodiment of an interface circuit, and Rajan’s accompanying disclosures provide more details as to how the interface circuit functions. Ex. 1015, 14:6–62.

Furthermore, we have considered the cited testimony from Dr. Brogioli, which focuses on how NAND flash and SRAM both differ from DRAM. Ex. 2023 ¶¶ 173–177. When considered in view of the disclosure of Rajan, the listed differences do not indicate or suggest difficulties beyond the skill of an ordinarily skilled artisan to implementing JEDEC-compliant interfaces with non-DRAM memories. For example, Dr. Brogioli testifies that NAND “signaling, electrical and physical interface are drastically different than that of DRAM memory devices, or DRAM memory modules that are JEDEC or JEDEC adjacent.” Ex. 2023 ¶ 175. But Rajan discloses that memory characteristics emulated by the interface circuit “may be electrical in nature, physical in nature, logical in nature, pertaining to a protocol, etc.” Ex. 1015, 3:36–38. As to signaling in particular, Rajan discloses that “[t]he interface circuit may emulate the number of signals, type of signals, duration of signal assertion, and so forth” and “may combine multiple signals to emulate another signal.” Ex. 1015, 3:48–51. Dr. Brogioli also testifies about “vastly different timing” between NAND and DRAM (Ex. 2023 ¶ 175), but Rajan discloses that “[a]n example of an emulated protocol characteristic might be a timing” (Ex. 1015, 3:42–43) and that the interface circuit’s “emulation logic may, in various embodiments, alter a *timing*, value, latency, etc. of any of the address, control, clock, and/or data signals it sends to or receives from the system and/or the physical memory” (Ex. 1015, 14:55–59 (emphasis added)). Therefore, Dr. Brogioli’s testimony fails to account for Rajan’s disclosure. Dr. Brogioli’s

testimony as to the “different signaling mechanism” that SRAM uses (Ex. 2023 ¶ 177) is similarly unpersuasive.

Furthermore, Dr. Brogioli’s testimony addressing NAND flash and SRAM differences fails to appreciate that Rajan is not so limited and lists other non-DRAM memories that can be used, including NOR flash. *See* Ex. 1015, 15:3–9.

For the foregoing reasons, we find that the combination of Kim and Rajan teaches “array dies.”

*(2) Electrical communication*

Patent Owner argues that Petitioner has not shown that the combination of Kim and Rajan teaches a “first die interconnect in electrical communication with” multiple array dies. PO Resp. 20–44.

Patent Owner argues that Kim discloses that “each TSV is in communication with only a single chip.” PO Resp. 20 (citing Ex. 1014, Fig. 5; Ex. 2023 ¶ 84). Although Kim’s Figure 5 depicts TSV1 in communication with chip C1 and TSV2 in communication with chip C2, Kim discloses that “any number of . . . chips may be used” (Ex. 1014 ¶ 48), as Petitioner points out. *See* Pet. 26 (citing Ex. 1003 ¶¶ 158–160; Ex. 1014 ¶¶ 48, 50). Petitioner argues that this disclosure would have motivated a person of ordinary skill in the art “to look at Rajan for the details about adding more memory chips in the stack (resulting, e.g., in four chips . . . in two groups . . .).” Pet. 26. We agree with Petitioner, and we find, that Kim’s disclosure suggests having more than the two chips depicted in Figure 5. *See* Ex. 1014 ¶¶ 48 (“Although only one main chip and two slave chips are shown, it is to be understood that any number of main and slave chips may be used.”), 50 (“While the semiconductor memory apparatuses

having two ranks are explained with reference to FIGS. 2 and 5, a person having ordinary skill in the art will appreciate that the technical concept of the present invention can be applied to a semiconductor memory apparatus which are divided into three or more ranks.”); *see also* Ex. 1003 ¶ 159 (“A Skilled Artisan would have understood from this disclosure that Kim’s invention can include more slave chips (and more ranks) and that some of those additional slave chips proposed by Kim can be in the same group as chips C1 or C2 and connected to TSV1 or TSV2, respectively.”).

Petitioner contends that a person of ordinary skill in the art would have been motivated to implement a shared data bus for multiple memory chips as taught by Rajan (e.g., 415A and 415B, above) using, e.g., Kim’s TSV interconnects . . . , in part because a [person of ordinary skill in the art] would have understood that there were a finite number of known ways to connect additional dies, including having a subset of dies sharing a through-silicon via (TSV).

Pet. 28 (citing Ex. 1003 ¶¶ 161–163; Ex. 1015, 5:36–43, Figs. 2–6). Thus, Petitioner’s contention is that, given Kim’s express disclosure of having more than two memory chips, a person of ordinary skill would have looked to Rajan for details about adding more memory chips. Petitioner argues that Rajan teaches “a shared data bus for multiple memory chips” as depicted in Rajan’s Figure 4 (Pet. 28), and Patent Owner agrees. *See* PO Resp. 21 (“In Rajan, Figure 4, Rajan’s DRAM circuits 417A and 417B share one data bus and DRAM circuits 417C and 417D share another data bus.”).

Patent Owner does not dispute Petitioner’s contention (Pet. 28) that there are a finite number of known ways to connect additional dies, one of which being to have dies share a TSV. *See* PO Resp. In support of its contention that “[s]haring a TSV among two or more dies was a known

option,” as confirmed by other references, such as Foster, Petitioner provides the annotated figure below.




FIG. 5

Pet. 28 (citing Ex. 1025 (“Foster”),<sup>10</sup> Fig. 5). Foster’s Figure 5 shows a cross-sectional view of die stack 20 including first substack 78 of dies 21 and 22 and second substack 80 of dies 23 and 24. Ex. 1025, 7:28–45. Petitioner identifies first substack 78 as a “first group of array dies” in yellow and second substack 80 as a “second group of at least one array die” in orange. Pet. 28. In Figure 5, TSVs 30 (identified by Petitioner in light green) connect to the first substack but do not extend to the second substack, and pass-through vias (PTVs) 29 (identified by Petitioner in teal) go through the first stack and connect to TSVs in the second substack. Ex. 1025, 7:45–54. Foster explains that PTVs are “conductive pathways through each die with no connections to any circuitry on the die” and that TSVs are

<sup>10</sup> US 8,258,619 B2, filed Nov. 12, 2009, issued Sept. 4, 2012.

“conductive pathways through the dies that also connect to electronic circuitry(36) on the die.” Ex. 1025, 2:53–56.

Patent Owner and Dr. Brogioli do not disagree with Petitioner’s assertion that Foster shows that “[s]haring a TSV among two or more dies was a known option.” *See* Pet. 28. Rather, they address different issues. Patent Owner argues that Petitioner cites Foster “for die grouping.” PO Resp. 21 n.3. And Dr. Brogioli testifies that “Foster does not specify the relationship of the dies forming the substack that shares a TSV or PTV” or suggest a certain die grouping. Ex. 2023 ¶ 164. Regardless of whether Foster teaches a certain die grouping, we agree with Petitioner that Foster shows that “[s]haring a TSV among two or more dies was a known option.” *See* Pet. 28.

Petitioner’s contention, as noted above, is that a person of ordinary skill in the art “would have been motivated to implement a shared data bus for multiple memory chips as taught by Rajan . . . using, e.g., Kim’s TSV interconnects.” Pet. 28. Rajan’s Figure 4 is shown below.




**FIG. 4**

Rajan’s Figure 4 above “shows a buffered stack of DRAM circuits having one address, control, and clock bus and two data busses.” Ex. 1015, 2:6–7; *see also* Ex. 1015, 4:51–55 (“Although the embodiments described here show the stack consisting of multiple DRAM circuits, a stack may refer to any collection of memory circuits (e.g. DRAM circuits, flash memory circuits, or combinations of memory circuit technologies, etc.).”). As shown in Figure 4, memory chips 417A and 417B share data bus 415A, and memory chips 417C and 417D share data bus 415B. Ex. 1015, 5:39–42.

Dr. Brogioli testifies that “Rajan uses wire bonds for connecting its DRAM circuits to the data busses, and wire bonds ordinarily cannot be shared in the way that TSVs are shared.” Ex. 2023 ¶ 143 (citing Ex. 1011 ¶ 50). In support of this statement, Dr. Brogioli does not cite the Rajan reference on which Petitioner’s challenge is based; rather, he cites Rajan137, which states that “data signals may be wired as one common bus, several busses or as an individual bus to each DRAM circuit.” Ex. 1011 ¶ 50, *quoted in* Ex. 2023 ¶ 143. Rajan, however, does not include this disclosure and does not state that its dies are connected with wire bonds. Thus, we do not credit Dr. Brogioli’s testimony that Rajan uses wire bonds.

Dr. Brogioli also testifies that

[t]he ‘060 patent discloses that an example of a die interconnect is a wire bond. But that wire bond is different from a wire bond dangling from the side of a DRAM circuit such as those in Rajan. Instead, the type of wire bonds referenced by the ‘060 patent are wires internal to a die, acting, for example, like a redistribution line.

Ex. 2023 ¶ 39 n.1 (citing “e.g., the red and blue wires in DDX4-48” (Ex. 2030)<sup>11</sup>). Even if Rajan’s Figure 4 shows wire bonds, an opinion that we do not credit as explained above, Dr. Brogioli does not explain why a wire bond in Rajan would be different from the wire bonds that can be die interconnects in the ‘060 patent. *See* Ex. 1001, 5:51–53 (“Examples of die interconnects include, but are not limited to, through-silicon vias (TSV), conducting rods, wire bonds, and pins.”). In addition, it is not clear what “the red and blue wires in DDX4-48” (Ex. 2023 ¶ 39 n.1) show because the cited page is about “Accused HBM Products” and does not mention wire

---

<sup>11</sup> Dr. Brogioli appears to be referring to Ex. 2030. *See* Ex. 2023 ¶ 57 (citing “EX2030, DDX4-53”).

bonds. *See* Ex. 2030, DDX4-48. Thus, this testimony is not helpful in explaining any alleged differences in wire bonds. *See* 37 C.F.R. § 42.65(a) (“Expert testimony that does not disclose the underlying facts or data on which the opinion is based is entitled to little or no weight.”). Furthermore, we do not agree with Dr. Brogioli’s testimony that wire bond die interconnects in the ’060 patent “are wires internal to a die” (Ex. 2023 ¶ 39 n.1) because wires that are only internal to one die would not interconnect with another die, and, thus, would not be a “die interconnect.”

In view of the foregoing, we find that Rajan does not disclose the particular structure in which the shared data busses in Figure 4 are implemented. In other words, Rajan’s disclosure is agnostic as to the particular interconnect technology by which the memories that share data busses are connected to the buffer chip.

Dr. Brogioli also testifies that a “data bus can be shared without sharing interconnects, especially in Rajan,” “because Rajan uses wire bonds for connecting its DRAM circuits to the data busses.” Ex. 2023 ¶ 143 (citing Ex. 1011 ¶ 50). As discussed above, we do not credit Dr. Brogioli’s opinion that Rajan uses wire bonds because it is based on a different reference. Furthermore, the fact that a data bus can be shared without sharing interconnects is irrelevant because Petitioner’s proposed combination is to share interconnects (Kim’s TSVs). *See* Pet. 28. Rajan’s Figure 4 discloses a stack of memory circuits, as discussed above (*see* Ex. 1015, 2:6–7, 4:51–55), and Kim discloses using TSVs in stacks of memory chips, as also discussed above (*see* Ex. 1014 ¶¶ 48, 50, Fig. 5). Thus, Petitioner’s combination is a straightforward one – “implement a shared data bus for multiple memory

chips as taught by Rajan . . . using, e.g., Kim’s TSV interconnects.” Pet. 28 (discussed above).

At oral argument, Patent Owner’s counsel argued that the “TSV is the individual data wire going down to the bus that’s shared.” Tr. 65:12–13.

According to Patent Owner’s counsel,

Rajan talks about [two chips] sharing a data bus. So, I don’t see how, from Rajan, you would get the idea of a shared TSV when Rajan speaks solely about sharing data buses. And the data bus, under JEDEC standard, involves individual data wires coming from each of the chips.

Tr. 65:26–66:4. By this argument, Patent Owner’s counsel appears to be suggesting that, even if TSVs are used, they would just connect each chip to a shared data bus. But this is not the proposed combination, as stated above, which is to implement a shared bus per Rajan using Kim’s TSVs. *See* Pet. 28. Implementing a bus using TSVs is consistent with the evidence of record. *See* Ex. 1025, 1:39–42 (“A more recent approach for wafer stacking is to connect the signals together with vias, effectively sending a bus of signal lines vertically through a stack of dies.”), 2:56–64 (“[A]ggregations of PTVs and TSVs will often be used to effect connection of bus signals from a substrate up through the die stack to circuitry somewhere in the die stack.”).

Based on the foregoing, we find persuasive Petitioner’s contention that a person of ordinary skill in the art

would have been motivated to implement a shared data bus for multiple memory chips as taught by Rajan (e.g., 415A and 415B, above) using, e.g., Kim’s TSV interconnects . . . , in part because a [person of ordinary skill in the art] would have understood that there were a finite number of known ways to connect additional dies, including having a subset of dies sharing a through-silicon via (TSV).

*See* Pet. 28. Furthermore, we find persuasive Petitioner’s reasoning that a person of ordinary skill in the art “would have been motivated to connect additional memory dies to Kim’s TSV1 and TSV2 because it would not require creating new TSVs (which would add space and circuitry).” Pet. 28–29 (citing Ex. 1003 ¶ 164). As discussed above, Kim suggests having more than two dies, Rajan discloses how to share a data bus among multiples dies, and the evidence of record shows that TSVs were used to implement data busses. Ex. 1014 ¶¶ 48, 50; Ex. 1025, 1:39–42, 2:56–64. We find that using the existing TSVs of Kim would be beneficial to avoid having to create new ones, “which would require space and additional circuitry.” Ex. 1003 ¶ 164.

Having found persuasive the above contentions, we turn to the issue of whether the proposed combination would have been within the ability of a person of ordinary skill in the art. *See KSR*, 550 U.S. at 421 (noting that “a person of ordinary skill has good reason to pursue the known options *within his or her technical grasp*” (emphasis added)); *see also id.* at 417 (“[I]f a technique has been used to improve one device, and a person of ordinary skill in the art would recognize that it would improve similar devices in the same way, using the technique is obvious unless its actual application is beyond his or her skill.”).

Patent Owner argues that the asserted combination could result in four possible configurations and that each configuration either would be inoperable or would not meet the claims. PO Resp. 22 (citing Ex. 2023 ¶ 86). Patent Owner asserts that the

four possibilities are: (1) four chips belong to two ranks and two chips from the same rank share a common TSV; (2) four chips belong to two ranks and a chip from one rank shares a TSV with a chip from another rank; (3) four chips belong to four ranks and

each of the four ranks has its own TSV; and (4) four chips belong to four ranks and every two ranks share a TSV.

PO Resp. 22 (citing Ex. 2023 ¶ 86).

We need not address the first three configurations because, for the reasons explained below, we find Petitioner’s combination persuasive in the fourth configuration. Patent Owner argues that a person of ordinary skill in the art would not combine the teachings of Kim and Rajan in the fourth configuration because it would result in unavoidable data collisions.

PO Resp. 32–44. Patent Owner asserts that, “[a]s Dr. Brogioli testified, handling data collisions involves ‘a lot of complexities’ that [are] ‘beyond the scope of a person of skill.’” PO Sur-reply 8 (citing Ex. 1052, 283:17–284:12). Patent Owner argues, therefore, that “the requirement that at least one die interconnect be in connection with a group of array dies (plural) would result in inoperable apparatuses for Kim’s intended purposes.” PO Resp. 43 (citing Ex. 2023 ¶¶ 149–150); *see* Ex. 2023 ¶ 150 (“A person of ordinary skill in the art would not have had a reason or a reasonable expectation of success to modify Kim so that multiple dies share the same TSV signaling path.”).

Petitioner argues that techniques for avoiding collisions on shared data lines were well-known to persons of ordinary skill in the art, citing as one example Rajan’s teachings of timing solutions to avoid collisions. Pet. Reply 2–7 (citing, *inter alia*, Ex. 1015, 8:59–12:13, Figs. 9–15). For the reasons explained below, we find persuasive Petitioner’s contentions that persons of ordinary skill in the art knew how to deal with collisions.

The specification of the ’060 patent mentions collisions only once:

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.

Ex. 1001, 21:39–56 (emphasis added). The fact that the '060 patent mentions *reducing or preventing* signal collisions suggests that there is no prohibition on signal collisions. This is in contrast with Dr. Brogioli's more limited view that if there are any data collisions the resulting structure would be inoperable. Ex. 2023 ¶ 150. Moreover, the challenged claims do not recite that data collisions must be avoided, and "[t]he reasonable expectation of success requirement refers to the likelihood of success in combining references *to meet the limitations of the claimed invention.*" *Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd.*, 821 F.3d 1359, 1367 (Fed. Cir. 2016) (emphasis added). Thus, Dr. Brogioli's testimony about the lack of "a reasonable expectation of success" (Ex. 2023 ¶ 150) is not commensurate in scope with the claims. Furthermore, the above-quoted passage from the '060 patent does not provide any particular details on how to control timing to reduce or prevent collisions, suggesting that this was within the knowledge and skill of a person of ordinary skill in the art, as Petitioner

asserts. *See* Pet. Reply 2–7. Thus, “the [’060] patent itself does not disclose the level of detail that [Patent Owner] would have us require of the prior art.” *See Lockwood v. Am. Airlines, Inc.*, 107 F.3d 1565, 1570 (Fed. Cir. 1997).

At oral argument, Patent Owner’s counsel argued that the ’060 patent “avoids the collisions” by “dramatically remov[ing] the latency on the TSVs,” which “dramatically increases the speed at which the load and reads can be cleared.” Tr. 40:4–11. According to Patent Owner’s counsel, the ’060 patent’s collision solution is “a physical solution to a timing problem, not a timing solution to a physical problem.” Tr. 39:15–40:11 (citing ’160 patent at 8:1–21, which corresponds to 7:63–8:18 in the ’060 patent). The cited disclosure is below.

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 232a is less than the maximum load and the load of the data conduit 232b is less than the maximum load. Further, in many implementations, the combined load of the data conduit 232a and the data conduit 232b is less than the maximum load of a single data conduit. In other words, it is possible to design the data conduit 232a and the data conduit 232b 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.


Ex. 1001, 7:63–8:18. This passage says nothing about avoiding collisions. As discussed above, the '060 patent discloses that “[c]ontrolling the command/address time slots and the data bus time slots enables the control dies 822 to reduce or prevent signal collisions.” Ex. 1001, 21:51–54. Thus, the '060 patent suggests that there is a timing solution to address collisions, contrary to Patent Owner’s counsel’s argument. *See* Tr. 39:15–40:11.

Petitioner asserts that Rajan “provides much more detail than the 060 Patent (EX1001, 21:51–:56) with respect to timing solutions for avoiding collisions.” Pet. Reply 6 (citing Ex. 1015, 8:59–12:13, Figs. 9–15). We agree because, as discussed above, the '060 patent simply mentions controlling timing to reduce or prevent collisions but does not provide any particular details on how to do this. Ex. 1001, 21:51–54. Rajan, however, actually explains how to avoid collisions through the timing of various operations. *See, e.g.*, Ex. 1015, 11:12–12:13, Figs. 14, 15. For example, Rajan discloses that “FIGS. 14 and 15 are a timing diagram 1400 and a timing diagram 1500 illustrating methods of avoiding such collisions.” Ex. 1015, 11:12–14; *see also* Ex. 1052, 200:15–19 (Dr. Brogioli testified as follows: “It looks like Figure 14 says it’s a timing diagram illustrating methods of avoiding such collisions. These look like command collisions of things like activate commands, read and write commands.”).

Patent Owner argues that Rajan does not teach how to avoid collisions for operations that “involved cross-rank read-after-write operations and certainly not for Kim’s memory systems” and that “[d]elaying write commands will just further delay the write data, and does not solve Kim’s data collision problem.” PO Sur-reply 8–9 (citing Ex. 1014 ¶ 42, Fig. 4A).

Rajan’s disclosure is more robust than simply delaying write operations. For example, Rajan discloses that the buffer chip can delay activate operations by one to three clock cycles and that “[t]he actual delay selected may depend on the presence or absence of other DRAM operations that may conflict with the activate operation, and may optionally change from one activate operation to another. In other words, the delay may be dynamic.” Ex. 1015, 11:40–47.

In support of its arguments regarding the fourth configuration, in which four chips belong to four ranks and every two ranks share a TSV, Patent Owner refers to Dr. Brogioli’s modification of Kim’s Figure 2 below.



PO Resp. 33; Ex. 2023 ¶ 120. Kim’s Figure 2 illustrates a memory configuration including shared data input/output (I/O) section 1000 connected via first and second global I/O lines GIO\_Rank0 and GIO\_Rank1 to first and second I/O driving sections 100 and 200, which are connected, respectively, to Rank0 and Rank1. Ex. 1014 ¶¶ 25–26. The modified figure above includes a third I/O driving section connected between GIO\_Rank0 and Rank 2 and a fourth I/O driving section connected between GIO\_Rank1 and Rank 3. Ex. 2023 ¶ 120. Dr. Brogioli testifies that this illustrates

Dr. Wolfe’s proposal that “Kim could be modified by having four ranks sharing two data busses, for instance, in a rank multiplication scenario.”

Ex. 2023 ¶¶ 119–120. At deposition, Dr. Brogioli confirmed that he “assume[d] that Kim’s circuits would operate the same way as in Kim’s original Figure 2.” Ex. 1052, 320:5–11.

Dr. Brogioli’s assumption of no change in operation with the addition of two ranks fails to account for the combined teachings of Kim and Rajan and the knowledge and skill of a person of ordinary skill in the art. As the Federal Circuit instructs,

*KSR* does not require that a combination only unite old elements without changing their respective functions. Instead, *KSR* teaches that “[a] person of ordinary skill is also a person of ordinary creativity, not an automaton.” And it explains that the ordinary artisan recognizes “that familiar items may have obvious uses beyond their primary purposes, and in many cases a person of ordinary skill will be able to fit the teachings of multiple patents together like pieces of a puzzle.”

*ClassCo, Inc. v. Apple, Inc.*, 838 F.3d 1214, 1219 (Fed. Cir. 2016) (citations omitted; alteration in original). The Federal Circuit further cautioned in *ClassCo* that the “rationale of *KSR* does not support [the] theory that a person of ordinary skill can only perform combinations of a puzzle element A with a perfectly fitting puzzle element B.” *Id.* Here, Dr. Brogioli alleges that the modified figure reflects “a rank multiplication scenario” (Ex. 2023 ¶ 119), but when asked about the fact that there are four ranks but only two chip select signals shown in the modified figure, Dr. Brogioli testified that “Dr. Wolfe didn’t articulate anything about additional chip select signals that would relate to what this figure is visualizing.” Ex. 1052, 320:12–24.

Dr. Wolfe, however, testifies that Rajan teaches “emulation techniques, including rank multiplication techniques” (Ex. 1003 ¶ 165), and that Rajan

teaches rank multiplication by generating chip select signals for each chip in the stack. Ex. 1003 ¶¶ 141, 287 (quoting Ex. 1015, 6:34–38). In particular, Rajan discloses that “extra address bits may be decoded by the buffer chip to individually select the DRAM chips, utilizing separate chip select signals (not shown) to each of the DRAM chips in the stack.” Ex. 1015, 6:34–38. Thus, Dr. Brogioli’s modified version of Kim’s Figure 2 does not reflect “a rank multiplication scenario” (Ex. 2023 ¶ 119) in view of the teachings of Rajan.

Patent Owner also argues that “the Petition does not present any evidence that memory systems other than JEDEC-compliant SDRAMs would use chip-selection signals to select chips in the stack or process chip-selection signals on a device level (such that there would be a need to send CS signals to the device).” PO Resp. 49 (citing Ex. 2023 ¶¶ 169–171, 193; Ex. 1015, 6:34–38). Dr. Brogioli identifies NAND and SRAM datasheets that he says do not show “pins for chip-select signals.” Ex. 2023 ¶¶ 169–170 (citing Ex. 2043, 8; Ex. 2044, 5–6). But at deposition, Dr. Brogioli acknowledged that these devices use chip enable signals. Ex. 1052, 236:23–239:3 (discussing Ex. 2043), 239:10–241:6 (discussing Ex. 2044). We agree with Petitioner that the chip enable signals are used “to select the chip that performs the read/write operation.” See Pet. Reply 15; Ex. 2043, 16 (“CE# is used to enable the device.”); Ex. 2044, 5 (pin for “Chip enable input”). The use of chip enable signals is consistent with the district court construction, which Patent Owner urges “should be applied by the Board,” that ““excludes situations in which a chip select signal could *enable* multiple array dies at once.”” See PO Resp. 18–19 (quoting Ex. 2004, 33–34 (emphasis added)).

Patent Owner asserts that “[t]here is also no evidence that the concept of ‘rank’/‘rank multiplication’ . . . would apply to non-DRAM memory systems.” PO Resp. 49 n.8 (quoting Ex. 1022, 413 (“[T]he word *rank* is now used to denote a set of DRAM devices that operate in lockstep to respond to a given command in a memory system.”)). We disagree for the reasons discussed above in section II.D.2.d.1, which explains that Kim does not even mention DRAM and thus is not limited to DRAM and that Rajan states that “any type of memory whatsoever” may be used in its disclosure and lists non-DRAM memories. Ex. 1015, 15:3–9.<sup>12</sup> Thus, we find that the combined teachings of Kim and Rajan teach ranks and rank multiplication with non-DRAM array dies.

*(3) Summary of findings for 1.c, 1.d.1, 1.d.2*

For the reasons discussed above, we find that the combination of Kim and Rajan teaches the subject matter of limitations 1.c, 1.d.1, and 1.d.2 and that a person of ordinary skill in the art would have combined the teachings of Kim and Rajan in the manner asserted with a reasonable expectation of success.

*e) Control die (1.e.1), data conduits (1.e.2, 1.e.3), and control circuit (1.e.4)*


Claim 1 recites a “control die” (1.e.1) comprising “at least a first data conduit between the first die interconnect and a first terminal of the plurality of input/output terminals” (1.e.2) and “at least a second data conduit

---

<sup>12</sup> We further note that Patent Owner’s argument is conclusory and appears only in a footnote. *Cf. CommScope Techs. LLC v. Dali Wireless Inc.*, 10 F.4th 1289, 1296 (Fed. Cir. 2021) (“[A]n argument that is only made in a footnote of an appellant’s brief is forfeited.”).

between the second die interconnect and the first terminal, the first terminal being a data terminal” (1.e.3).

Petitioner argues that Kim’s main chip C0 is a control die and that Kim’s Figure 5 (annotated version below) shows first and second data conduits between the TSVs and the data terminals DQ. Pet. 41–45.



Pet. 41. In the annotated figure above, Petitioner identifies the line from TSV1 to shared data I/O section 1000 as “first data conduit” in pink and identifies the line from TSV2 to shared data I/O section 1000 as “second data conduit” in rose. Pet. 42–46. As shown and as Petitioner asserts, the conduits are between the interconnects (TSVs) and the data terminal (DQ). See Pet. 42–45.

Claim 1 also recites “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” (1.e.4). Petitioner argues that Kim teaches a control circuit in rank selecting unit 1100 that controls the states of the conduits to the TSVs to be active or not active based on signals received

from an external device. Pet. 46–47 (citing Ex. 1014 ¶¶ 32, 35–36, 38, Fig. 3; Ex. 1003 ¶¶ 265–272; Ex. 1019, 6–14, 18, 33; Ex. 1023, 9, Fig. 16; Ex. 1022, 318–20, 332–35).

Patent Owner does not dispute Petitioner’s contentions that the combination of Kim and Rajan teaches limitations 1.e.1–1.e.4. Based on Petitioner’s persuasive contentions and evidence, summarized above, we find that the combination of Kim and Rajan teaches limitations 1.e.1–1.e.4.

*f) Objective evidence of non-obviousness*

We must consider any evidence of objective indicia of non-obviousness before reaching our conclusion on obviousness. *WBIP, LLC v. Kohler Co.*, 829 F.3d 1317, 1328 (Fed. Cir. 2016). This is one of the factual considerations underlying an obviousness determination. *Graham*, 383 U.S. at 17.

Patent Owner argues that the “history of commercial 3DS [(3D-Stacked)] and HBM [(high bandwidth memory)] in particular is objective evidence that the patented design is not obvious.” PO Resp. 1. Patent Owner, however, does not make any arguments to establish the nexus required for consideration of secondary considerations. *See ClassCo*, 838 F.3d at 1220 (“For objective evidence of secondary considerations to be accorded substantial weight, its proponent must establish a nexus between the evidence and the merits of the claimed invention.” (quotations and citation omitted)); *see also* Tr. 42:9–25 (Patent Owner confirming that it did not present an argument for secondary considerations).

Thus, the record does not contain any supportable arguments for objective indicia of non-obviousness.

*g) Conclusion for Claim 1*

We have considered the full trial record, and, for the reasons discussed above and based on Petitioner's contentions and evidence, we conclude that the subject matter of claim 1 would have been obvious to a person of ordinary skill in the art based on the combined teachings of Kim and Rajan.

*3. Dependent Claim 6*

Claim 6 depends from claim 1 and recites "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."

Petitioner argues that it would have been obvious to include chip-select conduits and interconnects to conduct the chip-select signals, relying on Rajan's disclosure of using separate chip select signals for each chip. Pet. 57–62 (citing Ex. 1015, 3:27–30, 6:30–7:67; Ex. 1014 ¶¶ 28–30, 32, 49, Figs. 2, 5; Ex. 1019, 12–13, Fig. 2; Ex. 1022, 319, Ex. 1023, 2–4, 9; Ex. 1016 ¶ 38; Ex. 1003 ¶¶ 315–334).

Patent Owner argues that "the Petition does not present any evidence that memory systems other than JEDEC-compliant SDRAMs would use chip-selection signals to select chips in the stack or process chip-selection signals on a device level (such that there would be a need to send CS signals to the device)." PO Resp. 49 (citing Ex. 2023 ¶¶ 169–171, 193; Ex. 1015, 6:34–38). Patent Owner also asserts that "[t]here is also no evidence that the concept of 'rank'/'rank multiplication' . . . would apply to non-DRAM memory systems." PO Resp. 49 n.8 (quoting Ex. 1022, 413 ("[T]he word *rank* is now used to denote a set of DRAM devices that operate in lockstep

to respond to a given command in a memory system.”)). We address these arguments above in § II.D.2.d.2 and disagree with them for the reasons explained above.

Patent Owner also argues that “TSVs that penetrate through multiple chips were difficult and expensive to form at the time of the invention” and that “one failed TSV would cause the entire die stack to be discarded.” PO Resp. 50 (citing Ex. 2023 ¶¶ 195–196; Ex. 2007; Ex. 2046). Patent Owner includes a parenthetical regarding “a JEDEC presentation that indicated commercially viable speed bin yield would not be viable until 2016.” PO Resp. 50 (citing Ex. 2046). According to Patent Owner, a person of ordinary skill in the art “would not have incurred the cost or the risk to build TSVs that are not needed for the memory chips’ operation.” PO Resp. 50 (citing Ex. 2023 ¶¶ 196–197). The evidence of record shows that a known configuration was to have multiple dies sharing TSVs, as discussed above in § II.D.2.d.2. *See* Ex. 1025 (Foster); *see also* Ex. 2023 ¶ 164 (Dr. Brogioli testifying with respect to Foster about “the dies forming the substack that shares a TSV”). The record evidence shows that using TSVs was within the skill of a person of ordinary skill in the art. *See* Exs. 1014 (Kim), 1025 (Foster). We find that this evidence shows that the use of TSVs was an obvious option for a person of ordinary skill in the art, notwithstanding Patent Owner’s commercial viability arguments. *See Uber Technologies, Inc. v. X One, Inc.*, 957 F.3d 1334, 1340 (Fed. Cir. 2020) (“Because a person of ordinary skill ‘has good reasons to pursue the known options within his or her technical grasp,’ § 103 bars the patentability of such obvious variations.” (quoting *KSR*, 550 U.S. at 421)); *cf. CFMT, Inc. v. Yieldup Int’l Corp.*, 349 F.3d 1333, 1338 (Fed. Cir. 2003) (“Title 35 does not require that a patent

disclosure enable one of ordinary skill in the art to make and use a perfected, commercially viable embodiment absent a claim limitation to that effect.”).

We find Petitioner’s contentions for claim 6 persuasive. Rajan discloses the following:

If the buffer chip is emulating a memory device which has a larger capacity than each of the physical DRAM chips in the stack, the buffer chip may receive from the host system’s memory controller more address bits than are required to address any given one of the DRAM chips. In this instance, the extra address bits may be decoded by the buffer chip to individually select the DRAM chips, utilizing separate chip select signals (not shown) to each of the DRAM chips in the stack.

Ex. 1015, 6:30–38. Thus, in Petitioner’s combination, each chip receives a chip select signal, which supports Petitioner’s position that conduits and interconnects conduct those signals to the chips. *See* Pet. 57–62.

Based on Petitioner’s persuasive contentions and evidence, summarized above, and for the reasons discussed above, we conclude that claim 6 is unpatentable as obvious over the combined teachings of Kim and Rajan.

*4. Claims 2–5, 8–14, 16–19, and 29–34*

Below we address Petitioner’s contentions for claims 2–6, 8–14, 16–19, and 29–34, which we find persuasive. Patent Owner does not raise arguments for these claims in addition to those already addressed above.

Claim 2 depends from claim 1 and recites “wherein the control signals include data path control signals for controlling the first and second data conduits.” Petitioner argues that chip select and read/write command signals would be used per JEDEC to control the direction of data transfer. Pet. 47–49 (citing Ex. 1014 ¶¶ 31–38, Fig. 3; Ex. 1003 ¶¶ 273–280).

Claim 3 depends from claim 1 and recites “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.” As with claim 2, Petitioner cites chip select and read/write signals as controlling the direction of data transfer. Pet. 49–51 (citing Ex. 1014 ¶ 38, Fig. 3; Ex. 1015, 3:27–30, 6:30–7:67, Fig. 18; Ex. 1003 ¶ 287).

Claim 4 depends from claim 3 and recites “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.” Petitioner argues that the combination of Kim and Rajan teaches that command (write and read) and address signals are received at the memory module and that address signals are sent to the dies to identify where to store or retrieve data. Pet. 51–56 (citing Ex. 1014 ¶ 29–30, 32, 38, Ex. 1015, 14:11–18, 14:55–60, Figs. 4, 18; Ex. 1019, 6–14, 18, 33; Ex. 1022, 318–20, 332–35; Ex. 1023, 9, Fig. 6; Ex. 1003 ¶¶ 291–307).

Claim 5 depends from claim 1 and recites “wherein the first die interconnect comprises a first through-silicon via and wherein the second die interconnect comprises a second through-silicon via.” As discussed for claim 1, Petitioner relies on Kim’s TSVs as the die interconnects. Pet. 56–57 (citing Ex. 1014 ¶¶ 46–49, Fig. 5; Ex. 1003 ¶¶ 308–313).

Claim 8 recites,

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.

Petitioner asserts that the combination of Kim and Rajan teaches data path control signals for the reasons given for claims 2 and 3. Pet. 62–63 (citing Ex. 1014, Fig. 3; Ex. 1003 ¶¶ 352–356). Petitioner argues that the combination of Kim and Rajan teaches a control die that is configurable to operate in two modes, one in which the number of ranks seen by the host system equals the actual number of ranks and the other in which there is rank multiplication. Pet. 62–66 (citing Ex. 1014, Fig. 3, Ex. 1015, code (57), 7:4–67, 14:51–62, Fig. 18; Ex. 1003 ¶¶ 352–376).

Claim 9 depends from claim 1 and recites “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.” Petitioner argues that the combination of Kim and Rajan teaches that the control die has conduits for the command/address signals discussed for claim 4 and that those signals include a memory cell address. Pet. 66–67 (citing Ex. 1015, 6:46–49, Fig. 18 (box labeled “Physical Storage Memory Cells”)).

Claim 10 depends from claim 1 and recites “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.” Petitioner argues that, in the combination of Kim and Rajan, conduits would provide power to the chips to operate and that a corresponding ground signal would be required to close the circuit. Pet. 67–69 (citing Ex. 1014 ¶¶ 26, Fig. 5; Ex. 1019, 14; Ex. 1016 ¶ 36, Fig. 2; Ex. 1003 ¶¶ 388–396).

Independent claim 11 is directed to a “memory package” and recites limitations substantially similar to subject matter recited in claims 1 and 6. *See* Pet. 69–70 (mapping similar claim limitations). Petitioner argues that the combination of Kim and Rajan teaches the additional subject matter of “driv[ing] a data signal to an array die selected by at least one of the chip-select signals” by sending data through Kim’s write selecting unit 1110 to a selected array die. Pet. 70–71 (citing Ex. 1014 ¶¶ 35, 40, Fig. 3; Ex. 1003 ¶¶ 429–432).

Claim 12 depends from claim 11 and recites “wherein the chip select conduits pass through the control die.” Petitioner argues that it would have been obvious for chip select signals that are received at the bottom of the memory package to pass through the control die to the array dies above the control die. Pet. 71 (citing Ex. 1015, 6:34–38, Fig. 4; Ex. 1014 ¶¶ 28, 32, Fig. 5; Ex. 1019, 6–11; Ex. 1021, 1–2; Ex. 1022, 374, 476; Ex. 1003 ¶¶ 437–443).

Claim 13 depends from claim 11 and recites “wherein the chip select conduits include drivers to drive the chip select signals to the respective array dies.” Petitioner argues that it would have been obvious to use drivers to produce enough current to transmit the chip select signals. Pet. 73 (citing Ex. 1014 ¶ 32; Ex. 1030, 135–36; Ex. 1038, 68; Ex. 1017, 1:14–20; Ex. 1003 ¶ 447).

Claims 14 and 16–19 recite subject matter that is substantially similar to subject matter recited in claims 1–5 and 8, and Petitioner refers to its contentions for similarly-recited subject matter in claims 1–5 and 8. Pet. 74.

Independent claim 29 is directed to 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.” Claim 29 recites that each DRAM package comprises subject matter that is substantially similar to limitations recited in claims 1, 4, 6, and 11.

Petitioner refers to its contentions for similar subject matter recited in other claims and argues that the combination of Kim and Rajan teaches the additional subject matter of claim 29. Pet. 74–78. In particular, Petitioner argues that the combination teaches a JEDEC memory module that has a register device (Rajan’s register 804). Pet. 74–77 (citing Ex. 1015, 1:28–32, 3:5–7, 4:20–24, 8:52–58, 15:3–12, Fig. 8; Ex. 1014 ¶ 50; Ex. 1022, 43–44, 316–20, 418–19, Fig. 7.2; Ex. 1037, 1:34–37, Fig. 1; Ex. 1003 ¶¶ 591–604, 642–649). As discussed above with respect to claim 1, we find that Rajan teaches using non-DRAM memories in a JEDEC-compliant memory package. We also find that the combination of Kim and Rajan teaches “a plurality of DRAM packages” and “a register device configured to receive command/address signals from the memory control hub and to generate control signals” because Rajan discloses “high capacity DIMM 800 using a plurality of buffered stacks of DRAM circuits 802 and a register device 804,” which “performs the addressing and control of the buffered stacks.” Ex. 1015, 8:52–56; *see* Pet. 76–77 (discussing this disclosure and Fig. 8 of Rajan). We find Petitioner’s contentions persuasive for the remaining subject matter of claim 29 for the reasons discussed for similar limitations in claims 1, 4, 6, and 11.

Claim 30 depends from claim 29 and recites “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.” Petitioner relies on Rajan’s disclosure of using separate chip select signals for each chip to show rank multiplication. Pet. 78 (citing Ex. 1015, 3:27–30, 6:30–7:67, 8:56–58; Ex. 1003 ¶¶ 642–649); *see also* Pet. 8–11 (explaining rank multiplication as known in the art).

Claim 31 depends from claim 29 and recites “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.” Petitioner refers to its contentions for similar subject matter in claims 1 and 2. Pet. 79.

Claim 32 depends from claim 29 and recites “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.” Petitioner refers to its contentions for similar subject matter in claim 16. Pet. 79.

Claim 33 depends from claim 29 and recites “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.” Petitioner argues that Rajan teaches delaying signals by two clock cycles and that a person of ordinary skill in the art would have implemented such a delay in Kim to emulate JEDEC. Pet. 79–80 (citing Ex. 1015, 9:46–10:27, Fig. 11; Ex. 1019, 23–24; Ex. 1003 ¶¶ 669–676).

Claim 34 depends from claim 29 and recites “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.”

Petitioner refers to its contentions for similar subject matter in claims 1 and 3. Pet. 80–81.

As noted above, Patent Owner does not raise additional arguments for these claims. We have reviewed Petitioner’s arguments and evidence, and we find them persuasive. Therefore, having considered the full record developed during the trial, we conclude that claims 2–5, 8–14, 16–19, and 29–34 are unpatentable as obvious over the combined teachings of Kim and Rajan.


*E. Obviousness over Kim, Rajan, and Riho  
(Claims 1–14, 16–19, 29–34)*

In this ground, Petitioner substantively addresses only claim 7 and relies on its contentions based on Kim and Rajan for the remaining claims. Pet. 81–83.

Claim 7 recites,

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.

Petitioner provides the annotated version of Riho's Figure 4 below.



Pet. 82. Riho's Figure 4 shows a stack of memory chips in two groups, D0–D7 and D8–D15, and stacked above logic chip 20 from bottom to top in the following order: D0, D8, D1, D9, D2, D10, D3, D11, D4, D12, D5, D13, D6, D14, D7, and D15. Ex. 1016 ¶¶ 62. In the annotated version of Figure 4 above, Petitioner identifies TSV08 as the "first die interconnect" in light green and TSV15 as the "second die interconnect" in teal. Pet. 82.

Petitioner argues that pair D0/D8 is coupled to TSV08 and that pair D7/D15 is coupled to TSV15 and that such pairing can reduce the load. Pet. 82–83 (citing Ex. 1016 ¶¶ 13, 62, 103, 132, Fig. 4; Ex. 1014 ¶¶ 48, 50, 53–54; Ex. 1003 ¶¶ 172–173, 336–350). Petitioner argues that a person of ordinary skill in the art "would have been motivated with a reasonable expectation of success (when adding chips to Kim) to use Riho's technique described above to further reduce the load (allowing for higher-frequency operations) and to compensate for phase variations, as taught by Riho, without the need to create additional TSVs." Pet. 82 (citing Ex. 1014 ¶¶ 48, 50; Ex. 1016 ¶¶ 119–120; Ex. 1003 ¶¶ 172–173).

Patent Owner raises several arguments in response. PO Resp. 51–54; PO Sur-reply 20–22. Patent Owner argues that analogousness is not a sufficient reason to combine the references. PO Resp. 51 (citing Pet. 81). The cited contention is not Petitioner’s reason to combine (included in discussion above) but is simply an assertion that Rajan and Riho are analogous art, which is a required showing to prove obviousness. *See Netflix, Inc., v. DIVX, LLC*, 80 F.4th 1352, 1358 (Fed. Cir. 2023) (“[F]or an obviousness determination, a reference may only qualify as prior art that a person of ordinary skill in the art would look to if it is ‘analogous to the claimed invention.’” (quoting *In re Bigio*, 381 F.3d 1320, 1325 (Fed. Cir. 2004))). As explained above in § II.D, we find that Kim and Rajan are analogous art to the ’060 patent. Patent Owner does not dispute that Riho is analogous art to the ’060 patent. *See* PO Resp. We agree with Petitioner that Riho is analogous art because it is in the same field of endeavor as the ’060 patent—memory packages and in particular those involving stacked memory. Ex. 1001, 1:20–21 (“The present disclosure relates to memory devices and memory modules.”), claim 1 (“memory package” with “stacked array dies”); Ex. 1016 ¶ 2 (“This invention relates to a semiconductor device incorporating DRAM or other chips and, in particular, relates to a semiconductor device formed by stacking a plurality of chips.”).

Patent Owner also argues that “unlike Kim where different ranks of memories cannot share a same TSVs, in Riho, pairs of dies from two different ranks are connected to the same TSV,” a “material difference” that Petitioner ignores according to Patent Owner. PO Resp. 51–52 (citing Ex. 1014 ¶¶ 42–45, Fig. 5; Ex. 1016 ¶¶ 28–29, 62; Ex. 2023 ¶ 202). We disagree with this argument because the combination of Kim and Rajan

teaches different ranks sharing a TSV, as discussed for claim 1. *See* § II.D.2 above.

According to Patent Owner, “[t]he ‘060 patent teaches that because loads of TSVs could be significant, the number of array dies in a group might need to be adjusted to account for the TSV load, in order to achieve load balancing on different data conduits.” PO Resp. 53 (citing Ex. 1001, 14:3–11, Table 1). Patent Owner argues that “Riho does not do that under Petitioner’s own theory” and, therefore, that “[c]ombining Riho and Kim would also not result in the claimed inventions.” PO Resp. 53 (citing Ex. 2023 ¶¶ 205–209).

Patent Owner’s arguments are not commensurate in scope with claim 7, which does not recite load balancing. Rather, claim 7 recites that the number of array dies in the first and second groups “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” (emphasis added). Thus, the prior art need not teach a balanced load to render obvious the claimed subject matter. Indeed, the ’060 patent discloses that in some embodiments, “the load of each conduit 332 and/or driver 334 may differ” and that “[t]his difference in the load of the conduits 332 and/or drivers 334 may be a design decision.” Ex. 1001, 14:14–17.

Patent Owner contends that Petitioner “simply announces that two chips per group would satisfy 7[a], without explaining why this resulted from ‘consideration of’ the TSV loads.” PO Resp. 53 (citing Pet. 82–83). We disagree. Claim 7 is directed to a memory package that is claimed in part by the process by which numbers of dies are selected, specifically that

the numbers of dies “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.”<sup>13</sup> Claim 7, however, does not recite that load is the only consideration, nor does it recite the starting point from which “to reduce a difference between” loads. Riho discloses the following:

[W]hen the controlled chips are divided into a plurality of groups and the penetrating through substrate vias are commonly used by the controlled chips of the respective groups, it is possible to reduce the load of the interconnections as compared with the case where the penetrating through substrate vias are not commonly used.

Ex. 1016 ¶ 13; *see* Pet. 83 (citing this disclosure). This disclosure makes clear that chip grouping is based on TSV load considerations.

Petitioner argues that Riho’s disclosure of minimizing “skew between data signals DQ and data strobe signals DQS/B” (Ex. 1016 ¶ 50) teaches the subject matter reciting “so as to reduce a difference between a first load on the first data conduit and a second load on the second data conduit.” Pet. 83 (citing Ex. 1016 ¶¶ 50, 53–54; Ex. 1003 ¶¶ 336–343). Patent Owner counters that, “in Riho, the skew between DQ and DQS/B is minimized by making the corresponding TSVs ‘substantially equal in length to each other.’” PO Resp. 52 (quoting Ex. 1016 ¶ 50; citing Ex. 1016 ¶ 53; Ex. 2023

---

<sup>13</sup> Neither party asserts that this is a product-by-process claim. *Kamstrup A/S v. Axioma Metering UAB*, 43 F.4th 1374, 1381 (Fed. Cir. 2022) (“A product by process claim is one in which a product is claimed, at least in part, by the ‘process by which it is made.’” (quoting *In re Thorpe*, 777 F.2d 695, 697 (Fed. Cir. 1985))). Because we conclude that the subject matter of claim 7 would have been obvious, we do not address whether this is a product-by-process claim element that does not impart patentable weight to the claim. *See id.* at 1382.

¶ 203). But as Petitioner points out, Riho discloses using an equal number of dies in each group as part of its structure of using equal length TSVs. *See* Pet. Reply 17. We agree with Petitioner, and we find, that Riho’s die grouping teaches that numbers of dies are selected for groups based on load considerations, even if other considerations are used or are even more prominent. *See* Ex. 1016 ¶¶ 13 (disclosing that dividing chips into a plurality of groups and sharing TSVs makes it “possible to reduce the load of the interconnections as compared with the case where the penetrating through substrate vias are not commonly used”), 103 (“[B]y employing the structure in which the stacked SDRAM chips are divided into the groups (chip selection groups) and the through-silicon vias (TSVs) are each shared by the groups, it is possible to reduce by half the load of each SDRAM chip as compared with the case where the SDRAM chips are not divided into the groups.”).

Patent Owner also argues that Petitioner never asserted that, in the proposed combination, the TSVs of Kim, which are not of equal length, would be made equal length according to the teachings of Riho. PO Sur-reply 20–21. We disagree because Petitioner asserts that a person of ordinary skill in the art “would have been motivated with a reasonable expectation of success (when adding chips to Kim) to use Riho’s technique described above,” referring to Riho’s Figure 4 showing pairs of chips connected to equal length TSVs. Pet. 82 (citing Ex. 1016 ¶¶ 62, 103, Fig. 4; Ex. 1014 ¶¶ 48, 50; Ex. 1003 ¶¶ 172–173). But even if the proposed combination uses Kim’s unequal length TSVs, Riho still provides a teaching of considering load when grouping dies, as discussed above. As the Supreme Court stated in *KSR*, “[w]hat matters is the objective reach of the

claim. If the claim extends to what is obvious, it is invalid under § 103.” *KSR*, 550 U.S. at 419. Here, claim 7 does not recite any particular amount of reduction of load difference nor the starting point from which to determine whether there is a reduction. We find that Riho’s teachings show that a person of ordinary skill in the art would consider TSV load when grouping dies “so as to reduce a difference” between loads on conduits. Whether or not there actually is a reduction in the difference in loads from a particular starting point is not a question we must resolve as it is not required by claim 7.

Having considered the full record developed during the trial, we conclude that claim 7 is unpatentable as obvious over the combined teachings of Kim, Rajan, and Riho for the reasons discussed above and given by Petitioner. We also conclude that claims 1–6, 8–14, 16–19, and 29–34 are unpatentable as obvious for the reasons given above in § II.D.

*F. Obviousness over Kim, Rajan, and Wyman  
(Claims 1–6, 8–34)*

In this ground, Petitioner substantively addresses the additional subject matter pertaining to drivers recited in claims 15 and 20–28 and relies on its contentions based on Kim and Rajan for the remaining claims. Pet. 84–91.

Claim 15 recites,


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.

Petitioner argues that it would have been obvious for the data conduits to have drivers of different strengths for the TSVs of different lengths in Kim based on Wyman’s teachings of using different amounts of drive that are adequate for different path lengths. Pet. 84–89 (citing Ex. 1017, 1:22–24, 1:45–48, 6:15–50, 6:57–60, 7:13–18, Fig. 8; Ex. 1014 ¶¶ 3, 46, Figs. 3, 5; Ex. 1030, 135–38; Ex. 1003 ¶¶ 174–181, 456–464). Petitioner contends that

Kim’s TSV1 and TSV2 are each coupled to circuitry in some, but not all, of the slave chips in the stack, requiring less driver strength than if *all* the slave chips were connected, and Wyman teaches to choose a driver size that will “adequately drive” the signal but is still less than the “total drive,” which would be “overkill” and inefficient.

Pet. 88–89 (citing Ex. 1003 ¶¶ 465–471; Ex. 1017, 1:45–47, 6:57–60); *see* Ex. 1017, 6:57–60 (disclosing that “a designer could assess the amount of drive needed . . . to adequately supply the signal to the devices (902, 904, 906) as needed while using less than the total drive a drive circuit of the prior art might supply”).


Below is Petitioner's annotated version of a portion of Kim's Figure 3.



Pet. 86, 88. The portion of Kim's Figure 3 above shows rank selecting unit 1100 as part of shared data I/O section 1000 including write selecting unit 1110 and read selecting unit 1120. Ex. 1014 ¶¶ 34–35. Petitioner modifies the figure above to include a navy blue triangle representing a “smaller driver” between write selecting unit 1110 GIO\_Rank0 (first die interconnect) and a larger dark green triangle representing a “larger driver” between write selecting unit 1110 and GIO\_Rank1 (second die interconnect). Pet. 85–88.

Patent Owner presents two main arguments in response to Petitioner's reliance on Wyman: (1) that Wyman teaches using a single driver such that a person of ordinary skill in the art would not have used different drivers as recited in the claim; and (2) that Wyman and Kim do not suggest that TSVs of different lengths in Kim would warrant using different amounts of drive. PO Resp. 54–55.

We disagree with Patent Owner’s second argument because it ignores the express disclosure of Wyman. Figure 8 of Wyman is below.



Wyman’s Figure 8 above shows “different through-chip via connections” for stacked chips 700-4 (the “mother chip”) and 700-1, 700-2, and 700-3 (the “daughter chips”). Ex. 1017, 2:21–23, 6:9–14. Figure 8A shows through-chip via connection 802 from mother chip 700-4 to the closest daughter chip labeled 700-3, and Figure 8B shows through-chip via connection 804 from mother chip 700-4 to the farthest daughter chip labeled 700-1. Ex. 1017, 6:15–23. Wyman discloses that “[t]he increased resistance, capacitance and impedance of such a connection (804) might require additional drive than that referred to in connection with FIG. 8A” and that the drive used for connection 802 “would be inadequate.” Ex. 1017, 6:24–28. Based on this disclosure of Wyman, we agree with Petitioner that “Wyman discloses that shorter paths (e.g., 802 . . . ) have less

load and thus require less drive, while longer paths (e.g., 804 . . . ) have more load and thus require a larger drive.” *See* Pet. 84.

We also have considered the following testimony from Dr. Brogioli:


Third, Wyman and Kim do not suggest that the two TSVs with different lengths would result in sufficiently different load to even warrant using different taps in a Wyman-style driver. For instance, load contribution from TSV “may be negligible compared to load contribution from the array dies.” *See* EX1001, 4:36-38. In such situations, same driver strengths can be used and a driver circuitry like that shown in Wyman, Figure 2 will suffice.

Ex. 2023 ¶ 216. The quoted passage from the ’060 patent states in full: “*In some embodiments*, the load contribution from a die interconnect may be negligible compared to the load contribution from the array dies.” Ex. 1001, 4:36–38 (emphasis added). Dr. Brogioli may be correct that, “[i]n such situations, same driver strengths can be used.” Ex. 2023 ¶ 216. Wyman, however, expressly discloses a different situation in which the signal drive for a shorter connection “would be inadequate” for a “longer through-chip via connection” due to “increased resistance, capacitance and impedance of” the longer connection. Ex. 1017, 6:21–30. Wyman also discloses that output drive can “meet the requirements of a particular output load.” Ex. 1017, 7:13–18.

Based on the foregoing, we find that Wyman teaches using different amounts of signal drive for different lengths of TSVs and that, in the proposed combination in which different sets of chips are connected to different length TSVs, each drive amount would be less than the total amount needed to drive a signal to all chips.

We next turn to Patent Owner’s argument that Wyman teaches using a single driver. More particularly, Patent Owner argues that Wyman teaches a

single driver circuit with different taps at which different drive levels can be selected. PO Resp. 54–55 (citing Ex. 1017, 1:36–37, 2:66–3:1, 6:15–50; Ex. 2023 ¶ 211). Patent Owner’s contentions refer to a configuration such as the one below in Figure 5 of Wyman.



Wyman’s Figure 5 above shows current drive portion 500 including driver circuits 502-1, 502-2, 502-3, and 502-4, each of which, via respective enabling leads 504-1, 504-2, 504-3, and 504-4, can be enabled, disabled, or placed in a tri state where minimal power is drawn. Ex. 1017, 3:37–41, 4:66–5:11. Wyman explains that “a designer may now utilize and tap-off at five locations (506-1 [T0], 506-2 [T1], 506-3 [T2], 506-4 [T3], 508) depending on the drive requirements for a specific element or device.” Ex. 1017, 5:11–14.

According to Patent Owner, “Wyman’s driver circuitry, once fabricated, would be regarded as a single driver of a given size regardless of the number of driver components used.” PO Resp. 55 (citing Ex. 2024, 141:18–145:14, 145:24–146:7, 149:24–150:8; Ex. 2023 ¶ 215). Patent Owner argues, therefore, that a person of ordinary skill in the art would not

have used different drivers but, instead, would have used a single driver with selectable drive levels. PO Resp. 55 (citing Ex. 1017, 6:44–50; Ex. 2023 ¶¶ 213–214).

Patent Owner’s proposed combination of using a single driver may be an acceptable option and may be even a better option than having separate drivers. The proposed combination, however, need not be the best option. *See Intel Corp. v. Qualcomm Inc.*, 21 F.4th 784, 800 (Fed. Cir. 2021) (“Our caselaw is clear. It’s not necessary to show that a combination is ‘the *best* option, only that it be a *suitable* option.” (quoting *PAR Pharm., Inc. v. TWI Pharms., Inc.*, 773 F.3d 1186, 1197–98 (Fed. Cir. 2014))). Petitioner proposes using separate drivers in its annotation of Kim’s Figure 3, which is reproduced above in this section. *See* Pet. 85–86 (asserting that a person of ordinary skill in the art would have been motivated to use “a smaller driver (navy blue) for the shorter TSV1 (light green) and a larger driver (dark green) for the longer TSV2 (teal)”). Dr. Wolfe testifies that “you can provide a separate copy of [Wyman’s] circuit 500 for each thing that needs to be driven after you analyze the current requirements.” Ex. 2025, 129:2–7, *cited in* Pet. Reply 18–19. As noted above, Wyman teaches using different amounts of signal drive for different lengths of TSVs. We see no patentable distinction between whether the different amounts of drive come from one set of transistors that can be tapped for different drive strengths or from two different sets of transistors. *See In re Harza*, 274 F.2d 669, 671 (CCPA 1960) (“It is well settled that the mere duplication of parts has no patentable significance unless a new and unexpected result is produced . . .”).

According to Patent Owner, Petitioner’s proposal to use different driver circuits for different TSVs “relies on nothing more than appeal to

common sense, which is improper under *Arendi*.” PO Sur-reply 23–24 (citing *Arendi SARL v. Apple Inc.*, 832 F.3d 1355 (Fed. Cir. 2016)). We disagree that Petitioner’s position relies on nothing more than common sense because Wyman teaches using different amounts of signal drive for different TSVs, as discussed above. That physically separate drivers could be used to produce different amounts of drive in different paths is an unremarkable observation. *See KSR*, 550 U.S. at 421 (“Rigid preventative rules that deny factfinders recourse to common sense . . . are neither necessary under our case law nor consistent with it.”).

Patent Owner and Dr. Brogioli also rely on testimony from Dr. Harold Stone, a technical expert for the Micron entities in district court litigation. *See* Ex. 2024, 141:18–145:14, 145:24–146:7, 149:24–150:8, *cited in* PO Resp. 55 and Ex. 2023 ¶ 215. In this line of questioning, Dr. Stone confirmed that his testimony was based on a hypothetical in which you could “turn on all ten transistors,” and he explained that “[i]f you can turn on all ten, the driver size is related to the load that you can drive. And if you’re going to drive ten transistors, you have to have enough area to dissipate the power they generate.” Ex. 2024, 147:3–11. This testimony, therefore, does not concern the situation in which different and separate drivers drive separate loads, as in Petitioner’s proposed combination discussed above.

Furthermore, although Wyman discloses that different taps can be provided to provide different drive levels, Wyman also discloses an alternative configuration in which taps are not used, as Petitioner correctly points out. *See* Pet. Reply 19 (citing Ex. 1017, 5:34–44). More particularly, Wyman discloses that its “approach could be used with a conventional drive circuit made up of multiple stages by using a via approach.” Ex. 1017,

5:35–44. In this configuration, Wyman discloses forming an “electrically conductive connection to an intermediate point between stages where the current drive is adequate, as needed, and to disable any remaining downstream stage(s) by using a via to break one or more connections and leaving the via unfilled or filling the via with an insulator.” Ex. 1017, 5:35–44. Thus, Wyman expressly discloses that a suitable option involves disabling unneeded transistors and not using taps to provide selectable drive levels. *See Intel*, 21 F.4th at 800. Furthermore, we find that Wyman’s disclosure of providing different drive strengths with different numbers of transistors would result in drivers of different physical size, consistent with the disclosure of the ’060 patent that “[t]he size of the driver may be adjusted by the selection of the transistor size and/or number of transistors included in the driver.” Ex. 1001, 17:43–45; *see* Pet. Reply 19 (discussing this disclosure); *see also* Ex. 2024, 147:7–9 (Dr. Stone’s testimony that “the driver size is related to the load that you can drive”).

Based on the foregoing, we find that Wyman teaches using different amounts of signal drive for different lengths of TSVs and that a person of ordinary skill would have recognized that different, physically separate drivers can be used to provide those different amounts of signal drive. We also find that a person of ordinary skill in the art would have combined Wyman’s teachings with Kim and Rajan “to meet the requirements of a particular output load” and to avoid using more drive power than necessary for each path. Ex. 1017, 7:13–18; *see also* Ex. 1017, 1:22–24 (“In instances where chips are stacked, creating shorter runs, utilizing the full capacity of a driver would be both wasteful and inefficient since far less power is required.”); Ex. 1003 ¶ 179 (“A Skilled Artisan would have understood, and

been motivated, to turn to analogous art, like Wyman teaching about efficiently driving signals between stacked devices, in order to implement the stacked devices of Kim.”). We also find that this combination teaches providing sufficient drive “without significant signal degradation,” as recited in claim 15, because Wyman discloses meeting the requirements of the load. Ex. 1017, 7:13–18; Ex. 1003 ¶¶ 469–471).

Having considered the full record developed during the trial, we conclude that claim 15 is unpatentable as obvious over the combined teachings of Kim, Rajan, and Wyman for the reasons discussed above and given by Petitioner.

Independent claim 20 is directed to a “method for optimizing load in a memory package” and recites structure and method steps that correspond to limitations discussed for claims 1, 4, 6, 11, and 15, and Petitioner refers to its contentions for similarly-recited subject matter in those claims. Pet. 89–90. For claims 21–28, which depend directly or indirectly from claim 20, Petitioner refers to its contentions for similarly-recited subject matter in claims 3–5, 11, 12, 15, and 16. Pet. 90–91. Patent Owner does not raise arguments in addition to those already addressed above.

We have reviewed Petitioner’s arguments and evidence, and we find them persuasive for the reasons discussed above for claims 1–6, 11, 12, 15, and 16. Therefore, having considered the full record developed during the trial, we conclude that claims 20–28 are unpatentable as obvious over the combined teachings of Kim, Rajan, and Wyman.

We also conclude that claims 1–6, 8–14, 16–19, and 29–34 are unpatentable as obvious for the reasons given above in § II.D.

*G. Patent Owner's Procedural Arguments*

Patent Owner makes various procedural arguments with which we disagree as discussed below.

*1. Alleged Improper Incorporation by Reference*

Patent Owner asserts that, “[t]hroughout the Petition, Petitioner tries to evade the word limit by including no substantive analysis in the Petition itself, but attempting to incorporate paragraphs and paragraphs of its expert’s declaration.” PO Resp. 19. As example, Patent Owner notes that “Petitioner devotes two paragraphs in Ground 1” for limitations 1.d.1. and 1.d.2 but cites eleven paragraphs of its expert declaration. PO Resp. 19. Patent Owner asserts that “[t]his is improper, and such arguments should be deemed as absent from the Petition itself.” PO Resp. 19.

We disagree with Patent Owner that there is “no substantive analysis” for certain limitations, and we refer to our findings above as to the Petition’s analysis. Furthermore, there is no prohibition on expert declarations that provide more discussion of claim limitations than a petition. Indeed, our Rules provide that “[e]xpert testimony that does not disclose the underlying facts or data on which the opinion is based is entitled to little or no weight.” 37 C.F.R. § 42.65(a). Thus, an expert declaration should provide the necessary factual bases underlying the declarant’s opinion, and the expert is not limited to only what a petitioner asserts in a petition.

*2. Alleged Improper New Arguments*

Patent Owner argues that having each die in the stack in a separate rank with its own chip select signal is a “new argument” that we should not consider. PO Sur-reply 19. At oral argument, Patent Owner’s counsel stated, “I heard so many times, I couldn’t count, oh, our combination is a

chip select for each die. So, in Riho, that would be 16 chip selects, and in Kim plus Rajan, that would be a chip select for each separate chip. That is nowhere in their petition.” Tr. 36:25–37:2.

We disagree with Patent Owner because the Petition cites Rajan’s disclosure of providing a separate chip select signal for each chip in the stack for claim 6. *See* Pet. 61 (citing Ex. 1015, 6:34–38). Thus, Patent Owner’s contention that this is “nowhere in their petition” is incorrect. Petitioner’s reliance on this disclosure for claim 1 in the Reply was an appropriate response to Patent Owner’s arguments about collisions. *See Apple Inc. v. Andrea Elecs. Corp.*, 949 F.3d 697, 706 (Fed. Cir. 2020) (“[A]ny ambiguity as to whether Apple raised a new argument on reply is eliminated when we consider whether Apple’s reply arguments are responsive to arguments raised in Andrea’s Patent Owner Response.”).

Patent Owner is also incorrect in asserting that having each die in the stack in a separate rank is a new argument. The Petition states the following:

[W]hen implementing “rank multiplication” (discussed above, pp.8-11), it was common to use a “fork-in-the-road” arrangement with two data buses for two ranks of memory devices (shown below left, yellow and orange), with the option of adding two additional ranks of memory devices (*resulting in four ranks*) to the existing data busses (shown below right, two yellow, two orange), meaning two memory dies (e.g., two orange) would share a given data bus.

Pet. 29 (emphasis added; citing Ex. 1003 ¶ 165; Ex. 1026, Figs. 12–13).

This passage refers to annotated versions of Figures 12 and 13 of U.S. Patent Application Publication No. 2006/0277355 A1 (Ex. 1026 (“Ellsberry”)).

Thus, the Petition expressly proposes four separate ranks in the combination.

*See also* Pet. 8 (discussing rank multiplication and citing Ex. 1015, 6:30–

7:67 (Rajan’s disclosure of using separate chip select signals for each chip)). Indeed, Patent Owner’s declarant Dr. Brogioli testifies that, “in Ellsberry, the 4 ranks are in four separate physical ranks.” Ex. 2023 ¶ 111. Petitioner, therefore, put Patent Owner on notice of its contention of using separate ranks. The passage quoted above also expressly states that “it was common to use a ‘fork-in-the-road’ arrangement.” Pet. 29. Thus, we also disagree with Patent Owner’s assertions that “fork in the road” is not in the Petition. *See, e.g.*, Tr. 37:4–7 (“The concept of a fork in the road is the magical elixir that makes clear that Rajan has a shared data interconnect, data terminal, and that the prior art references would share that data terminal. That is nowhere in the petition or the declaration.”); 70:16–17 (“Where does the fork in the road . . . appear in their petition? It does not.”).

Patent Owner also asserts that Petitioner’s arguments about known solutions to collisions are “new arguments that should be disregarded.” PO Sur-reply 2–3; *see also* PO Sur-reply 16–17 (asserting that Petitioner’s “suggest[ion] that Kim’s timing for read and write operations would change in the Kim-Rajan combination . . . is a new theory”). We disagree. Patent Owner asserts in its Response that having dies share a TSV would result in collisions. PO Resp. 32–44. In this case, Petitioner’s arguments that there were known solutions to data collision issues are directly responsive to Patent Owner’s collision arguments. *See Apple*, 949 F.3d at 706; *see also Rembrandt Diagnostics, LP v. Alere, Inc.*, 76 F.4th 1376, 1385 (Fed. Cir. 2023) (finding proper a “reply argument discussing cost and time savings [that] has a nexus to [a] prior argument and is responsive”).

Furthermore, Petitioner’s discussion of the skilled artisan’s knowledge of collision avoidance solutions is proper because it “it is used ‘to document

the knowledge that skilled artisans would bring to bear in reading the prior art identified as producing obviousness.” *Anacor Pharms., Inc v. Iancu*, 889 F.3d 1372, 1380–81 (Fed. Cir. 2018) (quoting *Genzyme Therapeutic Prods. Ltd. P’ship v. Biomarin Pharm. Inc.*, 825 F.3d 1360, 1369 (Fed. Cir. 2016)).

We also find Petitioner’s reliance on *Lockwood* to be appropriate in pointing out that the prior art contains more detail on avoiding collisions than the ’060 patent, which is directly responsive to Patent Owner’s criticisms of the asserted combination. *See* Pet. Reply 12 (citing *Lockwood*, 107 F.3d at 1570); *but see* Tr. 39:3–9 (Patent Owner’s counsel’s assertion that “[t]his Lockwood argument, an argument just like essentially every argument that was made by Samsung in this oral argument was an argument in reply. And the reason why it was put in reply is so their expert can’t have been tested on it, because if he would have been tested upon it, he would have told the truth, which is that there’s a basic physics behind the speed at which a given memory die can clear a load or clear a read.”).

Patent Owner also asserts that Petitioner’s argument that a person of ordinary skill in the art “could use different driver circuits for different TSVs” rather than using “a single driver” is “a new argument.” PO Sur-reply 23. Petitioner’s reliance on different drivers of different sizes is not a new theory because the Petition expressly asserts this and illustrates it, as discussed above in § II.D.2.e. In particular, Petitioner asserts that a person of ordinary skill in the art “would have been motivated to implement Wyman’s teachings in Kim to improve power efficiency by using a smaller driver (navy blue) for the shorter TSV1 (light green) and a larger driver (dark green) for the longer TSV2 (teal), as shown below,” referring to an

annotated version of a portion of Kim’s Figure 3. Pet. 85–86 (illustrating separate drivers). Patent Owner acknowledges as much by arguing that a person of ordinary skill in the art “would not have even used different drivers to drive TSV1 and TSV2 *as asserted*.” PO Resp. 55 (emphasis added; citing Pet. 85). Thus, our Decision relies on the Petition’s express assertions of separate drivers.

Patent Owner further asserts that reliance on non-DRAM devices such as NAND flash and SRAM is improper and “a complete violation of due process” because “the petition makes no discussion of NAND Flash and no discussion of SRAM as being a basis for the combination.” Tr. 55:12–17. We disagree. Before institution, we authorized Petitioner to file a preliminary reply addressing the district court’s construction of “array dies,” and we authorized Patent Owner to file a preliminary sur-reply. Ex. 3001. In the preliminary reply, Petitioner cited disclosure in Rajan that allows for “any type of memory whatsoever” to be used, including NAND flash and SRAM. Paper 9 at 3 (quoting Ex. 1015, 15:3–9). The Petition also cites this disclosure. *See* Pet. 75 (citing Ex. 1015, 15:3–12). The Decision on Institution discussed this disclosure in view of Patent Owner’s argument that array dies must be different from Rajan137’s DRAM circuits. Inst. Dec. 23–24. Thus, Patent Owner was provided with notice and an opportunity to be heard. *See Apple Inc. v. Corephotonics, Ltd.*, 81 F.4th 1353, 1361 (Fed. Cir. 2023) (stating that the Board’s “decisions must be reached only after the parties have been provided fair notice and an opportunity to be heard”).

#### *H. Remaining Grounds*

Because we determine that all challenged claims are unpatentable as discussed above, we need not separately assess the remaining grounds of

unpatentability based on the combinations of Riho and Rajan and Riho, Rajan, and Riho<sup>2</sup>. 35 U.S.C. § 318(a) (“If an inter partes review is instituted and not dismissed under this chapter, the Patent Trial and Appeal Board shall issue a final written decision with respect to the patentability of any patent claim challenged by the petitioner and any new claim added under section 316(d).”); *Bos. Sci. Scimed, Inc. v. Cook Grp. Inc.*, 809 F. App’x 984, 990 (Fed. Cir. 2020) (nonprecedential) (“We agree that the Board need not address issues that are not necessary to the resolution of the proceeding.”).

### III. PETITIONER’S MOTION TO EXCLUDE

Petitioner filed a Motion to Exclude Exhibits 2024, 2026, 2027, and 2029–2031. Paper 34. Because we do not rely on this evidence in a manner adverse to Petitioner, we dismiss the Motion to Exclude as moot.

### IV. CONCLUSION<sup>14</sup>

For the reasons discussed above, we determine that Petitioner has proven, by a preponderance of the evidence, that claims 1–34 of the ’060 patent are unpatentable, as summarized in the following table:

---

<sup>14</sup> Should Patent Owner wish to pursue amendment of the challenged claims in a reissue or reexamination proceeding subsequent to the issuance of this decision, we draw Patent Owner’s attention to the April 2019 *Notice Regarding Options for Amendments by Patent Owner Through Reissue or Reexamination During a Pending AIA Trial Proceeding*. See 84 Fed. Reg. 16,654 (Apr. 22, 2019). If Patent Owner chooses to file a reissue application or a request for reexamination of the challenged patent, we remind Patent Owner of its continuing obligation to notify the Board of any such related matters in updated mandatory notices. See 37 C.F.R. § 42.8(a)(3), (b)(2).

| <b>Claim(s)</b>                     | <b>35<br/>U.S.C. §</b> | <b>Reference(s)/Basis</b> | <b>Claims<br/>Shown<br/>Unpatentable</b> | <b>Claims<br/>Not Shown<br/>Unpatentable</b> |
|-------------------------------------|------------------------|---------------------------|------------------------------------------|----------------------------------------------|
| 1–6, 8–<br>14, 16–<br>19, 29–<br>34 | 103(a)                 | Kim, Rajan                | 1–6, 8–14,<br>16–19, 29–34               |                                              |
| 1–14,<br>16–19,<br>29–34            | 103(a)                 | Kim, Rajan, Riho          | 1–14, 16–19,<br>29–34                    |                                              |
| 1–6, 8–<br>34                       | 103(a)                 | Kim, Rajan,<br>Wyman      | 1–6, 8–34                                |                                              |
| 1–14,<br>16–19,<br>29–34            | 103(a)                 | Riho, Rajan <sup>15</sup> |                                          |                                              |
| 1–34                                | 103(a)                 | Riho, Rajan, Riho2        |                                          |                                              |
| <b>Overall<br/>Outcome</b>          |                        |                           | 1–34                                     |                                              |

## V. ORDER

Accordingly, it is

ORDERED that claims 1–34 of the '060 patent have been shown to be unpatentable;

FURTHER ORDERED that Petitioner's Motion to Exclude (Paper 34) is *dismissed*; and

FURTHER ORDERED that, because this is a Final Written Decision, parties to the proceeding seeking judicial review of the Decision must comply with the notice and service requirements of 37 C.F.R. § 90.2.

---

<sup>15</sup> As explained above, because we determine that the challenged claims are unpatentable based on the combinations with Kim and Rajan, we decline to address the remaining grounds.

For PETITIONER:

Eliot D. Williams  
Theodore W. Chandler  
Ferenc Pazmandi  
Michael E. Knierim  
Brianna L. Potter  
BAKER BOTTS L.L.P.  
eliot.williams@bakerbotts.com  
ted.chandler@bakerbotts.com  
ferenc.pazmandi@bakerbotts.com  
michael.knierim@bakerbotts.com  
brianna.potter@bakerbotts.com

For PATENT OWNER:

Hong Annita Zhong  
IRELL & MANELLA  
hzhong@irell.com

UNITED STATES PATENT AND TRADEMARK OFFICE

---

BEFORE THE OFFICE OF THE DEPUTY UNDER SECRETARY OF  
COMMERCE FOR INTELLECTUAL PROPERTY AND DEPUTY  
DIRECTOR OF THE UNITED STATES PATENT AND TRADEMARK  
OFFICE

---

SAMSUNG ELECTRONICS CO., LTD., MICRON TECHNOLOGY,  
INC., MICRON SEMICONDUCTOR PRODUCTS, INC., and MICRON  
TECHNOLOGY TEXAS LLC,<sup>1</sup>  
Petitioner,

v.

NETLIST, INC.,  
Patent Owner.

---

IPR2022-01427 (Patent 9,318,160 B2)  
IPR2022-01428 (Patent 8,787,060 B2)

---

Before DERRICK BRENT,<sup>2</sup> *Deputy Under Secretary of Commerce for  
Intellectual Property and Deputy Director of the United States Patent and  
Trademark Office.*

---

<sup>1</sup> Micron Technology, Inc., Micron Semiconductor Products, Inc., and Micron Technology Texas LLC filed petitions in IPR2023-00882 and IPR2023-00883 and were joined as parties to these proceedings. See IPR2022-01427, Paper 25; IPR2022-01428, Paper 25.

<sup>2</sup> Katherine K. Vidal, Under Secretary of Commerce for Intellectual Property and Director of the United States Patent and Trademark Office, took no part in this decision. The Director's authority is delegated to Derrick Brent, Deputy Under Secretary of Commerce for Intellectual Property and Deputy Director of the United States Patent and Trademark Office, by operation of

## ORDER

The Office received a request for Director Review of the Final Written Decision in each of the above-captioned cases. *See* IPR2022-01427, Paper 50; Ex. 3100; IPR2022-01428, Paper 50; Ex. 3100. Director Vidal is recused from this matter, and the request was referred to me.

Upon consideration of the request, it is:

ORDERED that the requests for Director Review are denied; and

FURTHER ORDERED that the Patent Trial and Appeal Board's Final Written Decision in these cases is the final decision of the agency.

---

the Director's Memorandum § II. c. *See* Director's Memorandum, Procedures for Recusal to Avoid Conflicts of Interest and Delegations of Authority (Apr. 20, 2022) (Recusal Procedure Memo), available at [www.uspto.gov/sites/default/files/documents/Director-Memorandum-on-Recusal-Procedures.pdf](http://www.uspto.gov/sites/default/files/documents/Director-Memorandum-on-Recusal-Procedures.pdf).

IPR2022-01427 (Patent 9,318,160 B2)

IPR2022-01428 (Patent 8,787,060 B2)

For PETITIONER:

Eliot D. Williams

Theodore W. Chandler

Ferenc Pazmandi

Michael E. Knierim

Brianna L. Potter

BAKER BOTTS LLP

Eliot.williams@bakerbotts.com

Ted.chandler@bakerbott.com

Ferenc.pazmandi@bakerbotts.com

michael.knierim@bakerbotts.com

Brianna.potter@bakerbotts.com

Matthew Hopkins

WINSTON & STRAWN LLP

mhopkins@winston.com

For PATENT OWNER:

Hong Annita Zhong

IRELL & MANELLA LLP

hzhong@irell.com



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**

(56) **References Cited**

U.S. PATENT DOCUMENTS

(71) Applicant: **Netlist, Inc.**, Irvine, CA (US)

5,345,412 A 9/1994 Shiratsuchi  
5,537,584 A 7/1996 Miyai et al.

(72) Inventor: **Hyun Lee**, Ladera Ranch, CA (US)

(Continued)

(73) Assignee: **Netlist, Inc.**, Irvine, CA (US)

FOREIGN PATENT DOCUMENTS

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

EP 1816570 A2 8/2007  
JP 09237492 9/1997

(Continued)

This patent is subject to a terminal disclaimer.

OTHER PUBLICATIONS

(21) Appl. No.: **14/337,168**

International Search Report and Written Opinion, International Application No. PCT/US14/48517, filed Jul. 28, 2014, dated Oct. 27, 2014.

(22) Filed: **Jul. 21, 2014**

(Continued)

(65) **Prior Publication Data**

US 2015/0070959 A1 Mar. 12, 2015

*Primary Examiner* — Huan Hoang

(74) *Attorney, Agent, or Firm* — Jamie J. Zheng, Esq.

**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.

(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 Jaganathan," 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 Jaganathan," 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 Jaganathan," 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.*, 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




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6A




FIG. 6B




FIG. 7




FIG. 8

1

# 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. 1A and 1B 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. 1A 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. 1A 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. 1A, 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. 1B 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

2

FIG. 1B, 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

3

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. 1A and 1B 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. 6A and 6B 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.

4

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

5

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 **210a-210d**). 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 **210a-210d** 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 **220a**, **220b**). 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 **220a** may be in electrical communication with a data port from array die **210a** and a

6

data port from array die **210b** (as illustrated by the darkened circles in FIG. 2) and not in electrical communication with any data ports from array die **210c** or any data ports from array die **210d**. The data ports of array dies **210a** and **210b** in electrical communication with the die interconnect **220a** 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 **210a** and **210b**. These other die interconnects can be electrically isolated from the corresponding data ports of array dies **210c** and **210d**.

However, continuing this example, the die interconnect **220b** may be in electrical communication with a data port from array die **210c** and a data port from array die **210d** (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 **210a** and array die **210b**. 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 **210c** and **210d**. Despite not being in electrical communication with any data ports from array die **210a** and **210b**, in some implementations, the die interconnect **220b** may pass through the array dies **210a** and **210b** (as illustrated by the unfilled circles) e.g., through through-holes or vias of array dies **210a** and **210b**. 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 **232a** and **232b**. 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 **232a** may be configured to transmit a data signal to the die interconnect **220a** without transmitting the data signal to the die interconnect **220b**. Conversely, the data conduit **232b** may be configured to transmit the data signal to the die interconnect **220b** without transmitting the data signal to the die interconnect **220a** (e.g., data conduit **232b** is electrically isolated from die interconnect **220a** and data conduit **232a** is electrically isolated from die interconnect **220b**).

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 **232a** and the data conduit **232b** 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 **232a** may include loads of the die interconnect **220a**, the array die **210a**, and the array die **210b**. Similarly, the load of the data conduit **232b** may include loads of the die interconnect **220b**, the array die **210c**, and the array die **210d**.

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 **232a** and the load of the data conduit **232b** 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 **232a**, **232b**. In some implementations, the difference between the load of the data conduit **232a** and the load of the data conduit **232b** 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 **232a** and the load of the data conduit **232b** 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 **232a** in this example is 2.5 and the load on the data conduit **232b** in this example is 3. Thus, in this example, the difference between the load of the data conduit **232a** and the load of the data conduit **232b** is at the threshold load difference value of 0.5. However, an alternative configuration that places the die interconnect **220a** in electrical communication with only the array die **210a**, and the die interconnect **220b** in electrical communication with the array dies **210b-210d** 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 **232a** would be 1.25 and the load on the data conduit **232b** would be 4. Thus, in the alternative configuration, the difference between the load of the data conduit **232a** and the load of the data conduit **232b** 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 **232a** is less than the maximum load and the load of the data conduit **232b** is less than the maximum load. Further, in many implementations, the combined load of the data conduit **232a** and the data conduit **232b** is less than the maximum load of a single data conduit. In other words, it is possible to design the data conduit **232a** and the data conduit **232b** 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 **220a** may be connected with data ports of the array dies **210a** and **210b** 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 **220b** may pass through array dies **210a** and **210b** using TSVs that do not enable electrical communication between the die interconnect **220b** and data ports of the array dies **210a** and **210b** (illustrated by the unfilled circles). In this way, the array dies **210a**, **210b** are not responsive to the data signal being transmitted by the die interconnect **220b**. However, the die interconnect **220b** may be connected with at least one data port of each of the array dies **210c** and **210d** 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 control 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 **250a-250d** as shown in FIG. 2). Further, the control die **230** may include corresponding die interconnects **252** (e.g., die interconnects **252a-252d**) 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 **252a** may be in electrical communication with the array die **210a** and the die interconnect **252b** may be in electrical communication with the array die **210b**. 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 **310a-310h**). 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 **332a-332f**). 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 **334a-334f**). Each of the drivers **334** may be configured to drive a signal along a corresponding die interconnect **320**. For instance, the driver **334a** of the conduit **332a** may be configured to drive a signal along the die interconnect **320a** to one or more of the array dies **310a** and **310b**. As a second example, the driver **334b** of the conduit **332b** may be configured to drive a signal along the die interconnect **320b** to one or more of the array dies **310c** and **310d**. Although the die interconnect **320b** may pass through the array dies **310a** and **310b**, because the die interconnect **320b**, in the example illustrated in FIG. 3, is not configured to be in electrical communication with the array dies **310a** and **310b**, the driver **334b** does not drive the signal to the array dies **310a** and **310b**.

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

11

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 332a may include the loads of the die interconnect 320a, the array die 310a, and the array die 310b. Similarly, the loads of the conduit 332b may include the loads of the die interconnect 320b, the array die 310c, and the array die 310d. The loads of both conduits 332a and 332b 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 332c, which may include the loads of the die interconnect 320c and the array die 310e, includes a load of one array die 310e because the corresponding die interconnect 320c 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 332a and the load of the conduit 332b is less than the maximum load. As a second example, the difference between the load of the conduit 332f and anyone of the conduits 332a-332e 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 332a, the load of the conduit 332b, and the load of the conduit 332c 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 310a may be substantially equal to the load of the array die 310h. Thus, the load contribution from the array dies 310 for a specific conduit 332 may

12

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 332a would be 2 L because the die interconnect 320a corresponding to the conduit 332a is in electrical communication with two array dies 310, array dies 310a and 310b. In alternative embodiments, the load of each of the array dies 310a may differ. For example, the load of array die 310a may be L and the load of the array die 310b 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 320a may extend from the top of the array die 310a to the bottom of the array die 310b. 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 320a may extend from the center of array die 310a to the center of array die 310b. As a second example, one segment of the die interconnect 320a may extend from the top of the control die 330 to the top of the array die 310a, and therefore may include a portion of the die interconnect 320a extending from the bottom of the array die 310a to the top of the array die 310a. 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 320a may be equivalent to one quarter of the load of an array die 310. Thus, for example, the load of the conduit 332a 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 320a and a load contribution of 0.25 L per segment (two in this case) of the die interconnect 320a. As a second example, the load of the conduit 332f may be 3 L assuming the same load values as the previous example and a load contribution from one array die 310h and eight die interconnect 320f 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 332b and 332f.

13

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 332b and 332f 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 334f may be larger than the driver 334e. Alternatively, each driver 334 may be substantially equal and may be configured based on the drivers 334 with the largest load, which are drivers 334b and 334f 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).

$$CL=AD+S/M$$

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 332a, 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 320a; 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 320a, 0.25 L. Therefore, as can be seen from Table 1, the capacitive load of the driver 332a 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

14

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 334f to 334a). 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 334f to 334a). 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 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 10L, which is the load of a conduit in electrical communica-

15

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 232a and 232b. 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 404a and 404b. In turn, these drivers 404a and 404b can drive the signal to conduits 406a and 406b 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 404a and 404b, and 408a and 408b 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 404a and 404b 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 406a is in electrical communication with array dies one and two (not shown), and that the conduit 406b 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 404a may be configured to drive a signal along a die interconnect in electrical communication with the conduit 406a to array die one and two. In this example, the driver 404b 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 404a and 404b 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 406a is in electrical communication with two of the array dies, and if the die interconnect in electrical communication with the conduit 406b is in electrical communication with the other two array dies, then the drivers 404a and 404b may drive a signal to all four of the array dies of this example.

16

Similar to the input/output port 402, the conduits 406a and 406b can be configured to receive a signal from the die interconnects that are in electrical communication with the conduits 406a and 406b. In turn, the signal received at one of the conduits 406a and 406b can be provided to drivers 408a and 408b respectively. The drivers 408a and 408b can each be configured to drive a signal received from a respective data conduit 406a and 406b 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 406a and 406b and the drivers 408a and 408b. 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 404a and 404b 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 404a and 408a being in electrical communication with the same conduit 406a, each of the drivers 404a and 408a can be in electrical communication with a separate conduit and consequently a separate die interconnect. In such an example, the drivers 404a and 408a 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 404a in electrical communication with an input port, and the driver 408a 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

17

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

18

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. **6A**). 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. **6B**, 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 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 **310b**, then the driver **334a** will drive a signal along the die interconnect **320a** to the array die **310b**, and the remaining drivers (e.g., drivers **334b-334f**) 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.

21

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

22

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 mode 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

23

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 corre-

24

sponding 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;

25

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.

26

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.

\* \* \* \* \*



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, 39<sup>th</sup> Annual IEEE/ACM International Symposium on, Dec. 2006, 469-479, Orlando, FL.

Daneshtalab et al., "CMT—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, 40<sup>th</sup> 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 Co, 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 35<sup>th</sup> 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




FIG. 1A




FIG. 1B





FIG. 2






FIG. 4




FIG. 5




FIG. 6A




FIG. 6B




FIG. 7




FIG. 8

1

# 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. 1A and 1B 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. 1A 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. 1A 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. 1A, 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. 1B 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. 1B, each array die 160 includes a chip select port 174. However, because the control die 170 does not also serve as an

2

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

3

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. 1A and 1B 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. 6A and 6B 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.

4

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>.

5

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 **210a-210d**). 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 **210a-210d** 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 **220a, 220b**). 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 **220a** may be in electrical communication with a data port from array die **210a** and a data port from array die **210b** (as illustrated by the darkened circles in FIG. 2) and not in electrical communication with any data ports from array die **210c** or any data ports from array

6

die **210d**. The data ports of array dies **210a** and **210b** in electrical communication with the die interconnect **220a** can be corresponding to the same data bit (e.g., D0). Other die interconnects (not shown) can be in electrical communication with other data ports corresponding to other data bits (e.g., D1, D2, . . . ) of array dies **210a** and **210b**. These other die interconnects can be electrically isolated from the corresponding data ports of array dies **210c** and **210d**.

However, continuing this example, the die interconnect **220b** may be in electrical communication with a data port from array die **210c** and a data port from array die **210d** (as illustrated by the darkened circles in FIG. 2) (e.g., corresponding to the same data bit, e.g., D0) without being in electrical communication with any data ports from array die **210a** and array die **210b**. Other die interconnects (not shown) can be in electrical communication with other data ports corresponding to other data bits (e.g., D1, D2, . . . ) of array dies **210c** and **210d**. Despite not being in electrical communication with any data ports from array die **210a** and **210b**, in some implementations, the die interconnect **220b** may pass through the array dies **210a** and **210b** (as illustrated by the unfilled circles) e.g., through through-holes or vias of array dies **210a** and **210b**. 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 **232a** and **232b**. 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 **232a** may be configured to transmit a data signal to the die interconnect **220a** without transmitting the data signal to the die interconnect **220b**. Conversely, the data conduit **232b** may be configured to transmit the data signal to the die interconnect **220b** without transmitting the data signal to the die interconnect **220a** (e.g., data conduit **232b** is electrically isolated from die interconnect **220a** and data conduit **232a** is electrically isolated from die interconnect **220b**).

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 **232a** and the data conduit **232b** 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 **232a** may include loads of the die interconnect **220a**, the array die **210a**, and the array die **210b**. Similarly, the load of the data conduit **232b** may include loads of the die interconnect **220b**, the array die **210c**, and the array die **210d**.

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 **232a** and the load of the data conduit **232b** 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 **232a**, **232b**. In some implementations, the difference between the load of the data conduit **232a** and the load of the data conduit **232b** 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 **232a** and the load of the data conduit **232b** 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 **232a** in this example is 2.5 and the load on the data conduit **232b** in this example is 3. Thus, in this example, the difference between the load of the data conduit **232a** and the load of the data conduit **232b** is at the threshold load difference value of 0.5. However, an alternative configuration that places the die interconnect **220a** in electrical communication with only the array die **210a**, and the die interconnect **220b** in electrical communication with the array dies **210b-210d** 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 **232a** would be 1.25 and the load on the data conduit **232b** would be 4. Thus, in the alternative configuration, the difference between the load of the data conduit **232a** and the load of the data conduit **232b** 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 **232a** is less than the maximum load and the load of the data conduit **232b** is less than the maximum load. Further, in many implementations,

the combined load of the data conduit **232a** and the data conduit **232b** is less than the maximum load of a single data conduit. In other words, it is possible to design the data conduit **232a** and the data conduit **232b** 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 **220a** may be connected with data ports of the array dies **210a** and **210b** 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 **220b** may pass through array dies **210a** and **210b** using TSVs that do not enable electrical communication between the die interconnect **220b** and data ports of the array dies **210a** and **210b** (illustrated by the unfilled circles). In this way, the array dies **210a**, **210b** are not responsive to the data signal beings transmitted by the die interconnect **220b**. However, the die interconnect **220b** may be connected with at least one data port of each of the array dies **210c** and **210d** 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 **250a-250d** as shown in FIG. 2). Further, the control die **230** may include corresponding die interconnects **252** (e.g., die interconnects **252a-252d**) 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 **252a** may be in electrical communication with the array die **210a** and the die interconnect **252b** may be in electrical communication with the array die **210b**. 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 **310a-310h**). 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 **332a-332f**). 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 **334a-334f**). Each of the drivers **334** may be configured to drive a signal along a corresponding die interconnect **320**. For instance, the driver **334a** of the conduit **332a** may be configured to drive a signal along the die interconnect **320a** to one or more of the array dies **310a** and **310b**. As a second example, the driver **334b** of the conduit **332b** may be configured to drive a signal along the die interconnect **320b** to one or more of the array dies **310c** and **310d**. Although the die interconnect **320b** may pass through the array dies **310a** and **310b**, because the die interconnect **320b**, in the example illustrated in FIG. 3, is not configured to be in electrical communication with the array dies **310a** and **310b**, the driver **334b** does not drive the signal to the array dies **310a** and **310b**.

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.

11

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 **332a** may include the loads of the die interconnect **320a**, the array die **310a**, and the array die **310b**. Similarly, the loads of the conduit **332b** may include the loads of the die interconnect **320b**, the array die **310c**, and the array die **310d**. The loads of both conduits **332a** and **332b** 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 **332c**, which may include the loads of the die interconnect **320c** and the array die **310e**, includes a load of one array die **310e** because the corresponding die interconnect **320c** 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 **332a** and the load of the conduit **332b** is less than the maximum load. As a second example, the difference between the load of the conduit **334f** and any one of the conduits **334a-334e** 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 **332a**, the load of the conduit **332b**, and the load of the conduit **332c** 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

12

example, the load of the array die **310a** may be substantially equal to the load of the array die **310h**. 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 **332a** would be  $2L$  because the die interconnect **320a** corresponding to the conduit **332a** is in electrical communication with two array dies **310**, array dies **310a** and **310b**. In alternative embodiments, the load of each of the array dies **310a** may differ. For example, the load of array die **310a** may be  $L$  and the load of the array die **310b** 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 **320a** may extend from the top of the array die **310a** to the bottom of the array die **310b**. 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 **320a** may extend from the center of array die **310a** to the center of array die **310b**. As a second example, one segment of the die interconnect **320a** may extend from the top of the control die **330** to the top of the array die **310a**, and therefore may include a portion of the die interconnect **320a** extending from the bottom of the array die **310a** to the top of the array die **310a**. 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 **320a** may be equivalent to one quarter of the load of an array die **310**. Thus, for example, the load of the conduit **332a** 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 **320a** and a load contribution of  $0.25L$  per segment (two in this case) of the die interconnect **320a**. As a second example, the load of the conduit **332f** may be  $3L$  assuming the same load values as the previous example and a load contribution from one array die **310h** and eight die interconnect **320f** segments. Table 1 specifies the capacitive load values for each conduit **332** assuming, as with the previous two examples, that the load of

13

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 **332b** and **332f**.

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 **332b** and **332f** 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 **334f** may be larger than the driver **334e**. Alternatively, each driver **334** may be substantially equal and may be configured based on the drivers **334** with the largest load, which are drivers **334b** and **334f** 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 **332a**, 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 **320a**; and M=4, for the ratio of the load of an array die **310**, L, to the load of a segment of the die

14

interconnect **320a**, 0.25L. Therefore, as can be seen from Table 1, the capacitive load of the driver **332a** 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 **334f** to **334a**). 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 **334f** to **334a**). 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 **232a** and **232b**. 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 **404a** and **404b**. In turn, these drivers **404a** and **404b** can drive the signal to conduits **406a** and **406b** 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 **404a** and **404b**, and **408a** and **408b** 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 **404a** and **404b** 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 **406a** is in electrical communication with array dies one and two (not shown), and that the conduit **406b** 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 **404a** may be configured to drive a signal along a die interconnect in electrical communication with the conduit **406a** to array die one and two. In this example, the driver **404b** 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 **404a** and **404b** 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 **406a** is in electrical communication with two of the array dies, and if the die interconnect in electrical communication with the conduit **406b** is in electrical communication with the other two array dies, then the drivers **404a** and **404b** may drive a signal to all four of the array dies of this example.

Similar to the input/output port **402**, the conduits **406a** and **406b** can be configured to receive a signal from the die interconnects that are in electrical communication with the conduits **406a** and **406b**. In turn, the signal received at one of the conduits **406a** and **406b** can be provided to drivers **408a** and **408b** respectively. The drivers **408a** and **408b** can each be configured to drive a signal received from a respective data conduit **406a** and **406b** 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 **406a** and **406b** and the drivers **408a** and **408b**. 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 **404a** and **404b** 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 **404a** and **408a** being in electrical communication with the same conduit **406a**, each of the drivers **404a** and **408a** can be in electrical communication with a separate conduit and consequently a separate die interconnect. In such an example, the drivers **404a** and **408a** 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 **404a** in electrical communication with an input port, and the driver **408a** 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

17

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

18

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. 6A). These shortcomings include:

- The DIMM density limitation due to the fixed number of chip-select signals received by the DIMM from the system memory controller.
- 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.
- 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. 6B, 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. 6B. 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 **310b**, then the driver **334a** will drive a signal along the die interconnect **320a** to the array die **310b**, and the remaining drivers (e.g., drivers **334b-334f**) 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

21

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

22

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,"

23

“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

24

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.

25

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

26

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 via.

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

27

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.

28

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.

\* \* \* \* \*

**UNITED STATES COURT OF APPEALS  
FOR THE FEDERAL CIRCUIT****CERTIFICATE OF COMPLIANCE WITH TYPE-VOLUME LIMITATIONS****Case Number:** 24-2240, 24-2241**Short Case Caption:** Netlist, Inc. v. Samsung Electronics Co., Ltd.**Instructions:** When computing a word, line, or page count, you may exclude any items listed as exempted under Fed. R. App. P. 5(c), Fed. R. App. P. 21(d), Fed. R. App. P. 27(d)(2), Fed. R. App. P. 32(f), or Fed. Cir. R. 32(b)(2).

The foregoing filing complies with the relevant type-volume limitation of the Federal Rules of Appellate Procedure and Federal Circuit Rules because it meets one of the following:

- ☒ the filing has been prepared using a proportionally-spaced typeface and includes 13992 words.
- ☐ the filing has been prepared using a monospaced typeface and includes \_\_\_\_\_ lines of text.
- ☐ the filing contains \_\_\_\_\_ pages / \_\_\_\_\_ words / \_\_\_\_\_ lines of text, which does not exceed the maximum authorized by this court's order (ECF No. \_\_\_\_\_).

Date: 02/04/2025Signature: /s/ Jeffrey A. LamkenName: Jeffrey A. Lamken