

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

SAMSUNG ELECTRONICS CO., LTD.,	)	
SAMSUNG SEMICONDUCTOR, INC. and	)	
SAMSUNG ELECTRONICS AMERICA,	)	
INC.,	)	
	)	
Plaintiffs,	)	C.A. No. 25-626 (JLH)
	)	
v.	)	
	)	
NETLIST, INC.,	)	
	)	REDACTED - PUBLIC VERSION
Defendant.	)	

**PLAINTIFF’S MOTION FOR LEAVE
TO FILE AN AMENDED COMPLAINT**

Plaintiffs Samsung Electronics Co., Ltd., Samsung Semiconductor, Inc., and Samsung Electronics America, Inc. (“Samsung”) respectfully move the Court for leave to amend its Complaint (D.I. 1), pursuant to the Fed. R. Civ. P. 15(a), to add declaratory judgment claims regarding U.S. Patent Nos. 10,025,731 and 12,373,366 against Defendant Netlist, Inc. (“Netlist”) and to join Avnet, Inc. (“Avnet”) as a co-plaintiff.

Pursuant to D. Del. LR 15.1, a copy of the proposed Amended Complaint is attached as Exhibit A and a redlined version of the Amended Complaint, highlighting the changes from the Complaint, is attached as Exhibit B.

Pursuant to D. Del. LR 7.1.1, counsel for Samsung raised the subject of this motion with counsel for Netlist, and Netlist has not consented to Samsung’s motion.

MORRIS, NICHOLS, ARSHT & TUNNELL LLP

/s/ Rodger D. Smith

OF COUNSEL:

Brian Nester
Peter Swanson
COVINGTON & BURLING LLP
One CityCenter
850 Tenth Street, NW
Washington, DC 20001-4956
(202) 662-6000

Alice J. Ahn
COVINGTON & BURLING LLP
Salesforce Tower
415 Mission Street, Suite 5400
San Francisco, CA 94105-2533
(415) 591-6000

July 29, 2025

Rodger D. Smith II (#3778)
Anthony D. Raucci (#5948)
Ben Yenerall (#7132)
1201 North Market Street
P.O. Box 1347
Wilmington, DE 19899
(302) 658-9200
rsmith@morrisnichols.com
araucci@morrisnichols.com
byenerall@morrisnichols.com

Attorneys for Plaintiffs

CERTIFICATE OF SERVICE

I hereby certify that on July 29, 2025, I caused the foregoing to be electronically filed with the Clerk of the Court using CM/ECF, which will send notification of such filing to all registered participants.

I further certify that I caused copies of the foregoing document to be served on July 29, 2025, upon the following in the manner indicated:

Karen Elizabeth Keller
Emily DiBenedetto, Esquire
SHAW KELLER LLP
1105 N Market St, 12th Floor
Wilmington, DE 19801
302-298-0713
Attorneys for Defendant

VIA ELECTRONIC MAIL

/s/ Rodger D. Smith

Rodger D. Smith II (#3778)

EXHIBIT A

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

SAMSUNG ELECTRONICS CO., LTD.,	)	
SAMSUNG SEMICONDUCTOR, INC.,	)	
SAMSUNG ELECTRONICS AMERICA,	)	
INC., and AVNET, INC.,	)	C.A. No. 25-626-JLH
	)	
Plaintiffs,	)	
	)	DEMAND FOR JURY TRIAL
v.	)	
	)	
NETLIST, INC.,	)	
	)	
Defendant.	)	

**FIRST AMENDED COMPLAINT FOR DECLARATORY JUDGMENT OF
NON-INFRINGEMENT; BREACH OF CONTRACT**

Plaintiffs Samsung Electronics Co., Ltd. (“SEC”), Samsung Semiconductor, Inc. (“SSI”) and Samsung Electronics America, Inc. (“SEA,” and together with SEC and SSI, “Samsung”) and Avnet, Inc. (“Avnet”) (collectively, “Plaintiffs”) seek a declaration that Plaintiffs do not directly or indirectly infringe U.S. Patent No. 12,308,087 (the “’087 patent”) (Exhibit 1) either literally or under the doctrine of equivalents; and a declaration that Plaintiffs do not directly or indirectly infringe U.S. Patent Nos. 10,025,731 (the “’731 patent”) (Exhibit 2), and 12,373,366 (the “’366 patent”) (Exhibit 3), either literally or under the doctrine of equivalents, or, alternatively, a ruling that Defendant Netlist, Inc. (“Netlist”) has breached contractual obligations owed to Samsung, including obligations to license its allegedly essential ’731 and ’366 patents to Samsung and its affiliates on reasonable and non-discriminatory (“RAND”) terms and conditions as follows:

NATURE OF THE ACTION

1. This is an action for declaratory judgment and breach of contract arising under the patent laws of the United States, Title 35 of the United States Code, the Declaratory Judgment Act, 28 U.S.C. § 2201 *et seq.*, and state contract law.

2. This action arises out of Netlist’s repeated assertions of patents against Samsung—and recently against Avnet, Samsung’s customer—after Netlist unilaterally attempted to terminate an agreement in which Netlist granted Samsung a perpetual, paid-up, worldwide license to Netlist patents, including the ’087, ’731, and ’366 patents. Most recently, on May 19, 2025, the day before the ’087 patent issued, Netlist prematurely filed infringement suits in the Eastern District of Texas alleging that Samsung and Micron Technology, Inc. (“Micron”) infringe the ’087 patent. At 12:00 am on May 20, 2025—the moment the ’087 patent issued—Netlist filed a second infringement suit against Samsung, accusing Samsung’s DDR4 and DDR5 products of infringing the ’087 patent. *Netlist, Inc. v. Samsung Elecs., Co., Ltd., et al.*, No. 2-25-cv-557-JRG (E.D. Tex.) (the “EDTX III Action”), D.I. 1, as well as a second suit against Micron asserting the ’087 patent. On July 8, 2025, Netlist amended its complaint, adding, *inter alia*: (i) Avnet, a Samsung customer, as a defendant; (ii) allegations that both Samsung and Avnet infringe the ’731 patent; and (iii) allegations that Avnet infringes the ’087 patent. *Id.*, D.I. 14.

3. The U.S. Patent and Trademark Office (“Patent Office”) recently invalidated, in *inter partes* reviews (“IPR”) filed by Samsung, all claims of two patents related to the ’087 patent—U.S. Patent Nos. 8,787,060 (“the ’060 patent”) and 9,318,160 (“the ’160 patent”). IPR2022-01428, Paper No. 48; IPR2022-01427, Paper No. 48. Netlist had accused Samsung High Bandwidth Memory products (“HBM Products”), including Samsung HBM2, HBM2E, and newer products with eight or more stacked DRAM dies, of infringing the ’060 and ’160 patents. Netlist

recently obtained the '087 patent, a continuation patent from the same patent family as the invalidated patents. Plaintiffs seek a declaratory judgment that they do not infringe the '087 patent.

4. The Patent Office also recently invalidated, in IPRs also filed by Samsung, all claims of two patents related to the '366 patent, U.S. Patent Nos. 11,016,918 (the "'918 patent") and 11,232,054 (the "'054 patent"), that Netlist accused certain Samsung memory modules of infringing. IPR2022-00996, Paper No. 49; IPR2022-00999, Paper No. 51. Specifically, Netlist has accused Samsung DDR5 Load-Reduced Dual Inline Memory Modules ("LRDIMMs"), DDR5 Registered Dual Inline Memory Modules ("RDIMMs"), DDR5 Small Outline Dual Inline Memory Modules ("SODIMMs"), and DDR5 Unbuffered Dual Inline Memory Modules ("UDIMMs") of infringing the '918 and '054 patents. Netlist recently obtained the '366 patent, a continuation patent from the same patent family as the invalidated patents. Samsung seeks a declaratory judgment that Samsung does not infringe the '366 patent.

5. Plaintiffs also seek a declaratory judgment that they do not infringe the '731 patent.

6. Alternatively, Samsung seeks relief for Netlist's breaches of contractual obligations owed to Samsung, including obligations to license the '731 and '366 patents to Samsung and its affiliates on RAND terms and conditions. Netlist has taken the position that parent patents to the '731 and '366 patents are necessarily infringed by the practice of certain standards promulgated by the Joint Electron Device Engineering Council ("JEDEC") and implemented by the Samsung memory modules at issue in this case. Netlist's position on the essentiality of those related patents would likewise apply to the '731 and '366 patents.

THE PARTIES

7. Samsung Electronics Co., Ltd. is a corporation organized and existing under the laws of the Republic of Korea, with its principal place of business at 129, Samsung-ro, Yeongtong-gu, Suwon-si, Gyeonggi-do, 443-742, Republic of Korea.

8. Samsung Semiconductor, Inc. is a corporation organized and existing under the laws of the State of California, with its principal place of business at 3655 North First Street, San Jose, California 95134.

9. Samsung Electronics America, Inc. is a corporation organized and existing under the laws of the State of New York, with its principal place of business located at 85 Challenger Road, Ridgefield Park, New Jersey 07660.

10. Avnet, Inc. is a corporation organized and existing under the laws of the State of New York, with its principal place of business located at 2211 South 47th Street, Phoenix, Arizona 85034.

11. On information and belief, Netlist, Inc. is a corporation organized and existing under the laws of the State of Delaware, with its principal place of business at 175 Technology Drive, Suite 150, Irvine, California 92618.

JURISDICTION AND VENUE

12. This Court has subject matter jurisdiction over the claims for declaratory judgment, which assert claims for declaratory judgment of non-infringement (Counts I, II, and III), under 28 U.S.C. §§ 1331, 1338(a), and 2201(a).

13. This Court has subject matter jurisdiction over the breach of contract claim (Claim IV) pursuant to 28 U.S.C. § 1367. The breach of contract claim forms part of the same case or controversy as the claims for declaratory judgment of non-infringement asserted by Samsung in this action.

14. This Court has personal jurisdiction over Netlist, a corporation organized and existing under the laws of the State of Delaware.

15. Venue is proper in this District under 28 U.S.C. § 1391(b)–(c) because Netlist is subject to personal jurisdiction in this District.

16. An immediate, real, and justiciable controversy exists between Plaintiffs and Netlist as to whether Plaintiffs have infringed the '087 patent, including because Netlist attempted to assert the patent against Samsung before it even issued, and also because Netlist is asserting the patent against Samsung, Micron and Avnet in the Eastern District of Texas .

17. An immediate, real, and justiciable controversy exists between Plaintiffs and Netlist as to whether Samsung has infringed the '731 patent, including because Netlist is asserting the patent against Samsung, Micron and Avnet in the Eastern District of Texas.

18. An immediate, real, and justiciable controversy exists between Samsung and Netlist as to whether Samsung has infringed the '366 patent, including because Netlist asserted against Samsung the parent patents of the '366 patent, which have been invalidated at the PTAB, and Netlist has a history of asserting continuation patents after receiving decisions of invalidity on the parent patents.

A. Infringement Notice and Samsung's First Declaratory Judgment Action

19. Shortly after Netlist unilaterally declared that Samsung was no longer licensed to Netlist's patent portfolio, Netlist sent to Samsung a "Notice of Infringement" letter. In that letter Netlist asserted that "Netlist is the owner of over 100 patents and patent applications related to memory technologies," that Samsung "no longer is licensed to any of Netlist's portfolio of patents," and that certain Samsung memory modules, LRDIMMs and RDIMMs, infringe several of Netlist's patents allegedly related those memory modules.

20. As discussed in detail below, *see infra* ¶¶ 61–64, on October 15, 2021, Samsung filed a declaratory judgment action against Netlist in the District of Delaware. *Samsung Elecs. Co., Ltd. et al. v. Netlist, Inc.*, No. 1:21-cv-1453-JLH (D. Del.) (the “Delaware I Action”), D.I. 1. Netlist moved to dismiss Samsung’s declaratory judgment claims for lack of jurisdiction. *Id.*, D.I. 11 at 3–4. The Court concluded that it had and would exercise jurisdiction over Samsung’s claims relating to three patents,¹ *id.*, D.I. 37 at 4–5, and Netlist counterclaimed for infringement of those patents, *id.*, D.I. 40. Netlist also brought Google, LLC (“Google”), Samsung’s customer, into the case and asserted counterclaims against Google. *Id.*, D.I. 45.

B. Netlist’s Assertion of Patents Against Samsung

21. As detailed *infra* ¶¶ 65–70, on December 20, 2021, Netlist began a patent litigation campaign against Samsung in the U.S. District Court for the Eastern District of Texas. *Netlist, Inc. v. Samsung Elecs. Co., Ltd. et al.*, No. 2:21-cv-463-JRG (E.D. Tex.) (the “EDTX I Action”), D.I. 1 ¶ 36. Netlist asserted the ’918 patent against Samsung, among other patents. In the complaint, Netlist also indicated its intent to later assert an allowed but still-pending continuation of the ’918 patent, the ’054 patent, which later issued. Netlist subsequently amended its complaint in the EDTX I Action to assert the ’054 patent. The ’366 patent at issue in this case is a continuation patent of the ’918 and ’054 patents. In the EDTX I Action, Netlist accused Samsung DDR5 LRDIMMs, DDR5 RDIMMs, DDR5 SODIMMs, and DDR5 UDIMMs of infringing the ’918 and ’054 patents.

¹ The Court dismissed Samsung’s declaratory judgment claim of non-infringement related to U.S. Patent No. 7,619,912, which had been “the subject of litigation [against Google in the Northern District of California] for over a decade,” in favor of the parties litigating the patent in California. Delaware I Action, D.I. 37 at 6.

22. In its amended complaint in EDTX I, Netlist also added the '060 and '160 patents. The '087 patent at issue in this case is a continuation patent of the '060 and '160 patents. In the EDTX I Action, Netlist accused some of the same Samsung HBM products at issue in this case, i.e., HBM2, HM2E, and HBM3. Netlist advanced the '918, '054, '060 and '160 patents against Samsung HBM products and certain Samsung DIMM products through trial in the Eastern District of Texas, where the Texas jury rendered a verdict of infringement of the '918, '054, '060 and '160 patents in Netlist's favor and awarded damages. The case is currently on appeal. *Netlist, Inc. v. Samsung Elecs. Co., Ltd. et al.*, 24-2203 (Fed. Cir.).

23. Netlist later filed a second infringement suit against Samsung in the Eastern District of Texas, asserting patents that Netlist had not asserted in the EDTX I Action. *Netlist, Inc. v. Samsung Elecs. Co., Ltd. et al.*, No. 2:22-cv-293-JRG, D.I. 1 (E.D. Tex.) (the "EDTX II Action").

24. As noted, on May 19, 2025, Netlist filed a third infringement suit against Samsung in the Eastern District of Texas, asserting the '087 patent—the day before the '087 patent had issued from the Patent Office. *Netlist, Inc. v. Samsung Elecs. Co., Ltd. et al.*, No. 2:25-cv-553-JRG, D.I. 1 (E.D. Tex.). The next day, shortly after midnight, Netlist filed another infringement suit in the Eastern District of Texas, again alleging infringement of the '087 patent. EDTX III Action. Netlist later dismissed its premature infringement case asserting the '087 patent. *Netlist, Inc. v. Samsung Elecs. Co., Ltd. et al.*, No. 2:25-cv-553-JRG, D.I. 2 (E.D. Tex.). Netlist subsequently moved to amend its EDTX III complaint to add the '731 patent and Avnet. EDTX III, D.I. 14.

25. Netlist has also filed suit against Samsung in Germany. *Netlist, Inc. v. Samsung Elecs. Co., Ltd.*, Nos. 4c O 31/22, 4c O 45/22 (Düsseldorf Regional Court).

C. Netlist's Patent Infringement Suits Against Other Memory Module Manufacturers and Technology Companies

26. Before Netlist's patent dispute with Samsung, Netlist began a patent litigation campaign against a competitor of Samsung, SK hynix, alleging that certain memory products infringed Netlist's patents. In total, Netlist filed six patent infringement actions against SK hynix from 2016 to 2020. *See infra* ¶ 71. Throughout its litigation campaign against SK hynix, Netlist demonstrated a pattern of asserting patents, receiving rulings of non-infringement or invalidity, and then asserting in a new action continuation patents of the non-infringed or invalid patents. *See infra* ¶ 76.

27. Further, Netlist has recently asserted patents—including the '060 and '160 patents—against Micron and its affiliates, which are competitors of Samsung with respect to at least certain of Samsung's memory products. *See infra* ¶ 74. Netlist also filed suit against Micron in Germany.

28. Netlist has also filed infringement suits against several technology companies in California, *see infra* ¶ 74, including against Google in the Northern District of California, *Netlist, Inc. v. Google LLC*, No. 3:090-cv-05718-RS (N.D. Cal.), against Inphi Corporation in the Central District of California, *Netlist, Inc. v. Inphi Corp.*, No. 2:09-cv-6900 (C.D. Cal. 2009), and against Smart Modular Technologies, Inc.; Diablo Technologies, Inc.; SanDisk, LLC; Smart Storage Systems, Inc.; and Smart Worldwide Holdings, Inc. in the Northern District of California, *Netlist, Inc. v. Smart Modular Technologies, Inc. et al.*, 8-13-cv-00996 (N.D. Cal. 2013). Netlist asserted infringement counterclaims against Diablo Technologies, Inc. in the Northern District of California. *Diablo Technologies, Inc. v. Netlist, Inc.*, 4:13-cv-03901 (N.D. Cal. 2013). Netlist sued Google in Germany for patent infringement.

D. Samsung's IPRs on the '918 and '054 Patents

29. On May 17, 2022, Samsung filed petitions for *inter partes* review of the '918 and '054 patents in the PTAB. IPR2022-0996; IPR2022-0999. On December 5 and 6, 2023, the PTAB issued Final Written Decisions in both IPRs, invalidating all claims of both the '918 and '054 patents. IPR2022-0996, Paper No. 49; IPR2022-0999, Paper No. 51.

30. Netlist appealed the PTAB's Final Written Decisions on the '918 and '054 patents. *Netlist, Inc. v. Samsung Elecs. Co., Ltd.*, 2024-1859, 2024-1863 (Fed. Cir.). If the Federal Circuit affirms the Final Written Decisions on the '918 and '054 patents before the Federal Circuit issues its decision in the EDTX I appeal, Netlist's jury verdict against Samsung on the '918 and '054 patents will be vacated.

E. Samsung's IPRs on the '060 and '160 Patents

31. On August 26, 2022, Samsung also filed petitions for *inter partes* review of the '060 and '160 patents in the PTAB. IPR2022-01428; IPR2022-01427. On April 1, 2024, the PTAB issued Final Written Decisions in both IPRs, invalidating all claims of both the '060 and '160 patents. IPR2022-01427, Paper No. 48; IPR2022-01428, Paper No. 48.

32. Netlist appealed the PTAB's Final Written Decisions on the '060 and '160 patents. *Netlist, Inc. v. Samsung Elecs. Co., Ltd.*, 2024-2240, 2024-2241 (Fed. Cir.). Netlist's jury verdict against Samsung on the '060 and '160 patents will be vacated if the Federal Circuit affirms the Final Written Decisions on the '060 and '160 patents before the Federal Circuit issues its decision in the EDTX I Action appeal.

F. Netlist's Threats of Future Infringement Actions

33. Since the dispute between Netlist and Samsung arose, Netlist has made demands that Samsung take a second license to Netlist's portfolio of patents based on Netlist's

claim that it terminated the fully paid-up, worldwide patent license granted in the parties' earlier license agreement.

34. For instance, as noted above, on October 15, 2020, Netlist issued a "Notice of Infringement" to Samsung and demanded that Samsung take a second license to Netlist's patents. Specifically, Netlist demanded that "Samsung and its subsidiaries honor third-party intellectual property rights" and engage in "formal licensing discussions"—"reserv[ing] all rights and remedies" with respect to enforcing its patents against Samsung.

35. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

36. Netlist has also recently stated in a regulatory filing that the "memory market remains poised for significant growth driven by the industry's transition to HBM" and "Netlist is well-positioned to capitalize on this through new product development and its intellectual property portfolio." Netlist, Inc., Securities and Exchange Commission Form 8-K (May 6, 2025), <https://investors.netlist.com/websites/netlist/English/3210/us-sec->

filing.html?shortDesc=Current%20report%20filing&format=convpdf&secFilingId=e1eae036-68a8-4165-a279-3a75d35f1a7c (last visited July 28, 2025). Netlist has therefore stated publicly that it intends to enforce its HBM patent portfolio to “capitalize” on the growth of the market.

37. Accordingly, as set forth herein, Netlist has asserted patents in the same patent family as the '087 patent—*i.e.*, the '060 and '160 patents—against Samsung HBM products currently being sold and used throughout the United States, and Netlist has accused other Samsung memory products of infringing additional patents in multiple infringement actions. Netlist has also sought to enforce patents in the same patent family as the '366 patent—*i.e.*, the '918 and '054 patents—against Samsung DDR5 products currently being sold and used throughout the United States.

38. Netlist's decision to appeal the PTAB's Final Written Decisions invalidating the '918, '054, '060 and '160 patents demonstrates Netlist's intent to continue asserting these two patent families, which includes the '087 and '366 patents, against Samsung HBM products and certain Samsung DIMM products notwithstanding the PTAB's findings of invalidity. Netlist's prosecution of the applications that led to the issuance of the '087 and '366 patents and Netlist's history of asserting continuation patents in litigation after having the parent patents invalidated further suggests that Netlist's patent campaign against Samsung, including with regard to the '087 and '366 patent families and Samsung HBM products and certain Samsung DIMM products, is not over.

39. Indeed, as explained above, on May 19, 2025, the day before the '087 patent issued, Netlist prematurely filed infringement suits in the Eastern District of Texas alleging that Samsung and Micron infringe the '087 patent. At 12:00 am on May 20, 2025—the moment the

'087 patent issued—Netlist filed a second infringement suit against Samsung, accusing Samsung's HBM products of infringing the '087 patent. EDTX III Action, D.I. 1. On July 8, 2025, Netlist amended its complaint, adding the '731 patent and Avnet. *Id.*, D.I. 14 and accusing Samsung's DDR4 and DDR5 products.

40. Because this action presents an actual controversy with respect to the '087, '731, and '366 patents, the Court may and should grant the declaratory relief sought pursuant to 28 U.S.C. § 2201 *et seq.*

BACKGROUND

G. Netlist's Extraordinary License Demand and Infringement Threats

41. On November 12, 2015, Netlist and SEC entered into a Joint Development and License Agreement (the "Agreement" or "JDLA"). The Agreement contains cross-license, joint development, and product supply provisions.

42. In the Agreement, Netlist granted SEC and its subsidiaries, including SSI and SEA, a perpetual, paid-up, worldwide, non-exclusive license to all patents owned or controlled by Netlist or any of its subsidiaries having an effective first filing date on or prior to November 12, 2020. The license extends through the expiration of the last to expire of the licensed patents.

43. SEC similarly granted Netlist and its subsidiaries a perpetual, paid-up, worldwide, non-exclusive license to all patents owned or controlled by SEC or any of its subsidiaries having an effective first filing date on or prior to November 12, 2020.

44. The Agreement required SEC to make a payment to Netlist of \$8 million subject to the terms and conditions therein. SEC made this payment to Netlist, except for a portion that was remitted as tax withholding to the Korean tax authorities, which later refunded the withheld portion to Netlist.

45. Since 2020, Netlist has taken extraordinary actions to back out of its grant of a patent license to Samsung.

46. On May 27, 2020, Netlist’s Chief Licensing Officer, Marc J. Frechette, wrote to Mr. Seung Min Sung of SEC and alleged that Samsung materially breached the Agreement by “repeatedly fail[ing] to fulfill Netlist’s request for NAND and DRAM products throughout the term of the Agreement” and, allegedly, by improperly deducting withholding taxes. Netlist did not allege, and has never alleged, that Samsung breached the Agreement in any way related to the patent cross-license. In the same letter—which was the first time Netlist raised its breach allegations—Mr. Frechette informed Mr. Sung that Netlist had filed a complaint in the U.S. District Court for the Central District of California with respect to the alleged breach by Samsung and provided a copy of the complaint for reference.

47. On May 28, 2020, the U.S. District Court for the Central District of California docketed Netlist’s breach of contract complaint against SEC. *Netlist Inc. v. Samsung Elecs. Co., Ltd.*, No. 8:20-cv-00993-MCS (C.D. Cal.) (the “Breach Action”).

48. Netlist subsequently wrote to SEC to terminate the Agreement, including the patent license to SEC and its subsidiaries. On July 15, 2020, Netlist’s Chief Licensing Officer, Marc J. Frechette, wrote to Mr. Seung Min Sung of SEC and stated that “Netlist is hereby terminating, effective immediately, the Agreement including the patent license granted to Samsung”

49. One week later, on July 22, 2020, Netlist amended its complaint in the Breach Action, seeking a declaration that the license it granted Samsung, but not the license Samsung granted Netlist, was terminated. In the Breach Action, Netlist sought monetary damages

and a declaration “that Netlist has terminated the Agreement pursuant to Section 13.2 and that Samsung’s licenses and rights under the agreement have ceased.”

50. Netlist then issued a “Notice of Infringement” to Samsung and a demand that Samsung take a second license to Netlist’s patents. Specifically, on October 15, 2020, Netlist stated that because it had terminated the Agreement, “Samsung is no longer licensed to any of Netlist’s portfolio of patents.” Netlist further asserted that its patents relate to “memory technologies,” which include HBM, LRDIMMs, RDIMMs, SODIMMs, and UDIMMs. Netlist then alleged: “Despite the termination of the Joint Development and License Agreement, Samsung continues to unlawfully utilize Netlist’s innovations and to infringe Netlist’s patents.” Netlist concluded the letter by demanding that “Samsung and its subsidiaries honor third-party intellectual property rights” and engage in “formal licensing discussions.” Netlist further stated that it “reserves all rights and remedies” with respect to the alleged infringement.

51. Netlist resumed its communications with Samsung in 2021. On February 2, 2021, Netlist’s Chief Licensing Counsel, Mr. Marc J. Frechette, indicated by email that Samsung must take a second license, coupled with the demand that Netlist “be made whole” for any alleged breach asserted in the Breach Action.

52. Accordingly, Netlist seeks to double-dip, with a demand that Samsung take a second license to Netlist’s patents.

53. On October 14, 2021, the court in the Breach Action issued a confidential order granting summary judgment in favor of Netlist on Netlist’s Third Claim for Relief, which sought a declaration that Netlist terminated Samsung’s license effective July 20, 2020. *See Breach Action*, D.I. 186. The court also granted summary judgment in Netlist’s favor that Samsung was liable for certain alleged breaches of other provisions in the Agreement. *See id.*

54. Beginning on December 1, 2021, the court in the Breach Action held a jury trial on damages for one of the alleged breaches, the alleged failure to fulfill Netlist's orders for NAND and DRAM products. On December 3, 2021, the jury returned a verdict in Samsung's favor, finding that Netlist failed to prove that it suffered any damages for this alleged breach. *See id.*, D.I. 276. The court in the Breach Action concluded that Netlist was not entitled to damages for the remaining alleged breach (the withholding of taxes) as a matter of law and awarded Netlist nominal damages of \$1.00 for each alleged breach (for a total of \$2.00). *Id.*, D.I. 306, ¶¶ 1–2.

55. Despite the district court awarding only nominal damages for Samsung's alleged breaches, Netlist issued a press release claiming to be “very pleased with the overall outcome of the case as it confirmed that Samsung no longer has a valid license to Netlist patents and therefore requires a licensing agreement.” Netlist Provides Update on Action in Federal Court Against Samsung, <https://investors.netlist.com/websites/netlist/English/2120/us-press-release.html?airportNewsID=a661884d-0025-46ea-a055-bae8d8e4570b> (last visited July 28, 2025).

56. SEC challenged the summary judgment ruling in the Breach Action on appeal to the Ninth Circuit. *Netlist Inc. v. Samsung Elecs. Co., Ltd.*, Nos. 22-55209 & 22-55247 (9th Cir.). SEC maintained that it did not breach the Agreement, that the Agreement has not been terminated, and that SEC continues to hold a paid-up, worldwide, non-exclusive license to Netlist's patents, including the '087, '731, and '366 patents.

57. The Ninth Circuit agreed that patent cross-licensing was a distinct aim of the JDLA, independent of the joint development project: “[T]he JDLA has two stated purposes: (1) developing a new NVDIMM-P product (i.e., the [joint development project]), and (2) patent cross-licensing. The title and preamble of the agreement exclusively reference these two topics,

and each substantive section corresponds entirely to one of the two goals.” Breach Action, D.I. 192, Ex. A at 3–4. The Ninth Circuit remanded the case to (a) resolve an ambiguity in the supply provision in Section 6.2 of the JDLA and (b) determine whether any breach of that provision by Samsung was material such that Netlist could terminate the agreement. Following the Ninth Circuit’s remand, a jury trial was held. On May 17, 2024, the jury issued its verdict, adopting Netlist’s proffered meaning of the disputed provision and finding that Samsung’s breach of that provision was material. *See id.*, D.I. 556.

58. In post-trial briefing, Samsung requested a new trial. The district court granted Samsung’s motion on account of implied juror bias. *Id.*, D.I. 640 at 16.

59. On March 24, 2025, the district court held a new trial, and the jury again adopted Netlist’s proffered meaning of the disputed provision. *See id.*, D.I. 769. Samsung is challenging the jury verdict in post-trial motions. The court has scheduled an evidentiary hearing on Samsung’s motion for a new trial for July 30, 2025. *Id.*, D.I. 810. If those challenges are unsuccessful, Samsung intends to appeal the verdict and judgment to the United States Court of Appeals for the Ninth Circuit.

H. Netlist’s Patent Infringement Lawsuits

60. Over the past decade, Netlist has pursued a litigation campaign against Samsung and other suppliers of memory modules, component suppliers, and users of memory modules (aside from Samsung). During this campaign, Netlist has alleged infringement of numerous patents, including patents related to the ’087 and ’366 patents, and accused several memory products, including DDR4, DDR5 and HBM Products.

1. Netlist’s Lawsuits Against Samsung

61. As noted above, on October 15, 2021, Samsung filed a declaratory judgment action against Netlist in this District for a declaration that certain Samsung LRDIMM and RDIMM

products did not infringe four Netlist patents. Delaware I Action, D.I. 1. Netlist moved to dismiss Samsung's declaratory judgment claims, arguing that no case or controversy existed because there was no imminent threat of Netlist filing suit against Samsung on the patents-in-suit. *Id.*, D.I. 11 at 3–4. Samsung later amended its complaint to add additional Netlist patents, *id.*, D.I. 14, and Netlist moved to dismiss Samsung's First Amended Complaint, *id.*, D.I. 24.

62. On August 1, 2022, the Court in the Delaware I Action denied-in-part and granted-in-part Netlist's motion to dismiss. The Court found that it had jurisdiction on U.S. Patent Nos. 10,217,523 (the "'523 patent"), 9,858,218 (the "'218 patent"), and 10,474,595 (the "'595 patent"), three of the four patents included in the original declaratory judgment complaint because, once the Central District of California granted summary judgment in favor of Netlist in the Breach Action, "Netlist's threats became more immediate." *Id.*, D.I. 37 at 4. However, the Court dismissed, *inter alia*, Samsung's declaratory judgment claim of non-infringement related to U.S. Patent No. 7,619,912 (the "'912 patent."). *Id.*, D.I. 37. The Court reasoned that the patent had been "the subject of litigation [against Google in the Northern District of California] for over a decade," so Samsung's declaratory judgment claim regarding that patent in Delaware was not "an efficient use of judicial resources." *Id.* at 6. In reaching its decision, the Court accepted Netlist's argument that Samsung's "collateral action [in Delaware] should be dismissed in favor of the much more advanced Northern District case." *Id.* at 3; *see also id.*, D.I. 11 at 3–5; D.I. 28 at 1–3, 8–10 (Netlist arguing that litigation on the '912 patent should proceed in the Northern District of California).

63. Following the Court's order in the Delaware I Action, Netlist counterclaimed against Samsung for infringement of the '523, '218, and '595 patents. Netlist also sought treble damages based on Samsung's alleged willful infringement. *Id.*, D.I. 45 ¶ 101. Netlist

subsequently expanded its infringement allegations to include additional Samsung products. *Id.*, D.I. 132-1, Ex. 16 at 2.

64. Netlist later asserted counterclaims against Google in the Delaware I Action for infringement of the '523, '218, and '595 patents. *Id.*, D.I. 45. Netlist's infringement allegations relate to Google's use of Samsung memory products.

65. On December 20, 2021, Netlist filed a lawsuit against SEC, SSI, and SEA in the Eastern District of Texas alleging direct and indirect infringement of three Netlist patents. EDTX I Action, D.I. 1. Netlist alleged infringement based on Samsung's making, selling, using, and/or importing of certain memory modules. *See id.*, ¶¶ 40, 53, 68.

66. Netlist filed the EDTX I Action the same day that it moved to dismiss the Delaware I Action based on, *inter alia*, an alleged lack of any case or controversy over Samsung's declaratory judgment claims. *See* Delaware I Action, D.I. 11 at 3 (arguing that the case did not involve a "substantial controversy of sufficient immediacy").

67. Netlist's complaint in the EDTX I Action asserted the '918 patent, among other patents. Netlist also stated that it intended to assert against Samsung a then-pending continuation of the '918 patent once the patent issued from the Patent Office. *See* EDTX I Action, D.I. 1 ¶ 32. The application issued on January 25, 2022, as the '054 patent, and Netlist subsequently amended its complaint to include the newly issued patent. *Id.*, D.I. 23. Netlist's amended complaint also added two additional patents: the '060 and '160 patents. *Id.* Netlist alleged that Samsung HBM products infringe those patents. *Id.*, D.I. 23 ¶¶ 118, 128.

68. Netlist also filed a second suit against Samsung in Texas. The same day the court in the Delaware I Action dismissed Samsung's declaratory judgment claims involving the '912 patent—reasoning that litigating the patent in the Northern District of California would be

the most efficient use of judicial resources—Netlist filed suit against Samsung on the '912 patent in the Eastern District of Texas. *See* EDTX II Action, D.I. 1. Netlist later amended its complaint to include three additional patents. Netlist alleged that Samsung infringed its patents by making, using, selling, offering to sell, and/or importing certain memory modules. *Id.*, ¶ 34.

69. On May 19, 2025, Netlist prematurely filed an infringement suit against Samsung on the '087 patent before the patent had issued. *Netlist, Inc. v. Samsung Elecs. Co., Ltd. et al.*, No. 2:25-cv-00553 (E.D. Tex.). At 12:00 am on May 20, 2025—the moment the '087 patent issued—Netlist filed a second infringement suit against Samsung, accusing Samsung's DDR5 products of infringing the '087 patent. EDTX III Action, D.I. 1. In the complaint, Netlist alleges that Samsung “directly infringed and [is] currently infringing at least one of the approved claims of the '087 Patent,” with respect to “any Samsung HBM2, HBM2E, HMB3, HBM3E and newer products (e.g., HBM4) made, sold, used, offered for sale, and/or imported into the United States by Samsung,” *id.*, D.I. 1 ¶¶ 23, 27, collectively, along with HBM4E, “Samsung HBM Products.” Netlist later dismissed its premature infringement case asserting the '087 patent. *Netlist, Inc. v. Samsung Elecs. Co., Ltd. et al.*, No. 2:25-cv-553-JRG, D.I. 2 (E.D. Tex.). Netlist subsequently amended its EDTX III complaint to add the '731 patent and Avnet. EDTX III, D.I. 14.

70. Netlist also filed suit against Samsung in Germany, asserting two patents against Samsung's memory modules. *Netlist, Inc. v. Samsung Elecs. Co., Ltd.*, Nos. 4c O 31/22, 4c O 45/22 (Düsseldorf Regional Court).

2. Netlist's Lawsuits Against Third Parties

71. Netlist has also pursued a litigation campaign against other suppliers of memory products, component suppliers, and users of memory products. For instance, Netlist engaged in a years-long, multi-forum litigation campaign against SK hynix, a competitor of Samsung. Between 2016 and 2020, Netlist filed six patent infringement suits against SK hynix,

the first two in the Central District of California. *See Netlist, Inc. v. SK hynix America Inc. et al.*, No. 8:16-cv-1605 (C.D. Cal. 2016); *Netlist, Inc. v. SK hynix, Inc. et al.*, No. 8:17-cv-1030 (C.D. Cal. 2017). Netlist also filed two complaints against SK hynix with the International Trade Commission. *See In re Certain Memory Modules and Components Thereof, and Products Containing the Same*, No. 337-TA-1023 (U.S.I.T.C. 2016); *In re Certain Memory Modules and Components Therefore*, No. 337-TA-1089 (U.S.I.T.C. 2017). Finally, Netlist filed two infringement suits against SK hynix in the Western District of Texas. *See Netlist, Inc. v. SK hynix Inc. et al.*, No. 6:20-cv-194 (W.D. Tex. 2020); *Netlist, Inc. v. SK hynix America. et al.*, No. 6:20-cv-525 (W.D. Tex. 2020).

72. In these suits, Netlist accused certain SK hynix memory modules of infringing numerous Netlist patents.

73. Netlist has also accused Google of infringing the '912 patent in connection with LRDIMM and RDIMM products. *See Netlist, Inc. v. Google Inc.*, No. 3:09-cv-05718-RS (N.D. Cal.). Netlist seeks an injunction in that case, which would preclude Google from using Samsung JEDEC-compliant LRDIMM and RDIMM memory modules. Netlist also seeks an injunction against Google in the Delaware I Action. D.I. 127 at 11.

74. Netlist has also asserted infringement claims against Inphi Corporation, *see Netlist, Inc. v. Inphi Corp.*, No. 2:09-cv-6900 (C.D. Cal. 2009), and infringement counterclaims against Diablo Technologies, Inc., *Diablo Technologies, Inc. v. Netlist, Inc.*, 4-13-cv-03901 (N.D. Cal. 2013). Netlist also filed an infringement suit against Smart Modular Technologies, Inc.; Diablo Technologies, Inc.; SanDisk, LLC; Smart Storage Systems, Inc.; and Smart Worldwide Holdings, Inc. *Netlist, Inc. v. Smart Modular Technologies, Inc. et al.*, 8-13-cv-00996 (N.D. Cal. 2013).

75. As noted above, Netlist has also recently asserted patents in multiple litigations in Texas against Micron, which is a competitor of Samsung with respect to at least certain of the memory products, including HBM and DIMM products. *See Netlist, Inc. v. Micron Tech., Inc. et al.*, No. 6:21-cv-430 (W.D. Tex.); *Netlist, Inc. v. Micron Tech., Inc. et al.*, No. 6:21-cv-431 (W.D. Tex.); *Netlist, Inc. v. Micron Tech., Inc. et al.*, No. 2:22-cv-203-JRG (E.D. Tex.); *Netlist, Inc. v. Micron Tech., Inc. et al.*, No. 2:22-cv-294 (E.D. Tex.). As with Samsung, Netlist filed suit against Micron for infringement of the '087 patent before it issued. *Netlist, Inc. v. Micron Tech., Inc. et al.*, No. 2:25-cv-552 (E.D. Tex.). Netlist then filed another complaint against Micron the day the '087 patent issued. *Netlist, Inc. v. Micron Tech., Inc. et al.*, No. 2:25-cv-558 (E.D. Tex.). Netlist subsequently amended its complaint to add the '731 patent and Avnet. *Id.*, D.I. 15.

76. As already noted, throughout its litigation campaigns, Netlist has demonstrated a pattern of asserting patents, having IPRs instituted or receiving rulings of non-infringement or invalidity, and then asserting continuations of the non-infringed or invalid patents. For example, Netlist asserted the '837 patent against SK hynix in district court and at the ITC. *Netlist, Inc. v. SK hynix Inc.*, No. 8:16-cv-01605 (C.D. Cal. Aug. 31, 2016), D.I. 1; *In re Certain Memory Modules*, No. 337-TA-1023 (U.S.I.T.C. Sept. 1, 2016), Doc. No. 589609. Shortly before the PTAB determined that certain claims of the '837 patent were unpatentable (and the ITC determined that SK hynix did not infringe the patent), *see* IPR2017-00548, Paper No. 25 (May 3, 2018); *In re Certain Memory Modules*, No. 337-TA-1023, 2017 WL 11569255, at *135 (Nov. 14, 2017), Netlist asserted a continuation of the '837 patent against SK hynix in district court and at the ITC. *Netlist, Inc. v. SK hynix Inc.*, No. 8:17-cv-1030 (C.D. Cal.), D.I. 1; *In re Certain Memory Modules*, No. 337-TA-1089 (U.S.I.T.C. Oct. 31, 2017), Doc. No. 627249. The PTAB subsequently determined that all claims of that patent were unpatentable, *see* IPR2018-00303, Paper No. 42 (Mar. 21, 2019), so Netlist asserted two additional

continuation patents in the same patent family against SK hynix, *Netlist, Inc. v. SK hynix Inc.*, No. 6:20-cv-00194 (W.D. Tex. June 9, 2020), D.I. 33, essentially seeking a third bite at the apple.

I. Netlist Patents and Standard Essential Allegations

1. Standard Essential Allegations Regarding the '731 Patent

77. The '731 patent is a continuation patent of U.S. Patent No. 8,154,901, which issued on April 10, 2010, from U.S. Patent App. 12/422,853 (the "'853 application"). On information and belief, on November 22, 2010, during development of JEDEC standards related to earlier versions of the memory modules Netlist now accuses of infringement, Netlist disclosed the '853 application as essential to standards being promulgated by JEDEC committee JC-42.3 and JC-45. On September 7, 2011, Netlist again disclosed the '853 application as essential to standards under consideration by JEDEC committee JC-40.

78. In the present action, Samsung seeks declarations that Samsung's products that comply with JEDEC DDR4 and DDR5 standards, including Samsung's DDR5 RDIMM, DDR5 SODIMM, DDR5 UDIMM, DDR5 MRDIMM, DDR5 EUDIMM, DDR5 SOEDIMM, and DDR4 LRDIMM memory modules do not infringe the '731 patent.

79. If, as Netlist has contended, the '731 patent is essential to any of the JEDEC standards (which they are not), Netlist is obligated to license the patent on RAND terms. Netlist has failed to offer such a license to Samsung.

2. Standard Essential Allegations Regarding the '366 Patent

80. The '366 patent is a continuation patent of U.S. Patent Nos. 8,874,831 (the "'831 patent") and 8,301,833 (the "'833 patent"). On information and belief, on August 7, 2018, during development of JEDEC standards related to earlier versions of the memory modules Netlist

now accuses of infringement, Netlist disclosed the '831 and '833 patents as essential to standards being promulgated by JEDEC committee JC-40.

81. In the present action, Samsung seeks declarations that Samsung's DIMM products that comply with JEDEC DDR4 and DDR5 standards, including Samsung's DDR5 RDIMM, DDR5 SODIMM, DDR5 UDIMM, DDR5 MRDIMM, DDR5 EUDIMM, DDR5 SOEDIMM, and DDR4 LRDIMM memory modules do not infringe the '366 patent.

82. If, as Netlist has contended, the '366 patent is essential to any of the JEDEC standards (which they are not), Netlist is obligated to license the patent on RAND terms. Netlist has failed to offer such a license to Samsung.

J. Overview of the '087 Patent

83. The '087 patent issued on May 20, 2025, from U.S. Patent Application No. 17/694,649, which was filed on March 14, 2022. The '087 patent claims priority to a chain of continuation applications, including the applications leading to issuance of the aforementioned '060 and '160 patents. The first full patent application was U.S. Patent Application No. 13/288,850, which was filed on November 3, 2011, and claimed priority to a provisional application that was filed on November 3, 2010.

84. During prosecution, the examiner rejected the then-pending claim 1 and 2 of the '087 patent for double patenting over the claims of the '060 and '160 patents. Netlist filed a terminal disclaimer on April 5, 2023, thereby disclaiming any statutory term that would extend beyond the expiration date of multiple patents, including the '060 and '160 patents.

85. The '087 patent describes a "memory package having stacked array dies and methods for reducing driver load in the memory package." Ex. 1 ('087 Patent) at [0002]. According to the patent, "[m]emory modules may include a number of memory packages," and "[e]ach memory package may itself include a number of array dies that are packaged together."

Id. at [0003]. A “driver” is used for “driving signals” between dies, such as “driver 134” in Figure 1A. *Id.* at [0004]. The ’087 patent states that “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.” *Id.* at [0007]. The ’087 patent continues to state that “to drive a signal along a die interconnect, a driver typically must be large enough to overcome the load on the driver,” but “generally a larger driver not only consumes more space on the control die, but also consumes more power.” *Id.*

86. The ’087 patent explains that Figure 2, which is reproduced below, “schematically illustrates an example embodiment of a memory package 200 in accordance with the present disclosure.” *Id.* at [0025], Fig. 2. The patent describes memory package 200 by stating that it “may include a plurality of array dies 210 (*e.g.*, array dies 210a-210d),” “a plurality of die interconnects 220 (*e.g.*, die interconnects 220a, 220b),” and “a control die 230” that “may include a number of data conduits 232.” *Id.* at [0025], [0028], [0031], [0032].




FIG. 2

87. For at least the reasons explained in Count I, the Samsung HBM Products do not infringe the claims of the '087 patent.

K. Overview of the '731 Patent

88. The '731 patent issued on July 17, 2018, from U.S. Patent Application No. 14/715,491, which was filed on May 18, 2015. The '731 patent claims priority to a chain of continuation applications. The first full patent application was U.S. Patent Application No. 12/422,853, which was filed on April 13, 2009, and claimed priority to U.S. Provisional Application Nos. 61/044,839, 61/044,825, and 61/044,801, each of which were filed on April 14, 2008.

89. The '731 patent describes “a memory module having a printed circuit board including at least one connector configured to be operatively coupled to a memory controller of a computer system.” Ex. 2 ('731 patent) at Abstract. According to the patent, “a single memory controller may gradually degrade system performance and increase power consumption as the physical and logical size of the memory subsystem increases.” *Id.* at 4:35-38. The patent states that it “address[es] these speed and power issues,” by using “a circuit . . . [that] can comprise an isolation switch and can logically isolate inactive subsystems and reduce signal reflection and/or other noise at, for example, the interface between the system board and the memory board and at the interface between the memory board and the memory devices.” *Id.* at 4:49-55.

90. The '731 patent explains that Figure 3, which is reproduced below, “schematically illustrates an example system incorporating a memory module 100 in accordance with certain embodiments described herein.” *Id.* at 5:14-16. The patent states that “[t]he memory module 100 comprises a printed circuit board (PCB) 102,” which “includes at least one connector 104 configured to be operatively coupled to a memory controller 106 of a computer system 108,” as well as “a plurality of memory devices 110.” *Id.* at 5:16-21. The patent also states that “[t]he memory module 100 . . . further includes a circuit 112 comprising a first set of ports 122,” “a second set of ports 128,” “a switching sub-circuit 140 configured to selectively operatively couple one or more ports 128 . . . to one or more ports 124,” and that this “reduces noise in one or more signals transmitted between one or more ports 124 . . . and the one or more ports 128.” *Id.* at 5:22-40.




FIG. 3

91. For at least the reasons explained in Count II, Plaintiffs do not infringe the claims of the '731 patent.

L. Overview of the '366 Patent

92. The '366 patent issued on July 29, 2025, from U.S. Patent Application No. 17/582,797, which was filed on January 24, 2022. The '366 patent claims priority to a chain of continuation applications. The first full patent application was U.S. Patent Application No. 12/131,873, which was filed on June 2, 2008, and claimed priority to U.S. Provisional Application

Nos. 61/512,871, which was filed on July 28, 2011, and 60/941,586, which was filed on June 1, 2007.

93. The '366 patent relates to "Flash-DRAM Hybrid Memory Modules" that supply power to the memory system from a system power supply when the memory module is in the first state, and from an on-board power supply when the memory module is in a second state. Ex. 3 ('366 patent) at [0160]. Figure 12 illustrates a memory module disclosed in the '366 patent:




FIG. 12

94. The memory system (1010) includes a volatile memory subsystem (1030), a non-volatile memory subsystem (1040), and a controller (1062) coupled to the non-volatile memory subsystem. *Id.* at [0146]. The on-board power supply (1080) can be implemented as a

power module (1100) that provides multiple voltages to the circuitry on the memory module. *Id.* at [0167]. Figure 16 illustrates the power module disclosed in the '366 patent:




FIG. 16

95. The '366 patent states “[t]he conversion element 1120 may comprise a plurality of sub-blocks 1122, 1124, 1126 as schematically illustrated by FIG. 16, which can provide more voltages in addition to the second voltage 1104 to the memory system 1010.” *Id.* at [0173].

96. For at least the reasons explained in Count III, the Samsung DDR4 and DDR5 Products do not infringe the claims of the '366 patent.

M. Netlist Has Breached Contractual Obligations Owed to Samsung

1. JEDEC and Semiconductor Memory Standards

97. JEDEC is an independent, non-profit semiconductor engineering trade association and standardization body based in the United States. Over 300 companies in the

computer and electronics industries are currently members of JEDEC, including Samsung and Netlist.

98. JEDEC has become the global leader in developing open standards and publications for a broad range of semiconductor technologies, including memory products. As relevant here, JEDEC develops standards for semiconductor memory chips and modules, including standards implemented by Samsung's products.

99. JEDEC's semiconductor memory standards enable interoperability, *i.e.*, the ability of memory products made by different manufacturers to work together with computer and electronic devices made by others. JEDEC standards are widely implemented, and, for many types of semiconductor memory products, it is imperative that they comply with JEDEC standards in order to be commercially viable.

100. Member companies participate in JEDEC's standard-setting process through over 50 committees and subcommittees that address various technological areas. Through a collaborative process, members contribute to and vote on technical proposals for incorporation into JEDEC standards. Upon committee and board approval, new standards are promulgated and made available to the public.

101. In some cases, JEDEC members own patents covering the technology they or others seek to have included in a JEDEC standard. A patent that includes a claim or claims that would necessarily be infringed by the use, sale, offer for sale, or disposition of a portion of a product in order to be compliant with an industry standard is referred to as a "standard essential patent" or an "SEP."

102. Each owner of an SEP can, unless restrained, demand and obtain exorbitant royalties for the use of its patent, far in excess of the value, if any, that the patented technology

would command had it not been incorporated into a standard. If unwilling to pay such excessive royalties, firms wishing to implement the standard face the risk of being foreclosed from using any portion of the standard, including unpatented and public domain technologies. This threat of foreclosure, if left unchecked, puts a manufacturer's investment in developing standard-compliant products at risk and, in effect, permits the possibility that the SEP holder may capture up to the value of the standard as a whole for each unique implementer, rather than the true value of its specific contribution to the standard, independent of the incorporation of the technology into the standard. The exploitation of SEPs to extract unreasonable or discriminatory royalties is referred to as patent "hold-up."

103. The problem of "hold-up" is exacerbated in the context of standards, like those at issue here, as to which there are large numbers of SEPs and many different firms that claim to hold SEPs. The problem is likewise compounded when the products at issue, like those at issue here, may be required to implement multiple standards. The resulting problem of excessive aggregate royalties is referred to as "royalty stacking."

104. In order to mitigate these risks, JEDEC—like many other standard-setting organizations—has adopted a Patent Policy that seeks to ensure that owners of SEPs license those patents to manufacturers of standard-compliant products on RAND terms and conditions. Every firm that is a member of a JEDEC committee or a participant in a JEDEC committee meeting, including Netlist, agrees to abide by the JEDEC Patent Policy.

All Committee Members, as a condition of committee membership or committee Participation, agree to abide by JEDEC rules and procedures, including this JEDEC patent policy ("Patent Policy").

Ex. 4 (JEDEC Manual No. 21T) § 8.2.2.1.

105. Pursuant to the JEDEC Patent Policy, as a condition of committee membership or participation, all JEDEC committee members agree to disclose “Potentially Essential Patents” relevant to the work of the committee:

All Committee Members must Disclose Potentially Essential Patents, known to their Representative(s) to be Potentially Essential Patents that are owned or controlled by that Committee Member

Id. § 8.2.3.

106. The requirement to disclose “Potentially Essential Patents” extends to disclosing pending patent applications. *Id.* § 8.2.1.

107. JEDEC committee members also agree to license “Essential Patent Claims” on RAND terms and conditions:

All Committee Members, as a condition of committee membership or committee Participation, agree to Disclose Potentially Essential Patents, as set forth more fully in 8.2.3, and to offer to license their Essential Patent Claims to all Potential Licensees on RAND terms and conditions, if and as consistent with 8.2.3 and 8.2.4.

Id. § 8.2.2.1.

108. JEDEC committee members also agree to license on RAND terms any “Essential Patent Claims” included in their patent applications:

For any Essential Patent Claims held or controlled by the entity, pending[,] or anticipated to be filed, which are or may be required to implement a Standard that may result from the JEDEC Standard Activity, the entity hereby makes one of the following commitments: . . . A license will be offered . . . under reasonable terms and conditions that are free of any unfair discrimination.

Id. § 8.2.5 (emphasis added).

109. The JEDEC Patent Policy defines “Potentially Essential Patent” as a patent that is reasonably believed to contain one or more Essential Patent Claims. *Id.* § 8.2.1. “Essential Patent Claim” is further defined as those “Patent claims the use of which would necessarily be

infringed by the use, sale, offer for sale or other disposition of a portion of a product in order to be compliant with the required portions of a final approved JEDEC standard.” *Id.*

110. The JEDEC Patent Policy provides that the disclosure of Potentially Essential Patents “shall be made as early as reasonably possible” and documented by submission to JEDEC of either a “License Assurance/Disclosure Form” or a “Notice of Refusal.” *Id.* § 8.2.3.

111. A License Assurance/Disclosure Form (also known as a “Letter of Assurance” or “LOA”) states that the committee member commits to license its SEPs to all potential licensees on RAND terms. On the contrary, a Notice of Refusal states that the member is not “willing to offer to license Essential Patent Claims . . . on RAND terms to all Potential Licensees.” *Id.* § 8.2.3.1. In order to be effective, any such LOA or Notice of Refusal must be delivered to the JEDEC Legal Department within thirty (30) calendar days of a draft specification’s completion. *Id.*

112. Member companies who disclose and agree to license their patents on RAND terms do so via the aforementioned “LOA” identifying the relevant JEDEC standard(s) and making one of the two following commitments specified by the JEDEC Patent Policy:

For any Essential Patent Claims held or controlled by the entity, pending or anticipated to be filed, which are or may be required to implement a Standard that may result from the JEDEC Standard Activity, the entity hereby makes one of the following commitments:

- (i) A license will be offered, without compensation, under reasonable terms and conditions that are free of any unfair discrimination to applicants desiring to utilize the license for the purpose of implementing the JEDEC Standard, subject to the terms and conditions in 8.2.4; or*
- (ii) A license will be offered, to applicants desiring to utilize the license for the purpose of implementing the JEDEC Standard under reasonable terms and conditions that are free of any unfair discrimination, subject to the terms and conditions in 8.2.4.*

Id. § 8.2.5.

113. According to JEDEC's Patent Policy, disclosures and commitments with respect to one patent are deemed to include all patents claiming priority to the same filing. *Id.* § 8.2.1. Committee members are therefore obligated to license on RAND terms later-issuing SEPs that claim priority to patents disclosed to JEDEC.

114. If a committee member who owns or controls patents essential (or potentially essential) to JEDEC standards does not disclose and agree to license them before the ballot closes, or notify the committee of an intention not to license through a "Notice of Refusal," the committee member will be deemed to have agreed to offer licenses on RAND terms:

If a Committee Member, at its discretion, elects not to submit a License Assurance/Disclosure Form (see Annex A.3) at or before the time the ballot closes and does not otherwise provide notice of an unwillingness to license in accordance with 8.2.3.1, the Committee Member and its Affiliates will be deemed to have agreed to offer to grant licenses on RAND terms and conditions for all of its Essential Patent Claims of the balloted Standard, if and as consistent with 8.2.4.

Id. § 8.2.5.

115. The JEDEC Patent Policy provides that it is to be interpreted and governed under the laws of the State of New York. *Id.* § 8.2.10. Although the JEDEC Manual 21 has been updated and amended from time to time, upon information and belief, all of the relevant provisions of the JEDEC Patent Policy have been materially the same at all times relevant hereto.

2. Netlist's RAND Commitments

116. Netlist alleges that it owns over 100 patents and patent applications related to memory technologies.

117. Netlist is also a member of JEDEC and has joined a number of JEDEC technical committees over time. After years of membership in JEDEC, Netlist withdrew from

membership in JEDEC briefly in 2011 after submitting Notices of Refusal for a number of Netlist patents, including the '853 application. JEDEC reinstated Netlist later in 2011 following Netlist's submission of LOAs for the same patents Netlist had previously refused to license on RAND terms, including the '853 application.

118. Netlist again withdrew from membership in JEDEC committees JC-40, JC-42, and JC-45 in 2015, after submitting Notices of Refusal for many of its patents. In May 2018, Netlist requested from the JEDEC board of directors that Netlist be reinstated in those committees. The JEDEC board of directors approved the request on the condition that Netlist would agree to be bound by the JEDEC Patent Policy regarding work conducted in JC-40, JC-42, and JC-45 during the time Netlist had not participated in the committees, since and including February 27, 2015. Netlist agreed in writing to the JEDEC board of directors' terms, and JEDEC reinstated Netlist on August 14, 2018, retroactive to February 27, 2015.

119. Upon information and belief, therefore, at the time that the JEDEC Standards were drafted and approved by JEDEC, Netlist was a member, attended the meetings, and/or was deemed a member of the JC-40, JC-42, and JC-45 committees that developed those standards.

120. To date, and on information and belief, Netlist has disclosed that at least 42 of its patents are essential or potentially essential to JEDEC standards.

121. Netlist is obligated to license the '731 patent on RAND terms to implementers of certain DDR4 and DDR5 standards, to the extent the '731 patent is essential to those standards. In fact, Netlist committed to license the '731 patent on RAND terms to implementers of certain DDR4 and DDR5 standards. Upon information and belief, on November 22, 2010, and September 7, 2011, Netlist submitted Letters of Assurance for the '853 application,

related to JEDEC committees JC-40, JC-42, and JC-45. Ex. 5 (2010 Letter of Assurance); Ex. 6 (2011 Letter of Assurance). The '731 patent is a continuation of the patent issuing from the '853 application. Therefore, under the JEDEC Patent Policy, the commitments in the '853 application Letters of Assurance extend to the '731 patent. Ex. 4 (JEDEC Manual No. 21T) § 8.2.1.

122. Netlist is obligated to license the '366 patent on RAND terms to implementers of certain DDR4 and DDR5 standards, to the extent the '366 patent is essential to those standards. In fact, Netlist committed to license the '366 patent on RAND terms to implementers of certain DDR4 and DDR5 Standards. Upon information and belief, on August 7, 2018, Netlist submitted a Letter of Assurance for the '831 and '833 patents related to the JEDEC JC-40 committee. The '366 patent is a continuation of the '831 and '833 patents. Therefore, under the JEDEC Patent Policy, the commitments in the Letter of Assurance related to the '831 and '833 patents extends to the '366 patent. Ex. 4 (JEDEC Manual No. 21T) § 8.2.1.

123. Under JEDEC Patent Policy, Netlist is therefore obligated to license the '731 and '366 patents on RAND terms to implementers of certain DDR4 and DDR5 standards, to the extent the '731 and '366 patents are essential to those standards. By participating in JEDEC meetings and submitting a Letter of Assurance, Netlist committed to license the '731 and '366 patents on RAND terms to implementers of certain DDR4 and DDR5 standards.

124. The commitments that Netlist has made to license its SEPs on RAND terms constitute binding contractual obligations that may be enforced by firms seeking to implement the JEDEC standards, such as Samsung.

3. Netlist's Failure to Comply with Its RAND Obligations

125. Although Netlist accepted the benefits of its membership in JEDEC to induce JEDEC to incorporate technologies over which Netlist claims to have patents into JEDEC standards, Netlist has failed to fulfill its corresponding contractual commitments. In other words,

despite voluntarily undertaking the obligation to license its alleged SEPs on RAND terms, Netlist has violated its RAND commitment.

126. Netlist has never offered a RAND license to Samsung for the '731 patent. Instead, without notice to Samsung or any license offer, Netlist amended its complaint in the EDTX III Action to include the '731 patent. EDTX III, D.I. 14. This failure to offer Samsung a license on RAND terms before filing suit on an alleged standard-essential patent violates Netlist's commitment to license its alleged SEP on RAND terms and conditions.

127. Netlist has also never offered a RAND license to Samsung for the '366 patent. This failure to offer Samsung a license on RAND terms on an alleged standard-essential patent violates Netlist's commitment to license its alleged SEP on RAND terms and conditions.

128. Accordingly, and as discussed elsewhere herein, Netlist has breached its obligations under the JEDEC Patent Policy.

COUNT I
(Declaration of Non-Infringement of the '087 Patent)

129. Plaintiffs restate and incorporate by reference the preceding paragraphs as if fully set forth herein.

130. The Patent Office issued the '087 patent, titled "Memory Package Having Stacked Array Dies and Reduced Driver Load," on May 20, 2025. On information and belief, Netlist claims to own all rights, title, and interest in the '087 patent. A true and correct copy of the '087 patent is attached hereto as Exhibit 1.

131. The '087 patent has six independent claims: 1, 9, 13, 20, 23, and 24. For example, claim 1 reads as follows:

Element	Claim Language
preamble	1. A dynamic random access memory (DRAM) package, comprising:

Element	Claim Language
(a)	stacked DRAM dies including at least a first plurality of DRAM dies and a second plurality of DRAM dies, each DRAM die of the stacked DRAM dies including C/A ports, data ports and DRAM memory cells, wherein each DRAM die is configurable to transfer data between the data ports and the DRAM memory cells;
(b)	terminals including command and/or address (C/A) terminals and data terminals, wherein the DRAM package is configured to receive C/A signals via the C/A terminals and is further configured to receive or output data signals via the data terminals in response to the (C/A) signals, wherein the DRAM package is configured to output first data signals in response to a first set of C/A signals associated with a memory read operation and to receive second data signals in response to a second set of C/A signals associated with a memory write operation;
(c)	die interconnects including C/A interconnects and data interconnects, the C/A interconnects including at least first C/A interconnects and second C/A interconnects, the first C/A interconnects configured to conduct the first set of C/A signals and the second set of C/A signals, the data interconnects including at least first data interconnects and second data interconnects, the first data interconnects configured to conduct the first data signals and the second data signals, each of the die interconnects including one or more through silicon vias (TSVs) in one or more DRAM dies in the stacked DRAM dies and configured to conduct signals to and/or from the one or more DRAM dies in the stacked DRAM dies through the one or more TSVs;
(d)	a control die coupled between the terminals and the stacked DRAM dies, the control die including conduits including C/A conduits and data conduits, the C/A conduits including at least first C/A conduits coupled to respective ones of the first C/A interconnects and second C/A conduits coupled to respective ones of the second C/A interconnects, the data conduits including at least first data conduits coupled to respective ones of the first data interconnects and second data conduits coupled to respective ones of the second data interconnects;
(e)	wherein a first C/A interconnect of the first C/A interconnects is in electrical communication with corresponding C/A ports on the first plurality of DRAM dies and not in electrical communication with any C/A port on any of the second plurality of DRAM dies;
(f)	wherein a second C/A interconnect of the second C/A interconnects is in electrical communication with corresponding C/A ports on the second plurality of DRAM dies and not in electrical communication with any C/A port on any of the first plurality of DRAM dies;
(g)	wherein a first data interconnect of the first data interconnects is in electrical communication with corresponding data ports on the first plurality of DRAM dies and not in electrical communication with any data port on any of the second plurality of DRAM dies, each of the first data interconnects including a first respective set of TSVs, the first respective

Element	Claim Language
	set of TSVs including a TSV in each DRAM die of the first plurality of DRAM dies and at least one TSV in at least one DRAM die of the second plurality of DRAM dies, wherein the TSV in the each DRAM die of the first plurality of DRAM dies is in electrical communication with a corresponding data port on the each DRAM die, and wherein the at least one TSV in the at least one DRAM die of the second plurality of DRAM dies is not in electrical communication with any data port on the at least one DRAM die;
(h)	wherein a second data interconnect of the second data interconnects is in electrical communication with corresponding data ports on the second plurality of DRAM dies and not in electrical communication with any data port on any of the first plurality of DRAM dies;
(i)	wherein a first conduit of the first data conduits is coupled between the first data interconnect and a first data terminal of the data terminals, and a second conduit of the second data conduits is coupled between the second data interconnect and the first data terminal;
(j)	wherein the control die further includes control logic configurable to control respective states of the first and second conduits in response to one or more C/A signals received via one or more of the C/A terminals, wherein the one or more C/A signals do not include any chip select signal;
(k)	wherein the die interconnects further include first unidirectional interconnects configured to conduct signals from one or more DRAM dies of the stacked DRAM dies to the control die and not configured to conduct any signal from the control die to any of the stacked DRAM dies;
(l)	wherein the die interconnects further include second unidirectional interconnects configured to conduct signals from the control die to one or more DRAM dies of the stacked DRAM dies and not configured conduct any signal from any of the stacked DRAM dies to the control die;
(m)	wherein the control die is configured to receive signals from one or more DRAM dies of the stacked DRAM dies via the first unidirectional interconnects and is not configured to drive any signal to any of the stacked DRAM dies via any of the first unidirectional interconnects;
(n)	wherein the control die is configured to drive signals to one or more DRAM dies of the stacked DRAM dies via the second unidirectional interconnects and is not configured to receive any signal from any of the stacked DRAM dies via any of the second unidirectional interconnects; and
(o)	wherein the control die is configured to, in response to the first set of C/A signals, receive first signals associated with the memory read operation from a DRAM die of the stacked DRAM dies via the first unidirectional interconnects, and in response to the second set of C/A signals, drive second signals associated with the memory write operation to one or more DRAM dies of the stacked DRAM dies via the second unidirectional interconnects.

132. Plaintiffs have not directly or indirectly infringed any claim of the '087 patent, either literally or under the doctrine of equivalents, at least because the Samsung HBM Products do not employ, incorporate, or otherwise make use of all of the limitations of the claims of the '087 patent.

133. For example, claims 1, 9, 13, 20, 23, and 24 require “unidirectional interconnects” (claims 1, 9, 13, 20, and 23) or “unidirectional conduit[s]” (claim 24). Samsung HBM Products do not include “unidirectional interconnects” or “unidirectional conduits.” Claim limitations 1(k)-(o) provide additional requirements for the “unidirectional interconnect,” including, among other things, that “first unidirectional interconnects [are] configured to conduct signals from one or more DRAM dies of the stacked DRAM dies to the control die and not configured to conduct any signal from the control die to any of the stacked DRAM dies” and “second unidirectional interconnects [are] configured to conduct signals from the control die to one or more DRAM dies of the stacked DRAM dies and not configured conduct any signal from any of the stacked DRAM dies to the control die.” The Samsung HBM products do not include through-silicon vias (“TSVs”) that satisfy all claim limitations related to “unidirectional interconnect.” [REDACTED]

[REDACTED] This means the Samsung HBM products do not include at least the “second unidirectional interconnects,” which may only “conduct signals from the control die to the one or more DRAM dies” and may not “conduct any signal from any of the stacked DRAM dies to the control die.” Thus, the Samsung HBM Products do not meet all of the limitations of any claims of the '087 patent.

134. As another example, claim 1 requires that the “control die . . . is configured to receive signals from one or more DRAM dies of the stacked DRAM dies via the first

unidirectional interconnects and is not configured to drive any signal to any of the stacked DRAM dies via any of the first unidirectional interconnects” and “configured to drive signals to one or more DRAM dies of the stacked DRAM dies via the second unidirectional interconnects and is not configured to receive any signal from any of the stacked DRAM dies via any of the second unidirectional interconnects.” Independent claims 13, 20, and 23 include similar limitations directed to a “control die.” The Samsung HBM Products do not include a “control die” that meets the requirements of claims 1, 13, 20, and 23. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] For this additional reason, the Samsung HBM Products do not meet the limitations of at least claims 1–8 and 13–23.

135. The Samsung HBM Products do not infringe any of claims 2–8, 10–12, 14–19, 12–22, and 25–28 at least because these claims depend directly or indirectly from one of claims 1, 9, 13, 20, 23, and 24.

136. Plaintiffs, their customers, and their end users are not liable for infringement of the ’087 patent for any Samsung HBM Products made, imported, or sold under license from Netlist.

137. A substantial, immediate, and real controversy exists between Plaintiffs and Netlist regarding whether Plaintiffs or their customers or end users infringe the ’087 patent by making, using, selling, and/or offering for sale the Samsung HBM Products in the United States, or by importing the Samsung HBM Products into the United States. A judicial declaration is necessary to determine the parties’ respective rights regarding the ’087 patent.

138. Plaintiffs seek a judgment declaring that Plaintiffs and their customers and end users do not infringe the '087 patent, either literally or under the doctrine of equivalents, by making, using, selling, and/or offering for sale the Samsung HBM Products in the United States, or by importing the Samsung HBM Products into the United States, either directly under 35 U.S.C. § 271(a) or indirectly under 35 U.S.C. §§ 271(b)–(c).

COUNT II
(Declaration of Non-Infringement of the '731 Patent)

139. Plaintiffs restate and incorporate by reference the preceding paragraphs as if fully set forth herein.

140. The Patent Office issued the '731 patent, titled “Memory Module and Circuit Providing Load Isolation and Noise Reduction,” on July 17, 2018. Netlist claims to own all rights, title, and interests in the '731 patent. EDTX III, D.I. 14, ¶ 23. A true and correct copy of the '731 patent is attached hereto as Exhibit 2.

141. The '731 patent has one independent claim, claim 1, which reads as follows:

Element	Claim Language
preamble	1. A memory module operable to communicate data with a system memory controller via a memory bus in response to address and control signals from the system memory controller, comprising:
(a)	a printed circuit board comprising at least one connector configured to be operatively coupled to the memory bus, the at least one connector including a plurality of edge connections distributed along one or more edges of the printed circuit board;
(b)	a plurality of memory devices on the printed circuit board, the plurality of memory devices being arranged in multiple ranks, each rank comprising an independent set of memory devices that can be accessed by the system memory controller to access a full bit-width of the memory bus;
(c)	at least one circuit coupled between the at least one connector and the plurality of memory devices, the at least one circuit comprising a first set of port connections coupled to respective ones of the plurality of edge connections and a second set of port connections coupled to the plurality

Element	Claim Language
	of memory devices, each circuit of the at least one circuit further comprising a set of correction circuits, each correction circuit of the set of correction circuits being configured to make corrections in one or more signals transmitted between a respective port connection of the first set of port connections and a corresponding port connection in the second set of port connections, the each correction circuit of the set of correction circuits including at least one programmable impedance matching circuit; and
(d)	control circuitry configured to receive the address and control signals from the system memory controller and to control the plurality of memory devices in response to the address and control signals, wherein the control circuitry is further configured to dynamically control the at least one programmable impedance matching circuit in the each correction circuit of the set of correction circuits in the each circuit of the at least one circuit based on which of the multiple ranks is selected to communicate data with the system memory controller in response to the address and control signals.

142. Plaintiffs have not directly or indirectly infringed any claim of the '731 patent, either literally or under the doctrine of equivalents, at least because the Samsung DDR4 and DDR5 Products do not employ, incorporate, or otherwise make use of all of the limitations of the claims of the '731 patent.

143. For example, claim 1(d) requires “control circuitry [that] is further configured to dynamically control the at least one programmable impedance matching circuit in the each correction circuit of the set of correction circuits in the each circuit of the at least one circuit based on which of the multiple ranks is selected to communicate data with the system memory controller in response to the address and control signals.” Samsung DDR4 and DDR5 Products do not include circuitry that is configured as claimed. The accused “control circuitry” does not control the accused “programmable impedance matching circuit” or the accused “correction circuits.” Rather, the accused features are controlled by the system memory controller.

Thus, the Samsung DDR4 and DDR5 Products do not meet all of the limitations of any claims of the '713 patent.

144. The Samsung DDR4 and DDR5 Products do not infringe any of claims 2–18 at least because these claims depend directly or indirectly from claim 1.

145. Plaintiffs, their customers, and their end users are not liable for infringement of the '731 patent for any Samsung DDR4 and DDR5 Products made, imported, or sold under license from Netlist.

146. A substantial, immediate, and real controversy exists between Plaintiffs and Netlist regarding whether Plaintiffs or their customers or end users infringe the '731 patent by making, using, selling, and/or offering for sale the Samsung DDR4 and DDR5 Products in the United States, or by importing the Samsung DDR4 and DDR5 Products into the United States. A judicial declaration is necessary to determine the parties' respective rights regarding the '731 patent.

147. Plaintiffs seek a judgement declaring that Plaintiffs and their customers and end users do not infringe the '731 patent, either literally or under the doctrine of equivalents, by making, using, selling, and/or offering for sale the Samsung DDR4 and DDR5 Products in the United States, or by importing the Samsung DDR4 and DDR5 Products into the United States, either directly under 35 U.S.C. § 271(a) or indirectly under 35 U.S.C. §§ 271(b)–(c).

COUNT III
(Declaration of Non-Infringement of the '366 Patent)

148. Plaintiffs restate and incorporate by reference the preceding paragraphs as if fully set forth herein.

149. The Patent Office issued the '366 patent, titled “Memory with On-Module Power Management,” on July 29, 2025. On information and belief, Netlist claims to own all rights,

title, and interest in the '366 patent. A true and correct copy of the '366 patent is attached hereto as Exhibit 3.

150. The '366 patent has three independent claims, 1, 16, and 29. For example, claim 1 reads as follows:

Element	Claim Language
preamble	1. A dual in-line memory module (DIMM) configured to fit into a memory slot connector of a system board of a computer system, the memory slot connector including a set of power conduits, a set of data conduits, and a set of address and control conduits, wherein data signals are transferred between the system board and the DIMM via the set of data conduits, and address and control signals are delivered from the system board to the DIMM via the set of address and control conduits, the DIMM comprising:
(a)	a printed circuit board (PCB) having edge connections configured to fit into the memory slot connector of the system board, the PCB including:
(b)	a first set of edge connections configured to be electrically connected to the set of power conduits and operable to deliver to the DIMM power from power rails of the system board via the set of power conduits, wherein the power delivered to the DIMM from the power rails of the system board is the only power received by the DIMM;
(c)	a second set of edge connections configured to be electrically connected to the set of data conduits, and
(d)	a third set of edge connections configured to be electrically connected to the set of address and control conduits,
(e)	a set of voltage supply lines electrically connected to the first set of edge connections and configured to receive the power delivered to the DIMM from the power rails of the system board by way of the memory slot connector of the system board and the first set of edge connections of the DIMM,
(f)	a set of data lines electrically connected to the second set of edge connections, and
(g)	a set of address and control lines electrically connected to the third set of edge connections;

Element	Claim Language
(h)	a controller including a voltage monitor circuit and nonvolatile memory, the controller coupled to the PCB and to an input voltage supply line of the set of voltage supply lines, the voltage monitor circuit configured to: (i) monitor an input voltage supplied via the input voltage supply line; (ii) generate a trigger signal upon detecting a trigger condition; and (iii) transmit the trigger signal to at least one other portion of the controller, wherein the controller is configured to perform, in response to the trigger signal, a write operation to write data into the nonvolatile memory, and wherein the trigger condition occurs when the input voltage exceeds a first threshold voltage;
(i)	a set of components coupled to the PCB, the set of components including a plurality of double data rate (DDR) synchronous dynamic random access memory (SDRAM) devices coupled to the set of address and control lines and to the set of data lines, the DDR SD RAM devices being operable to receive or output data signals via the set of data lines based on address and control signals received from the system board via the set of address and control lines; and
(j)	first, second, third, and fourth converter circuits coupled to the PCB and to the input voltage supply line, and configured to receive power from the input voltage supply line and to deliver power via first, second, third, and fourth regulated voltage lines,
(k)	wherein the set of components includes: a first component configured to receive power via at least the first regulated voltage line; a second component configured to receive power via at least the second regulated voltage line; a third component configured to receive power via at least the third regulated voltage line; and a fourth component configured to receive power via at least the fourth regulated voltage line;
(l)	wherein each component of the set of components: (i) is electrically connected to one or more of the first, second, third, and fourth regulated voltage lines; and (ii) receives power only via the one or more of the first, second, third, and fourth regulated voltage lines,
(m)	wherein the plurality of DDR SDRAM devices includes: a first group of at least five DDR SDRAM devices that are each connected to a first chip select line configured to electrically conduct a first chip select signal; a second group of at least four DDR SDRAM devices that are each connected to a second chip select line configured to electrically conduct a second chip select signal;

Element	Claim Language
(n)	wherein when power from the power rails of the system board is provided to the DIMM and each group of the first and second groups of DDR SD RAM devices is not in self-refresh mode: the first group of DDR SD RAM devices is configured to be enabled, in response to the first chip select signal, to receive or output at least 40 1-bit DDR data signals in parallel via at least 40 data lines of the set of data lines independently of whether the second group of DDR SDRAM devices is enabled, in response to the second chip select signal, to receive or output an integer number of 1-bit DDR data signals in parallel via an integer number of data lines of the set of data lines; the at least 40 1-bit DDR data signals include at least 32 1-bit wide DDR data signals and at least eight 1-bit DDR error-correcting code (ECC) data signals; and
(o)	wherein the sum of the at least 40 data lines and the integer number of data lines equals a total number of data conduits of the set of data conduits of the memory slot connector, and the total number of data conduits is at least 72 data conduits.

151. Plaintiffs have not directly or indirectly infringed any claim of the '366 patent, either literally or under the doctrine of equivalents, at least because the Samsung DDR4 and DDR5 Products do not employ, incorporate, or otherwise make use of all of the limitations of the claim of the '366 patent.

152. For example, claims 1, 16, and 29 require “DDR SDRAM devices” (claims 1 and 29) or “DDR SDRAM packages” (claim 16) configured “to receive or output . . . 1-bit DDR data signals in parallel. . .” In the Samsung DDR5 Products, the accused “DDR SRAM devices” or “DDR SDRAM packages” are not configured to receive or output “1-bit DDR data signals.” Rather, the accused DDR SDRAM devices and packages are configured to receive or output serial DDR data signals.

153. As another example, claims 1, 16, and 29 requires a “first, second, third, and fourth converter circuits [that are] coupled to the PCB and to the input voltage supply line, and [are] configured to receive power from the input voltage supply line and to deliver power via

first, second, third, and fourth regulated voltage lines.” Netlist alleges that the claimed converters are in a power management IC, or PMIC. The Samsung DDR4 Products do not include a power management IC.

154. The Samsung DDR4 and DDR5 Products do not infringe any of claims 2–15, 17–28, and 30–38 at least because these claims depend directly or indirectly from claims 1, 16, and 29.

155. Plaintiffs, their customers, and their end users are not liable for infringement of the ’366 patent for any Samsung DDR4 and DDR5 Products made, imported, or sold under license from Netlist.

156. A substantial, immediate, and real controversy exists between Plaintiffs and Netlist regarding whether Plaintiffs or their customers or end users infringe the ’366 patent by making, using, selling, and/or offering for sale the Samsung DDR4 and DDR5 Products in the United States, or by importing the Samsung DDR4 and DDR5 Products into the United States. A judicial declaration is necessary to determine the parties’ respective rights regarding the ’366 patent.

157. Plaintiffs seek a judgment declaring that Plaintiffs and their customers and end users do not infringe the ’366 patent, either literally or under the doctrine of equivalents, by making, using, selling, and/or offering for sale the Samsung DDR4 and DDR5 Products in the United States, or by importing the Samsung DDR4 and DDR5 Products into the United States, either directly under 35 U.S.C. § 271(a) or indirectly under 35 U.S.C. § 271(b)–(c).

COUNT IV
(Alternative Claim for Breach of Contract)

158. Samsung restates and incorporates by reference the preceding paragraphs as if fully set forth herein.

159. If the JEDEC-compliant Samsung DDR4 and DDR5 memory modules infringe any claim of the '731 or '366 patents, then, in the alternative, the '731 and/or '366 patents are standard-essential patents, and Netlist has breached its contractual obligations to license any such patent claim to Samsung on RAND terms and conditions.

160. Under the JEDEC Patent Policy, Netlist has a contractual commitment to offer implementers of the JEDEC Standards, including Samsung, licenses to any Essential Patent Claims (as defined in the JEDEC Patent Policy) on RAND terms and conditions. This contractual commitment is on-going. Netlist continues to have an obligation to license its SEPs to Samsung on RAND terms, notwithstanding any breach of the Agreement or any termination of the license contained therein. Having entered the Agreement, Samsung showed that it is a willing licensee.

161. Regarding the '731 patent, Netlist submitted Letters of Assurance to the JC-40, JC-42, and JC-45 committees in which Netlist promised to make available to implementers of certain DDR4 and DDR5 standards a license to the '853 application "under reasonable terms and conditions that are demonstrably free of any unfair discrimination." Ex. 6 (2011 Letter of Assurance). Under the JEDEC Patent Policy, those commitments extend to the claims of the '731 patent, to the extent they are essential to the DDR4 and DDR5 standards.

162. Regarding the '366 patent, Netlist submitted a Letter of Assurance to the JC-40 committee in which Netlist promised to make available to implementers of certain DDR4 and DDR5 standards a license to the '831 and '833 patents "under reasonable terms and conditions that are demonstrably free of any unfair discrimination." Under the JEDEC Patent Policy, those commitments extend to the claims of the '366 patent, to the extent they are essential to the DDR4 and DDR5 Standards.

163. The JEDEC Patent Policy and/or Letters of Assurance constitute a valid and binding contract between Netlist and JEDEC.

164. Samsung is a third-party beneficiary to the JEDEC Patent Policy and/or Letters of Assurance, and to Netlist's RAND obligations thereunder. Participation in JEDEC committees and the standard-setting process is expressly conditioned on a commitment to license essential patents to third-party implementers on RAND terms. Ex. 4 (JEDEC Manual No. 21T) § 8.2.2.1. The JEDEC Patent Policy is therefore intended to benefit third-party implementers of JEDEC standards, such as Samsung, and the obligations under the JEDEC Patent Policy, including the RAND obligations, are intended to be enforceable by third-party implementers of the JEDEC standards. Third-party implementers, such as Samsung, are also the only parties who could recover for Netlist's breach of its obligations under the JEDEC Patent Policy.

165. Samsung has relied on the JEDEC Patent Policy and/or Letters of Assurance, including the RAND obligations set forth therein, in designing, manufacturing, and selling standard-compliant products, including the Samsung DDR4 and DDR5 memory modules, and Samsung has supported the JEDEC standard based on its understanding and expectation that JEDEC members, including Netlist, will abide by their obligations.

166. In ongoing litigation, Netlist and Samsung dispute whether Samsung has a license to Netlist's patents, including the '731 and '366 patents, under the Agreement. Netlist has taken the position that Samsung's license under the Agreement has been terminated, while Samsung maintains that the termination was not effective and that it has a license to the '731 and '366 patents. As set forth above, Netlist has demanded that Samsung enter a second license to patents that Samsung previously licensed from Netlist, specifically demanding that because of the

alleged JDLA termination “Samsung will require a new license to our patent portfolio,” which would include the ’731 and ’366 patents.

167. Although Netlist accepted the benefits of its membership in JEDEC, Netlist has failed to fulfill its corresponding contractual commitments. In particular, since the date of Netlist’s purported termination of the JDLA to the present, Netlist has failed to offer Samsung a license to patents that Netlist alleges are essential to the JEDEC Standards, including the ’731 and ’366 patents, on RAND terms and conditions. Filing an amended complaint in EDTX III on the ’731 without offering Samsung a RAND license to its alleged SEP, and failing to offer a RAND license to Samsung on the ’366 patent, violated Netlist’s RAND commitment. Thus, to the extent the ’731 or ’366 patent is essential to any of the JEDEC Standards, Netlist has breached its contractual obligations to license such patent to Samsung on RAND terms and conditions.

168. As a result of Netlist’s breaches of contract, Samsung has been injured in its business or property because it has been denied access to a license to Netlist’s claimed SEP on RAND terms, a license that can only be obtained from Netlist. Netlist’s conduct threatens Samsung with imminent loss of customers and potential customers, loss of goodwill and product image, loss of sales, litigation costs, uncertainty in business planning, and uncertainty among customers and potential customers.

JURY DEMAND

Samsung demands a jury trial on all issues and claims so triable.

PRAYER FOR RELIEF

WHEREFORE, Samsung prays for judgment and relief as follows:

(a) Declare that Plaintiffs do not directly or indirectly infringe the ’087 patent, either literally or under the doctrine of equivalents, and that they are not liable for damages or injunctive relief based on any claim in the ’087 patent;

(b) Declare that Plaintiffs do not directly or indirectly infringe the '731 patent, either literally or under the doctrine of equivalents, and that they are not liable for damages or injunctive relief based on any claim in the '731 patent;

(c) Declare that Plaintiffs do not directly or indirectly infringe the '366 patent, either literally or under the doctrine of equivalents, and that they are not liable for damages or injunctive relief based on any claim in the '366 patent;

(d) In the alternative, if the '731 or '366 patent is found to be infringed by Samsung, declare that Netlist is liable for breach of contract; and

- (1) Grant injunctive relief requiring that any Netlist demand that Samsung take a license to the '731 and/or '366 patent—to the extent the patent is found to be essential to the JEDEC standards—must be on reasonable terms and conditions that are demonstrably free from any unfair discrimination;
- (2) Enjoin Netlist from further demanding excessive, non-RAND royalties from Samsung for alleged infringement of the '731 and/or '366 patent;
- (3) Declare that Netlist is barred from seeking and/or enforcing injunctive relief against Samsung (including its affiliates) or its direct or indirect customers and end-users in any jurisdiction with respect to any alleged infringement of the '731 and/or '366 patent;
- (4) Enjoin Netlist from seeking and/or enforcing injunctive relief against Samsung (including its affiliates) or its direct or indirect

customers and end users in any jurisdiction with respect to any alleged infringement of the '731 and/or '366 patent;

(5) Compensate Samsung for all damages caused by Netlist's breaches of contract, including breaches of its RAND obligations;

(e) Declare that judgment be entered in favor of Plaintiffs and against Netlist on Plaintiffs' claims;

(f) Find that this is an exceptional case under 35 U.S.C. § 285;

(g) Award Plaintiffs pre-judgment and post-judgment interest;

(h) Award Plaintiffs their costs and attorneys' fees in connection with this action; and

(i) Such further and additional relief as the Court deems just and proper.

MORRIS, NICHOLS, ARSHT & TUNNELL LLP

/s/ Rodger D. Smith II

OF COUNSEL:

Brian Nester
Peter Swanson
COVINGTON & BURLING LLP
One CityCenter
850 Tenth Street, NW
Washington, DC 20001-4956
(202) 662-6000

Alice J. Ahn
COVINGTON & BURLING LLP
Salesforce Tower
415 Mission Street, Suite 5400
San Francisco, CA 94105-2533
(415) 591-6000

Rodger D. Smith II (#3778)
Anthony D. Raucci (#5948)
Ben Yenerall (#7132)
1201 North Market Street
P.O. Box 1347
Wilmington, DE 19899
(302) 658-9200
rsmith@morrisnichols.com
araucci@morrisnichols.com
byenerall@morrisnichols.com

Attorneys for Plaintiffs

July 29, 2025

EXHIBIT 1



US012308087B2

(12) **United States Patent**
Lee

(10) **Patent No.:** **US 12,308,087 B2**

(45) **Date of Patent:** ***May 20, 2025**

(54) **MEMORY PACKAGE HAVING STACKED
ARRAY DIES AND REDUCED DRIVER
LOAD**

(56) **References Cited**

U.S. PATENT DOCUMENTS

(71) Applicant: **Netlist, Inc.**, Irvine, CA (US)

5,345,412 A 9/1994 Shiratsuchi
5,398,210 A 3/1995 Higuchi

(Continued)

(72) Inventor: **Hyun Lee**, Ladera Ranch, CA (US)

(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.

CN 102576565 9/2015
EP 1816570 A2 8/2007

(Continued)

This patent is subject to a terminal disclaimer.

OTHER PUBLICATIONS

(21) Appl. No.: **17/694,649**

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'First Amended Complaint,' filed May 3, 2022, 71 pages.

(22) Filed: **Mar. 14, 2022**

(Continued)

(65) **Prior Publication Data**

US 2022/0208233 A1 Jun. 30, 2022

Primary Examiner — Huan Hoang

(74) Attorney, Agent, or Firm — USCH Law, PC

Related U.S. Application Data

(63) Continuation of application No. 17/157,903, filed on Jan. 25, 2021, now abandoned, which is a
(Continued)

(51) **Int. Cl.**

G11C 5/06 (2006.01)

G11C 7/10 (2006.01)

G11C 7/12 (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); **G11C 7/12** (2013.01)

(58) **Field of Classification Search**

CPC G11C 5/066; G11C 5/06; G11C 7/1057; G11C 7/1084; G11C 7/12


(Continued)

(57)

ABSTRACT

A DRAM packages comprises stacked array dies including at least a first array die and a stacked over at least the first array die, data terminals, via which the DRAM package receives or outputs data signals, first data interconnects between respective ones of the data terminals and the first array die, and second data interconnects distinct from the first data interconnects and between respective ones of the data terminals and the second array die. The DRAM package further comprises drivers configured to drive first data signals to the first array die via the first data interconnects and second data signals to the second array die via the second data interconnects. A first data signal is driven by one or more drivers having a first driver size, and a second data signal is driven by one or more drivers having a second driver size different from the first driver size.

28 Claims, 8 Drawing Sheets



US 12,308,087 B2

Page 2

Related U.S. Application Data

continuation of application No. 16/412,308, filed on May 14, 2019, now Pat. No. 10,902,886, which is a continuation of application No. 15/602,099, filed on May 22, 2017, now Pat. No. 10,290,328, which is a continuation of application No. 15/095,288, filed on Apr. 11, 2016, now Pat. No. 9,659,601, which is a continuation of application No. 14/337,168, filed on Jul. 21, 2014, now Pat. No. 9,318,160, which is a 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.
- (58) **Field of Classification Search**
USPC 365/63
See application file for complete search history.

- (56) **References Cited**

U.S. PATENT DOCUMENTS

5,537,584	A	7/1996	Miyai	7,098,541	B2	8/2006	Adelmann
5,563,838	A	10/1996	Mart et al.	7,102,905	B2	9/2006	Funaba
5,617,559	A	4/1997	Le	7,123,497	B2	10/2006	Matsui
5,649,159	A	7/1997	Le	7,130,308	B2	10/2006	Haddock
5,655,113	A	8/1997	Leung	7,133,960	B1	11/2006	Thompson
5,717,851	A	2/1998	Vishay	7,161,820	B2	1/2007	Funaba et al.
5,724,604	A	3/1998	Moyer	7,166,934	B2	1/2007	Dewey
5,729,716	A	5/1998	Lee	7,177,187	B2	2/2007	Ishii
5,784,705	A	7/1998	Leung	7,200,021	B2	4/2007	Raghuram
5,802,541	A	9/1998	Reed	7,209,376	B2	4/2007	Saito
5,905,401	A	5/1999	Sher	7,215,561	B2	5/2007	Park
5,973,392	A	10/1999	Senba	7,221,617	B2	5/2007	Flach et al.
5,991,192	A	11/1999	Wang et al.	7,242,635	B2	7/2007	Okuda
6,011,710	A	1/2000	Wiggers	7,254,036	B2	8/2007	Pauley
6,031,762	A	2/2000	Saitoh	7,269,042	B2	9/2007	Kinsley
6,070,217	A	5/2000	Connolly	7,269,764	B2	9/2007	Dart et al.
6,109,929	A	8/2000	Jasper	7,280,417	B2	10/2007	Choi et al.
6,173,357	B1	1/2001	Ju	7,286,436	B2	10/2007	Bhakta
6,243,283	B1	6/2001	Bertin	7,289,386	B2	10/2007	Bhakta
6,260,127	B1	7/2001	Olarig	7,334,150	B2	2/2008	Ruckerbauer
6,414,904	B2	1/2002	So et al.	7,375,970	B2	5/2008	Pauley
6,351,827	B1	2/2002	Co et al.	7,379,316	B2	5/2008	Rajan
6,381,140	B1	4/2002	Liao	7,379,361	B2	5/2008	Co
6,418,065	B2	7/2002	Sato et al.	7,426,649	B2	9/2008	Brittain et al.
6,434,044	B1	8/2002	Gongwer et al.	7,442,050	B1	10/2008	Bhakta
6,446,158	B1	9/2002	Karabatsos	7,460,418	B2	12/2008	Jung
6,466,496	B2	10/2002	Kuge	7,464,225	B2	12/2008	Tsern
6,551,857	B2	4/2003	Leedy	7,466,577	B2	12/2008	Sekiguchi
6,553,449	B1	4/2003	Dodd	7,495,943	B2	2/2009	Takemura et al.
6,553,450	B1	4/2003	Dodd	7,532,537	B2	5/2009	Solomon
6,556,475	B2	4/2003	Yamazaki et al.	7,558,096	B2	7/2009	Ikeda
6,594,167	B1	7/2003	Yamasaki	7,587,559	B2	9/2009	Brittain et al.
6,611,459	B2	8/2003	Shiromoto et al.	7,613,880	B2	11/2009	Miura et al.
6,618,791	B1	9/2003	Dodd	7,619,893	B1	11/2009	Yu
6,621,155	B1	9/2003	Perino et al.	7,619,912	B2	11/2009	Bhakta
6,656,767	B1	12/2003	King	7,630,202	B2	12/2009	Pauley
6,664,118	B2	12/2003	Nishihara et al.	7,633,165	B2	12/2009	Hsu
6,677,671	B2	1/2004	King	7,633,785	B2	12/2009	Kim
6,683,372	B1	1/2004	Wong	7,636,274	B2	12/2009	Solomon
6,704,910	B2	3/2004	Hong	7,644,216	B2	1/2010	Fahr et al.
6,717,855	B2	4/2004	Underwood	7,656,735	B2	2/2010	Kagan et al.
6,721,860	B2	4/2004	Klein	7,679,967	B2	3/2010	Chung et al.
6,747,887	B2	6/2004	Halbert	7,683,459	B2	3/2010	Ma
6,788,592	B2	9/2004	Nakata	7,760,533	B2	7/2010	Alzheimer
6,832,303	B2	12/2004	Tanaka	7,774,535	B2	8/2010	Nakamura
6,889,335	B2	5/2005	Hargis et al.	7,796,446	B2	9/2010	Ruckerbauer
6,948,084	B1	9/2005	Manapat et al.	7,811,097	B1	10/2010	Bhakta
7,019,553	B2	3/2006	Blodgett et al.	7,827,348	B2	11/2010	Lee
7,024,518	B2	4/2006	Halbert	7,830,692	B2	11/2010	Chung
7,034,569	B1	4/2006	Balasubramanian et al.	7,830,706	B2	11/2010	Hanzawa et al.
7,046,538	B2	5/2006	Thomas	7,834,440	B2	11/2010	Ito
7,078,793	B2	7/2006	Ruckerbauer et al.	7,839,645	B2	11/2010	Pauley
7,093,066	B2	8/2006	Klein	7,864,627	B2	1/2011	Bhakta
				7,865,674	B2	1/2011	Gower
				7,881,150	B2	2/2011	Solomon
				7,890,811	B2	2/2011	Rothman
				7,894,229	B2	2/2011	Lahtinen
				7,894,230	B2	2/2011	Kim
				7,907,434	B2	3/2011	Matsuzaki et al.
				7,969,192	B2	6/2011	Wyman
				7,978,721	B2	7/2011	Jeddeloh
				7,990,171	B2	8/2011	Chung
				7,990,746	B2	8/2011	Rajan
				7,999,367	B2	8/2011	Kang
				8,001,434	B1	8/2011	Lee et al.
				8,019,589	B2	9/2011	Rajan
				8,033,836	B1	10/2011	Bhakta
				8,041,881	B2	10/2011	Rajan
				8,064,222	B2	11/2011	Nishio
				8,081,536	B1	12/2011	Solomon
				8,089,795	B2	1/2012	Rajan
				8,094,504	B2	1/2012	Smolka
				8,120,958	B2	2/2012	Bilger
				8,130,560	B1	3/2012	Rajan
				8,174,115	B2	5/2012	Chung
				8,179,737	B2	5/2012	Byeon
				8,189,328	B2	5/2012	Kanapathippillai
				8,233,303	B2	7/2012	Best
				8,258,619	B2	9/2012	Foster, Sr.

US 12,308,087 B2

Page 3

(56) References Cited

U.S. PATENT DOCUMENTS

8,259,461 B2 9/2012 Hollis
 8,281,074 B2 10/2012 Jeddeloh
 8,295,070 B2 10/2012 Fukano
 8,310,841 B2 11/2012 Foster
 8,315,068 B2 11/2012 Foster
 8,325,539 B2 12/2012 Park
 8,341,336 B2 12/2012 Chang et al.
 8,352,896 B2 1/2013 Thayer
 8,421,237 B2 1/2013 Chia et al.
 8,432,027 B2 1/2013 Foster et al.
 8,369,122 B2 2/2013 Byeon
 8,384,417 B2 2/2013 Laisne et al.
 8,400,781 B2 3/2013 Gillingham
 8,411,478 B2 4/2013 Yun
 8,417,870 B2 4/2013 Lee
 8,437,163 B2 5/2013 Nakanishi et al.
 8,471,362 B2 6/2013 Lee
 8,473,653 B2 6/2013 Kondo et al.
 8,476,767 B2 7/2013 Lee
 8,488,399 B2 7/2013 Yu et al.
 8,516,185 B2 8/2013 Lee
 8,516,409 B2 8/2013 Coteus
 8,566,556 B2 10/2013 Rajan
 8,582,373 B2 11/2013 Hollis
 8,654,593 B2 2/2014 Oh et al.
 8,659,136 B2 2/2014 Youn
 8,681,546 B2 3/2014 Fai et al.
 8,689,064 B1 4/2014 Lee
 8,743,582 B2 6/2014 Kang et al.
 8,756,364 B1 6/2014 Bhakta
 8,782,350 B2 7/2014 Lee
 8,787,060 B2 7/2014 Lee
 8,803,545 B2 8/2014 Yoko et al.
 8,817,547 B2 8/2014 Veces
 8,817,549 B2 8/2014 Shin et al.
 8,830,715 B2 9/2014 Yu et al.
 8,866,303 B2 10/2014 Kim
 8,867,286 B2 10/2014 Wu
 8,873,282 B2 10/2014 Min
 8,891,278 B1 11/2014 Stephens, Jr.
 8,897,053 B1 11/2014 Stephens, Jr.
 8,930,647 B1 1/2015 Smith
 8,943,224 B2 1/2015 Ware et al.
 8,949,538 B2 2/2015 Jeddeloh
 9,087,555 B2 7/2015 Nomoto et al.
 9,123,552 B2 9/2015 Keeth
 9,142,262 B2 9/2015 Ware
 9,160,349 B2 10/2015 Ma
 9,171,824 B2 10/2015 Best
 9,177,609 B2 11/2015 Manuel et al.
 9,318,160 B2 4/2016 Lee
 9,484,326 B2 11/2016 Keeth et al.
 9,502,345 B2 11/2016 Sunpil et al.
 9,659,601 B2 5/2017 Lee
 10,290,328 B2* 5/2019 Lee G11C 7/1057
 10,902,886 B2* 1/2021 Lee G11C 5/06
 2001/0008006 A1 7/2001 Klein
 2002/0038405 A1 3/2002 Leddige
 2002/0048195 A1 4/2002 Klein
 2002/0112119 A1 8/2002 Halbert
 2003/0070052 A1 4/2003 Lai
 2004/0098528 A1 5/2004 Janzen
 2005/0010737 A1 1/2005 Ware
 2005/0257109 A1 11/2005 Averbuj
 2005/0281096 A1 12/2005 Bhakta
 2006/0117152 A1 6/2006 Amidi
 2006/0136765 A1 6/2006 Poisner et al.
 2006/0202317 A1 9/2006 Barakat et al.
 2006/0233012 A1 10/2006 Sekiguchi
 2006/0259678 A1 11/2006 Gervasi
 2006/0262586 A1 11/2006 Solomon et al.
 2006/0277355 A1 12/2006 Ellsberry
 2006/0294437 A1 12/2006 Washburn et al.
 2007/0070669 A1 3/2007 Tsern
 2007/0096332 A1 5/2007 Satoh

2007/0293094 A1 12/2007 Aekins
 2008/0025123 A1 1/2008 Rajan
 2008/0025134 A1 1/2008 Rajan
 2008/0025137 A1 1/2008 Rajan
 2008/0094808 A1 1/2008 Kanapathippillai
 2008/0104352 A1 5/2008 Talbot
 2008/0162790 A1 7/2008 Im
 2008/0253085 A1 10/2008 Soffer
 2008/0296779 A1 12/2008 Matsui
 2008/0307240 A1 12/2008 Dahan et al.
 2009/0020608 A1 1/2009 Bennett et al.
 2009/0027844 A1 1/2009 Chen et al.
 2009/0070727 A1 3/2009 Solomon
 2009/0103345 A1 4/2009 McLaren
 2009/0103387 A1 4/2009 Shau
 2009/0248969 A1 10/2009 Wu
 2009/0290442 A1 11/2009 Rajan
 2009/0296503 A1 12/2009 Chu
 2010/0020583 A1 1/2010 Kang
 2010/0090338 A1 4/2010 Lee
 2010/0091537 A1 4/2010 Best
 2010/0110745 A1 5/2010 Jeddeloh
 2010/0110748 A1 5/2010 Best
 2010/0125681 A1 5/2010 Patel
 2010/0174858 A1 7/2010 Chen
 2010/0195364 A1 8/2010 Riho
 2011/0006360 A1 1/2011 Ikebuchi
 2011/0016250 A1 1/2011 Lee
 2011/0016269 A1 1/2011 Lee
 2011/0026293 A1 2/2011 Riho
 2011/0050320 A1 3/2011 Gillingham
 2011/0085408 A1 4/2011 Solomon et al.
 2011/0103156 A1 5/2011 Kim
 2011/0108888 A1 5/2011 Or-Bach
 2011/0125966 A1 5/2011 Amidi
 2011/0125982 A1 5/2011 Choi
 2011/0156232 A1 6/2011 Youn
 2011/0169171 A1 7/2011 Marcoux
 2011/0175639 A1* 7/2011 Yoko G11C 5/04
 257/E23.011
 2011/0193226 A1 8/2011 Kirby
 2011/0208906 A1 8/2011 Gillingham
 2011/0211411 A1* 9/2011 Ide G11C 11/4076
 365/219
 2012/0106229 A1 5/2012 Kondo

FOREIGN PATENT DOCUMENTS

JP 09237492 9/1997
 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
 KR 20050073902 A 7/2005
 WO WO1999030240 6/1999
 WO WO2005117021 A 12/2005
 WO WO2010021410 A 2/2010
 WO WO2010138480 12/2010
 WO WO2011049710 4/2011
 WO WO2011094437 8/2011

OTHER PUBLICATIONS

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, Defendant's P.R. 3-3 Invalidity Contentions for U.S. Pat. Nos. 9,318,160 and 8,787,060, filed Jul. 13, 2022, 199 pages.
 USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 9,318,160 in View of Admitted Prior Art,' filed Jul. 13, 2022, 39 pages.
 USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 9,318,160 in View of U.S. Patent App. Pub. No. 2008/0025137 (Rajan 137),' filed Jul. 13, 2022, 42 pages.

US 12,308,087 B2

Page 4

(56) **References Cited**

OTHER PUBLICATIONS

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 9,318,160 in View of U.S. Patent App. Pub. No. 2011/0103156 (Kim),' filed Jul. 13, 2022, 43 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 9,318,160 in View of U.S. Pat. No. 9,142,262 (Ware),' filed Jul. 13, 2022, 33 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 9,318,160 in View of U.S. Patent App. Pub. No. 2011/0026293 (Riho 293),' filed Jul. 13, 2022, 53 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 9,318,160 in View of U.S. Patent App. Pub. No. 2010/0195364 (Riho 364),' filed Jul. 13, 2022, 46 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 9,318,160 in View of U.S. Patent App. Pub. No. 2010/0091537 (Best),' filed Jul. 13, 2022, 60 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 9,318,160 in View of U.S. Pat. No. 8,258,619 (Foster),' filed Jul. 13, 2022, 24 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 9,318,160 in View of U.S. Patent App. Pub. No. 2010/0110745 (Jeddeloh),' filed Jul. 13, 2022, 36 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 9,318,160 in View of U.S. Patent App. Pub. No. 2011/0208906 (Gillingham),' filed Jul. 13, 2022, 61 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 9,318,160 in View of U.S. Pat. No. 8,120,958 (Bilger),' filed Jul. 13, 2022, 45 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 9,318,160 in View of U.S. Pat. No. 9,123,552 (Keeth),' filed Jul. 13, 2022, 47 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 9,318,160 in View of U.S. Pat. No. 7,969,192 (Wyman),' filed Jul. 13, 2022, 31 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 9,318,160 in View of U.S. Pat. No. 9,160,349 (Ma),' filed Jul. 13, 2022, 66 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 9,318,160 in View of U.S. Pat. No. 7,796,446 (Ruckerbauer),' filed Jul. 13, 2022, 44 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 9,318,160 in View of Gabriel H. Loh, 3D-Stacked Memory Architectures for Multi-Core Processors, Georgia Institute of Technology, International Symposium on Computer Architecture (2008) (Loh),' filed Jul. 13, 2022, 28 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 9,318,160 in View of U.S. Pat. No. 8,041,881 (Rajan 881),' filed Jul. 13, 2022, 65 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 9,318,160 in View of JEDEC Proposals regarding High Bandwidth Memory (HBM) and Low Power Memories (JEDEC HBM and Low Power Proposals),' filed Jul. 13, 2022, 95 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 9,318,160 in View of Micron Hybrid Memory Cube,' filed Jul. 13, 2022, 30 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 9,318,160 in View of U.S. Pat. No. 8,471,362 (Lee),' filed Jul. 13, 2022, 104 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 8,787,060 in View of Admitted Prior Art,' filed Jul. 13, 2022, 96 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 8,787,060 in View of U.S. Patent App. Pub. No. 2008/0025137 (Rajan 137),' filed Jul. 13, 2022, 61 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 8,787,060 in View of U.S. Patent App. Pub. No. 2011/0103156 (Kim),' filed Jul. 13, 2022, 65 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 8,787,060 in View of U.S. Pat. No. 9,142,262 (Ware),' filed Jul. 13, 2022, 60 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 8,787,060 in View of U.S. Patent App. Pub. No. 2011/0026293 (Riho 293),' filed Jul. 13, 2022, 71 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 8,787,060 in View of U.S. Patent App. Pub. No. 2010/0195364 (Riho 364),' filed Jul. 13, 2022, 66 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 8,787,060 in View of U.S. Patent App. Pub. No. 2010/0091537 (Best),' filed Jul. 13, 2022, 76 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 8,787,060 in View of U.S. Pat. No. 8,258,619 (Foster),' filed Jul. 13, 2022, 43 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 8,787,060 in View of U.S. Patent App. Pub. No. 2010/0110745 (Jeddeloh),' filed Jul. 13, 2022, 53 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 8,787,060 in View of U.S. Patent App. Pub. No. 2011/0208906 (Gillingham),' filed Jul. 13, 2022, 81 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 8,787,060 in View of U.S. Pat. No. 8,120,958 (Bilger),' filed Jul. 13, 2022, 63 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 8,787,060 in View of U.S. Pat. No. 9,123,552 (Keeth),' filed Jul. 13, 2022, 65 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 8,787,060 in View of U.S. Pat. No. 7,969,192 (Wyman),' filed Jul. 13, 2022, 48 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 8,787,060 in View of U.S. Pat. No. 9,160,349 (Ma),' filed Jul. 13, 2022, 101 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 8,787,060 in View of U.S. Pat. No. 7,796,446 (Ruckerbauer),' filed Jul. 13, 2022, 67 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Invalidity of U.S. Pat. No. 8,787,060 in View of Gabriel H. Loh, 3D-Stacked Memory Archi-

US 12,308,087 B2

Page 5

(56) **References Cited**

OTHER PUBLICATIONS

tures for Multi-Core Processors, Georgia Institute of Technology, International Symposium on Computer Architecture (2008) (Loh),’ filed Jul. 13, 2022, 46 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, ‘Invalidity of U.S. Pat. No. 8,787,060 in View of U.S. Pat. No. 8,041,881 (Rajan 881),’ filed Jul. 13, 2022, 82 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, ‘Invalidity of U.S. Pat. No. 8,787,060 in View of JEDEC Proposals regarding High Bandwidth Memory (HBM) and Low Power Memories (JEDEC HBM and Low Power Proposals),’ filed Jul. 13, 2022, 114pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, ‘Invalidity of U.S. Pat. No. 8,787,060 in View of Micron Hybrid Memory Cube,’ filed Jul. 13, 2022, 50 pages.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, ‘Invalidity of U.S. Pat. No. 8,787,060 in View of U.S. Pat. No. 8,471,362 (Lee),’ filed Jul. 13, 2022, 124 pages.

JEDEC Proposal, ‘Future Mobile Memory (FMM) Wide I/O Refresh Scheme,’ Item #1777.18, Sep. 10, 2010, 3 pages.

JEDEC Proposal, ‘Wide I/O Ball Map Baseline Proposal,’ Item #1777.29, Jun. 2010, 4 pages.

JEDEC Proposal, ‘Read Clock Proposal,’ FMD—Wide IO TG, Mar. 2010, 6 pages.

JEDEC Proposal, ‘Advanced Memory Package Proposal,’ Item #1782.01, Mar. 2010, 7 pages.

JEDEC Proposal, ‘Wide-IO TG Report,’ FM-Mobile Wide IO TG, Item #1777.00, Dec. 2009, 13 pages.

JEDEC Proposal, ‘MIPI@M-PHY Future Mobile PHY Proposal,’ JC42.6 Item #1776.10, Sep. 2009, 38 pages.

JEDEC Proposal, ‘A Stackable, Configurable Memory Sheet for ASICs: A First Showing,’ 2011 TG, Item #1787.01, Sep. 13, 2010, 11 pages.

JEDEC Proposal, ‘Future High Bandwidth Memory TG,’ TG42_1: TG Report, Item #1797.00, Sep. 2011, 20 pages.

JEDEC Proposal, ‘TSV Tile Memory Clocking & Command,’ 2011 TG, Item #1787.05, Apr. 2021/22, 5 pages.

JEDEC Proposal, ‘Future High Bandwidth Memory TG,’ TG42_1: TG Report, Item #1797.00, Jun. 2011, 18 pages.

J. Thomas Pawlowski Presentation, ‘Hybrid Memory Cube (HMC),’ Hot Chips 23, © 2011 Micron Technology, Inc.

USDC, ED Texas, *Netlist, Inc. v. Micron Technology, Inc.*, Case 2:22-cn-00203-JRG-RSP, “Complaint,” filed Jun. 22, 2010, 71 pages.

International Search Report and Written Opinion, PCT/US2011/059209, dated Jan. 31, 2013.

“Posts Tagged ‘3D Stacking’,” from Chip Design Mag., <http://chipdesignmag.com/lpd/blog/tag/3d-stacking/> (Printed Oct. 13, 2011).

Ahmad et al., “Modeling of peak-to-peak switching noise along a vertical chain of power distribution TSV pairs in a 30 stack of ICs interconnected through TSVs,” Norchip Conference, Article No. 5669473, Nov. 15-16, 2010, IEEE Computer Society.

Altera, ACEX iK, Programmable Logic Device Family, Data Sheet, May 2003, Ver 3.4.

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.

Black et al., “Die Stacking (30) Microarchitecture,” MICR0-39, 39th Annual IEEE/ACM International Symposium on, Dec. 2006, 469-479, Orlando, FL.

Daneshtalab et al., “CMIT—A novel cluster-based topology for 30 stacked architectures,” 30 Systems Integration Conference (3DIC), 2010 IEEE International, Nov. 16-18, 2010, pp. 1-5.

Excerpts from the Authoritative Dictionary of IEEE Standard Terms, Seventh Edition, Standards Information Network, IEEE Press (2000), pp. 349-411.

Excerpts from the Authoritative Dictionary of IEEE Standard Terms, Seventh Edition, Standards Information Network, IEEE Press (2000), pp. 133-265.

Funaya et al., “Cache partitioning strategies for 3-D stacked vector processors,” IEEE 30 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 30 Die-Stacked DRAMs,” Microarchitecture, 2007, 40th Annual IEEE/ACM International Symposium, Dec. 1-5, 2007, pp. 134-145.

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.

Jacob, Bruce, et al. “Memory Systems—Cache, DRAM, Disk,” Elsevier Inc. 2008, Chapters 7 and 10.

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 30 Packaging Technology for High-Density Stacked DRAM,” VLSI Technology, Systems and Applications, 2007, Apr. 23-25, 2007, pp. 1-2.

Kim et al., A Quantitative Analysis of Performance Benefits of 30 Die Stacking on Mobile and Embedded Soc.

Kurita et al., “A 3-D packaging technology with highly-parallel memory/logic interconnect,” IEICE Transactions on Electronics, v E92-C, No. 12, pp. 1512-1522, 2009, Maruzen Col, Ltd.

Kurita et al., “Vertical Integration of Stacked DRAM and High-Speed Logic Device Using SMAFTI Technology,” Advanced Packaging, IEEE Transactions, Aug. 2009, vol. 32 Issue 3, pp. 657-665.

Loh, “3D-Stacked Memory Architectures for Multi-core Processors,” Proceedings of the 35th Annual International Symposium on Computer Architecture (ISCA ’08), 453-464, IEEE Computer Society, Washington DC, USA.

McCluskey, Edward J., *Logic Design Principles with Emphasis on Testable Semicustom Circuits*, Prentice Hall, 1986, pp. 104-107 and 119-120.

Microsoft Computer Dictionary, Fifth Edition, Microsoft Press, p. 334, 2002.

Reese, “Introduction to Logic Synthesis using Verilog HDL,” Morgan&Claypool Publishers, 2006, pp. 1-28.

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.

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.

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.

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.

US 12,308,087 B2

Page 6

(56)

References Cited

OTHER PUBLICATIONS

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. IPR2014-01029, 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.

Inter Partes Review Case No. 2014-00882, Exhibit 1023, 'Dr. Srinivasan Jagannathan Supplemental Declaration,' filed Jun. 19, 2015.

Inter Partes Review Case No. 2014-00882, Exhibit 2002, 'Declaration of Professor Carl Sechen,' filed May 8, 2015.

Inter Partes Review Case No. 2014-00882, Final Written Decision 35 USC 318 and 37 CFR 42.73.

Inter Partes Review Case No. 2014-00882, Patent Owner Response, filed May 8, 2015.

Inter Partes Review Case No. 2014-00882, Petitioner Diablo Technologies, Inc.'s Reply to Patent Owner Netlist, Inc.'s Response, filed Jun. 19, 2015.

Inter Partes Review Case No. 2014-00883, Exhibit 1026, 'Dr. Srinivasan Jagannathan Supplemental Declaration,' filed Jun. 19, 2015.

Inter Partes Review Case No. 2014-00883, Final Written Decision—35 USC 318 and 37 CFR 42.73, filed Dec. 14, 2015.

Inter Partes Review Case No. 2014-00883, Patent Owner Response, filed May 8, 2015.

Inter Partes Review Case No. 2014-00883, Petitioner Diablo Technologies, Inc.'s Reply to Patent Owner Netlist, Inc.'s Response, filed Jun. 19, 2015.

Inter Partes Review Case No. 2014-01011, Exhibit 1028, 'Dr. Srinivasan Jagannathan Supplemental Declaration,' filed Jun. 19, 2015.

Inter Partes Review Case No. 2014-01011, Final Written Decision—35 USC 318 and 37 CFR 42.73, filed Dec. 14, 2015.

Inter Partes Review Case No. 2014-01011, Patent Owner Response, filed May 8, 2015.

Inter Partes Review Case No. 2014-01011, Petitioner Diablo Technologies, Inc.'s Reply to Patent Owner Netlist, Inc.'s Response, filed Jun. 19, 2015.

Inter Partes Review Case No. IPR2014-00882, Corrected Petition for Inter Partes Review of U.S. Pat. No. 7,881,150, filed on Jul. 8, 2014.

Inter Partes Review Case No. IPR2014-00882, Exhibit 1007 to Petition for Inter Partes Review, "Declaration of Dr. Srinivasan Jagannathan," filed on Jun. 22, 2014.

Inter Partes Review Case No. IPR2014-00883, Corrected Petition for Inter Partes Review of U.S. Pat. No. 8,081,536, filed on Jul. 8, 2014.

Inter Partes Review Case No. IPR2014-00883, Exhibit 1011 to Petition for Inter Partes Review, "Declaration of Dr. Srinivasan Jagannathan," filed on Jun. 21, 2014.

Inter Partes Review Case No. IPR2014-01011, Corrected Petition for Inter Partes Review of U.S. Pat. No. 7,881,150, filed on Jul. 8, 2014.

Inter Partes Review Case No. IPR2014-01011, Exhibit 1007 to Petition for Inter Partes Review, "Declaration of Dr. Srinivasan Jagannathan," filed on Jun. 22, 2014.

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, Decision Denying Institution of Inter Partes Review 37 C.F.R. § 42.108, issued Mar. 9, 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 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 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. 7,881,150, IPR Case No. IPR2014-01011, Decision—Institution of Inter Partes Review 37 C.F.R. § 42.108, issued Dec. 16, 2014.

Inter Partes Review of U.S. Pat. No. 7,881,150, IPR Case No. IPR2014-01011, Exhibit 3001 to Decision—Institution of Inter Partes Review, Excerpts from IEEE Dictionary, issued Dec. 16, 2014.

Inter Partes Review of U.S. Pat. No. 7,881,150, IPR Case No. IPR2014-01011, Exhibit 3002 to Decision—Institution of Inter Partes Review, Excerpts from IEEE Dictionary, issued Dec. 16, 2014.

Inter Partes Review of U.S. Pat. No. 7,881,150, IPR Case No. IPR2014-01011, Exhibit 3003 to Decision—Institution of Inter Partes Review, Excerpts from Oxford English Dictionary, issued Dec. 16, 2014.

Inter Partes Review of U.S. Pat. No. 7,881,150, IPR Case No. IPR2014-01011, Exhibit 3004 to Decision—Institution of Inter Partes Review, Excerpts from Oxford English Dictionary, issued Dec. 16, 2014.

Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Exhibit 2001, "Declaration of Professor Carl Sechen," filed Feb. 22, 2016.

Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Patent Owner Response to Petition, filed Feb. 22, 2016.

Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Petitioner's Reply, filed May 19, 2016.

Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Supplemental Declaration of Dr. Jagannathan, filed May 19, 2016.

Inter Partes Review of U.S. Pat. No. 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. 8,081,536, Case No. IPR2015-01021, Exhibit 2001, "Declaration of Professor Carl Sechen," filed Feb. 22, 2016.

Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Patent Owner Response to Petition, filed Feb. 22, 2016.

Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Petitioner's Reply, filed May 19, 2016.

Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Supplemental Declaration of Dr. Jagannathan, filed May 19, 2016.

Inter Partes Review of U.S. Pat. No. 7,532,537 B2, Case No. IPR2017-00667, Final Written Decision, filed Jul. 18, 2018.

Inter Partes Review of U.S. Pat. No. 7,532,537 B2, Case No. IPR2017-00668, Final Written Decision, filed Jul. 18, 2018.

Inter Partes Review of U.S. Pat. No. 8,516,185 B2, Case No. IPR2017-00577, Final Written Decision, filed Jul. 5, 2018.

Inter Partes Review of U.S. Pat. No. 9,606,907 B2, Case No. IPR2018-00362, and Case No. IPR2018-00363, Decision Granting Institution of Inter Partes Review, filed Aug. 6, 2018.

Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 1003, 'Declaration of Harold S. Stone,' filed Jan. 5, 2017.

Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 1007, 'JEDEC JESD79 publication (Jun. 2000),' filed Jan. 5, 2017.

Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibits 1009, 'TI 74LS245 datasheet (2002),' filed Apr. 10, 2017.

US 12,308,087 B2

Page 7

(56) **References Cited**

OTHER PUBLICATIONS

Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Patent Owner's Preliminary Response, filed Apr. 10, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Patent Owner's Response, filed Oct. 25, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Petitioners' Reply to Patent Owner's Response, filed Feb. 6, 2018.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Trial Instituted Document, filed Jul. 7, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1002, 'Excerpts of FH for U.S. Pat. No. 8,756,364,' filed Dec. 30, 2016.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1003, 'Declaration of Harold Stone,' filed Dec. 15, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1007, 'JEDEC JESD79 publication (Jun. 2000),' filed Dec. 30, 2016.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1008, 'JEDEC Declaration for DDR Specification,' filed Dec. 30, 2016.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1015, 'Texas Instruments 74LS245 datasheet (2002),' filed Dec. 30, 2016.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1017, '2000-03 Samsung CMOS SDRAM data sheet,' filed Sep. 15, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1019, 'Final Decision [IPR2014-01011],' filed Sep. 15, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1020, 'Final Decision [IPR2014-00883],' filed Dec. 30, 2016.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1023, 'Gordon Moore 1965 article,' filed Dec. 30, 2016.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1026, 'Xilinx CoolRunner XPLA3 CPLD product specification (2000),' filed Dec. 15, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1027, 'Xilinx Programmable Logic Design Quick Start Handbook (2004),' filed Jan. 9, 2018.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1028, 'Xilinx TQFP package datasheet (2000),' filed Jan. 9, 2018.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1029, 'Xilinx VQFP package datasheet (2002),' filed Feb. 9, 2018.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1030, 'Xilinx CS280 package datasheet (1999),' filed Dec. 15, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 1031, 'Petitioners Demonstratives,' filed Apr. 9, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 2001, 'Exhibit 2001,' filed Dec. 15, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 2002, 'Exhibit 2002,' filed Dec. 15, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 2003, 'Exhibit 2003,' filed Feb. 9, 2018.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 2004, 'Exhibit 2004,' filed Dec. 30, 2016.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibits 2005, 'Exhibit 2005,' filed Sep. 15, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Final Written Decision, filed May 3, 2018.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Patent Owner's Preliminary Response, filed Apr. 9, 2017.

Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Patent Owner's Response, filed Sep. 15, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Petition for Inter Partes Review, filed Dec. 30, 2016.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Trial Instituted Document, filed May 15, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907 B2, Case No. IPR2018-00364, and Case No. IPR2018-00365, Decision Granting Institution of Inter Partes Review, filed Jun. 29, 2018.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1003, 'Declaration of Dr. Harold Stone,' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1009, 'JEDEC Standard Double Data Rate DDR SDRAM Specification, JESD79 Jun. 2000,' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1010, 'JEDEC Standard 21-C, DDR SDRAM Registered DIMM Design Specification Jan. 2002,' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1011, 'JEDEC Standard DDR2 SDRAM Specification, JESD79-2B Jan. 2005,' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1012, 'Declaration of John J. Kelly Regarding Records of Joint Electron Device Engineering Council JEDEC,' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1022, 'Decision Denying Institution of Inter Partes Review of U.S. Pat. No. 8,516,185, *SanDisk Corp. v. Netlist, Inc.*, IPR2014-01029, Paper No. 11 Dec. 16, 2014,' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1023, 'Decision Denying Institution of Inter Partes Review of U.S. Pat. No. 8,516,185, *Smart Modular Techs. Inc. v. Netlist, Inc.*, IPR2014-01369, Paper No. 12 Mar. 9, 2015,' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1024, 'Excerpts from the Hearing in Certain Memory Modules and Components Thereof, and Products Containing Same, Inv. No. 337-TA-1023 May 8, 2017,' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1025, 'Complainant Netlist, Inc.'s Initial Post-Hearing Brief, Certain Memory Modules and Components Thereof, and Products Containing Same, Inv. No. 337-TA-1023 May 30, 2017 excerpts relevant to '185 patent,' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1026, 'Respondents' Post-Hearing Brief, Certain Memory Modules and Components Thereof, and Products Containing Same, Inv. No. 337-TA-1023 May 30, 2017 excerpts relevant to '185 patent,' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1027, 'High-quality versions of demonstrative graphics included in Respondents' Post-Hearing Brief (Ex. 1026),' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1028, 'Complainant Netlist Inc.'s Reply Post-Hearing Brief, Certain Memory Modules and Components Thereof, and Products Containing Same, Inv. No. 337-TA-1023 (Jun. 9, 2017) (excerpts relevant to '185 patent),' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1029, 'Respondents' Reply Post-Hearing Brief, Certain Memory Modules and Components Thereof, and Products Containing Same, Inv. No. 337-TA-1023 (Jun. 9, 2017) (excerpts relevant to '185 patent),' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1030, 'High-quality versions of demonstrative graphics included in Respondents' Reply Post-Hearing Brief (Ex. 1029),' filed Dec. 22, 2017.
Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1031, 'Institution of Inter Partes Review of U.S. Pat. No. 8,516,185, *SK hynix Inc. et al. v. Netlist, Inc.*, IPR2017-00577, Paper No. 8 (Jul. 7, 2017),' filed Dec. 22, 2017.

US 12,308,087 B2

Page 8

(56)

References Cited

OTHER PUBLICATIONS

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1032, 'Final Written Decision, *Diablo Techs., Inc. v. Netlist, Inc.*, IPR2014-00882, Paper No. 33 (Dec. 14, 2015),' filed Dec. 22, 2017.

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1033, '*Netlist, Inc. v. Diablo Techs., Inc.*, No. 2016-1742 (Fed. Cir. Jul. 25, 2017),' filed Dec. 22, 2017.

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1034, 'Netlist's Infringement Claim Chart for U.S. Pat. No. 9,606,907 Jun. 14, 2017,' filed Dec. 22, 2017.

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1035, 'Stone, H.S. Microcomputer Interfacing, Reading, MA: Addison Wesley, 1982,' filed Dec. 22, 2017.

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1039, 'Intel E7525 Memory Controller Hub (MCH) Chipset Datasheet (Jun. 2004),' filed Dec. 22, 2017.

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1040, 'Initial Determination, Certain Memory Modules and Components Thereof, and Products Containing Same, Inv. No. 337-TA-1023 (Nov. 14, 2017) (redacted excerpts),' filed Dec. 22, 2017.

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1041, 'Bruce Jacob et al., Memory System: Cache, DRAM, Disk (2008) (excerpts),' filed Dec. 22, 2017.

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1042, 'Direct RDRAM datasheet (2000),' filed Dec. 22, 2017.

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Exhibits 1048, 'Certified translation of the Preliminary Opinion issued by the German Patent and Trademark Office (dated Jan. 8, 2019),' filed Feb. 11, 2019.

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Paper 14, 'Patent Owner's Response,' filed Oct. 19, 2018.

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Paper 23, 'Corrected Petitioners' Reply to Patent Owner's Response,' filed Feb. 20, 2019.

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Paper 29, 'Termination Decision Document,' filed Jun. 27, 2019.

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00363, Petition for Review of U.S. Pat. No. 9,606,907, filed Dec. 22, 2017.

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00363, Preliminary Response, filed Apr. 10, 2018.

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00364, Petition for Inter Partes Review of U.S. Pat. No. 9,606,907, filed Dec. 27, 2017.

Inter Partes Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00365, Petition for Inter Partes Review of U.S. Pat. No. 9,606,907, filed Dec. 27, 2017.

Examination Report, European Patent Application No. 10730021.2, dated Apr. 14, 2014.

Response to Examination Report, European Patent Application No. 10730021.2, dated Jun. 4, 2014.

Examination Report, European Patent Application No. 10730021.2, dated Apr. 29, 2015.

Response to Examination Report, European Patent Application No. 10730021.2, dated Apr. 29, 2015, filed Nov. 4, 2015.

English Translation of the Notice of Grounds for Rejection, Korean Patent Application No. 2012-7004038, dated May 11, 2016.

Notice of Allowance, U.S. Appl. No. 13/970,606, filed Aug. 20, 2013, dated Jun. 27, 2016.

Non-Final Office Action, U.S. Appl. No. 13/970,606, filed Aug. 20, 2013, dated Nov. 23, 2015.

Response to Non-Final Office Action, U.S. Appl. No. 13/970,606, filed Aug. 20, 2013, dated Mar. 23, 2016.

Response, Article 94.3 EPC, EO 18179414.0, Aug. 6, 2019, 23 pgs.

Netlist, Inc., Third Party Observation, EP 3404660, Aug. 6, 2019, 34 pgs.

Office Action, EP 18179414, dated Apr. 2, 2019, 6 pgs.

Patent document filed by third party, EP 18179414, Feb. 22, 2019, 10 pgs.

Observations by third parties, EP 18179414.0, Feb. 22, 2019, 1 pg.

Patent Document filed by a third party, Feb. 20, 2019, 10 pgs.

Netlist, Inc., Observations by third parties, EP 3404660, Feb. 20, 2019, 46 pgs.

Reply to Written Opinion prepared by the EPO, EP 18179414.0, dated Jan. 17, 2019, 20 pgs.

Netlist, Inc., Information on Search Strategy, EP 18179414, Oct. 23, 2018, 1 pg.

European Search Report, EP 18179414, dated Oct. 23, 2018, 3 pgs.

European Search Opinion, EP 18179414, dated Oct. 23, 2018, 3 pgs.

CN201080039043.0, 'Certified translation of the Examination Decision issued by the State Intellectual Property Office of the P.R.C. (dated May 30, 2018),' May 30, 2018, 43 pgs.

Inter Partes Review of U.S. Pat. No. 8,787,060, Case No. IPRIPR2022-01428, Petition for Inter Partes Review of U.S. Pat. No. 8,787,060, filed Aug. 26, 2022.

Inter Partes Review of U.S. Pat. No. 8,787,060, Case No. IPRIPR2022-01428, Patent Owner Preliminary Response, filed Jan. 19, 2023.

Inter Partes Review of U.S. Pat. No. 8,787,060, Case No. IPRIPR2022-01428, Petitioner's Authorized Reply to POPR, filed Feb. 17, 2023.

Inter Partes Review of U.S. Pat. No. 8,787,060, Case No. IPRIPR2022-01428, PO's Preliminary Sur-Reply, filed Mar. 1, 2023.

Inter Partes Review of U.S. Pat. No. 8,787,060, Case No. IPRIPR2022-01428, Institution Decision: Granting Institution of Inter Partes Review 35 U.S.C. sec. 314, filed Apr. 12, 2023.

Inter Partes Review of U.S. Pat. No. 9,318,160, Case No. IPRIPR2022-01427, Petition for Inter Partes Review of U.S. Pat. No. 9,318,160, filed Aug. 26, 2022.

Inter Partes Review of U.S. Pat. No. 9,318,160, Case No. IPRIPR2022-01427, Patent Owner Preliminary Response, filed Jan. 19, 2023.

Inter Partes Review of U.S. Pat. No. 9,318,160, Case No. IPRIPR2022-01427, Petitioner's Authorized Reply to POPR, filed Feb. 17, 2023.

Inter Partes Review of U.S. Pat. No. 9,318,160, Case No. IPRIPR2022-01427, Patent Owner's Preliminary Sur-Reply, filed Mar. 1, 2023.

Inter Partes Review of U.S. Pat. No. 9,318,160, Case No. IPRIPR2022-01427, Institution Decision: Granting Institution of Inter Partes Review 35 U.S.C. sec. 314, filed Apr. 12, 2023.

Inter Partes Review of U.S. Pat. No. 8,787,060, Case No. IPRIPR2022-01428, Exhibit 1003, 'Wolfe declaration,' filed Aug. 26, 2022.

Inter Partes Review of U.S. Pat. No. 9,318,160, Case No. IPRIPR2022-01427, Exhibit 1003, 'Wolfe declaration,' filed Aug. 26, 2022.

Inter Partes Review of U.S. Pat. No. 9,318,160, Case No. IPRIPR2022-01427, EX2009, 'Park, Sung Joo Deposition Transcript (Nov. 23, 2022),' filed Jan. 19, 2023.

IPR2018-00362, U.S. Pat. No. 9,606,907, Final Written Decision, entered: Jun. 27, 2019, 94 pgs.

IPR2017-00667, U.S. Pat. No. 7,532,537 B2, Final Written Decision 35 U.S.C. § 318(a) and 37 C.F.R. § 42.73, entered Jul. 18, 2018, 44 pgs.

IPR2017-00668, U.S. Pat. No. 7,532,537 B2, Final Written Decision 35 U.S.C. § 318(a) and 37 C.F.R. § 42.73, entered Jul. 18, 2018, 46 pgs.

USPTO Patent Trial and Appeal Board, Appeal 2014-007777, Inter Partes Reexamination Control Nos. 95/000,577 & 95/000,577, U.S. Pat. No. 7,289,386, Decision On Appeal, issued Feb. 25, 2015, 25 pgs.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Samsung's Motion for Summary Judgment,' filed Feb. 10, 2023, 19 pgs.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Defendants' P.R. 4-2(A) Preliminary Claim Construction and Extrinsic Evidence,' Aug. 5, 2022, 55 pgs.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Samsung Defendants' Proposed Constructions of Disputed Claim Terms and Supporting Evidence,' filed on Aug. 19, 2022, 69 pgs.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, 'Defendants' Responsive Claim Construction Brief,' filed Sep. 16, 2022, 38 pgs.

US 12,308,087 B2

Page 9

(56)

References Cited

OTHER PUBLICATIONS

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, ‘Claim Construction Order,’ filed Dec. 14, 2022, 35 pgs.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, ‘Plaintiff Netlist Inc.’s objections to the magistrate judge’s memorandum opinion and order regarding claim construction,’ filed on Dec. 28, 2022, 23 pgs.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, ‘Defendants’ Response to Netlist’s Objections to Claim Construction Memorandum Order,’ filed on Jan. 11, 2023, 11 pgs.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, ‘Plaintiff Netlist Inc.’s Opposition to Samsung’s Motion for Summary Judgment of Non-Infringement,’ filed Feb. 23, 2023, 23 pgs.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, ‘Samsung’s Fourth Amended Initial and Additional Disclosures,’ filed Dec. 6, 2022, 23 pgs.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, ‘Netlist Inc.’s Motion for Partial Summary Judgment on Samsung’s Invalidity Defenses,’ filed on Feb. 9, 2023, 22 pgs.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, ‘Samsung’s Opposition to Netlist’s Motion for Partial Summary Judgment on Samsung’s Invalidity Defenses,’ filed on Feb. 22, 2023, 29 pgs.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, ‘Complaint’ filed Dec. 20, 2021, 49 pgs.

USDC, ED Texas, *Netlist Inc. v. Samsung Electronics Co., Ltd.*, Civil Action No. 2:21-CV-463-JRG, ‘The Samsung Defendants’ Answer to First Amended Complaint,’ filed May 18, 2022, 53 pgs.

JEDEC Standard No. 21-C, GDDR4 Specific SGRAM Functions, p. 3.11.5.8, From JEDEC Board Ballot JCB-06-007, Formulated Under the Cognizance of the JC42.3 Sub-Committee on DRAM Parametrics, Release 1.0, Nov. 2005.

JEDEC Standard No. 21-C, ‘PC2100 and PC1600 DDR SDRAM Registered DIMM,’ Revision 1.3, Jan. 2002.

JEDEC Standard No. 79-3C, DDR3 SDRAM, JESD79-3C, Revision of JESD79-3B, Apr. 2008.

JEDEC Standard No. MO-207, ‘Square & Rectangular Die-Size, Ball Grid Array Family,’ JEDEC Solid State Product Outline, Dec. 10, 2022 pgs.

Declaration of Julie Carlson re JEDEC Standard No. 79-3C, DDR3 SDRAM, JESD79-3C (Revision of JESD79-3B, Apr. 2008), dated Sep. 16, 2021, 7 pgs.

Cho et al., ‘A 1.2V 64Gb 341GB/s HBM2 Stacked DRAM with Spiral Point-to-Point TSV Structure and Improved Bank Group Data Control,’ IEEE International Solid-State Circuits Conference, ISSCC 2018 / Session 12 / DRAM / 12.3, Article No. 978-1-5090-4940-0/18, 2018, 3pgs.

Jacob, Bruce, et al. ‘Memory Systems—Cache, DRAM, Disk,’ Elsevier Inc. 2008, 296 pgs. (uploaded in three parts).

Bruce L. Jacob, ‘Synchronous DRAM Architectures, Organizations, and Alternative Technologies’, Electrical & Computer Engineering Dept., University of Maryland, MD 20742, Dec. 10, 2002, URL: <http://www.ece.umd.edu/~blj/>.

Stephen Brown and Zvonko Vranesic, ‘Fundamentals of Digital Logic with Verilog Design,’ 2nd Edition, Chapter 3, Department of Electrical and Computer Engineering University of Toronto, 2008, McGraw-Hill Companies, Inc. 10 pgs.

Harold S. Stone, ‘Microcomputer Interfacing,’ University of Massachusetts, Amherst, 1982, Addison-Wesley Publishing Company, Inc. 399 pgs. (uploaded in three parts).

Inter Partes Review of U.S. Pat. No. 9,318,160, Case No. IPRIPR2022-01427, Patent Owner Response, filed Jul. 27, 2023.

Inter Partes Review of U.S. Pat. No. 8,787,060, Case No. IPRIPR2022-01428, Patent Owner Response, filed Jul. 27, 2023.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2023, ‘Declaration of Michael C. Brogioli, Ph.D. in Support of Patent Owner’s Response,’ Jul. 27, 2023.

IPR2022-01428 of U.S. Pat. No. 8,787,060, Exhibit 2023, ‘Declaration of Michael C. Brogioli, Ph.D. in Support of Patent Owner’s Response,’ Jul. 27, 2023.

Inter Partes Review of U.S. Pat. No. 9,318,160, Case No. IPRIPR2022-01427, Petitioner’s Objections to Patent Owner’s Evidence Submitted with Patent Owner’s Response, filed Aug. 3, 2023.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2024, USDC, Eastern District of Texas, Marshall Division, Case No. 2:22-cv-203-JRG-RSP, *Netlist, Inc. vs. Micron Technology, Inc., et al.*, Transcript of Jun. 26, 2023, Videotaped Deposition of Harold Stone, Ph.D.

IPR2022-01427 and IPR2022-01428, Exhibit 2025, Transcript of Jun. 27, 2023, Videoconference Deposition of Andrew Wolfe, Ph.D.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2027, USDC, Eastern District of Texas, Marshall Division, Case No. 2:21-CV-463-JRG, *Netlist, Inc. v. Samsung Electronics Co., Ltd, et al.*, Jury Trial Day 1—Apr. 14, 2023.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2028, USDC, Eastern District of Texas, Marshall Division, Case No. 2:21-CV-463-JRG, *Netlist, Inc. v. Samsung Electronics Co., Ltd, et al.*, Trial Transcript—vol. 3 Apr. 18, 2023 excerpt.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2029, USDC, Eastern District of Texas, Marshall Division, Case No. 2:21-CV-463-JRG, *Netlist, Inc. v. Samsung Electronics Co., Ltd, et al.*, Trial Transcript—vol. 4—Apr. 19, 2023 excerpt.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2030, USDC, Eastern District of Texas, Marshall Division, Case No. 2:21-CV-463-JRG, *Netlist, Inc. v. Samsung Electronics Co., Ltd, et al.*, Robins Direct Testimony Demonstratives.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2031, USDC, Eastern District of Texas, Marshall Division, Case No. 2:21-CV-463-JRG, *Netlist, Inc. v. Samsung Electronics Co., Ltd, et al.*, Milton Direct PDX2 Demonstratives—excerpts.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2032, Netlist Presentation on U.S. Pat. No. 8,787,060, undated.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2036, JEDEC, ‘Next Big Thing: DDR4 3DS,’ Server Forum 2014.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2037, SK hynix Newsroom, ‘Continuing to Make HBM History: The Story of SK hynix’s HBM Development,’ Sep. 8, 2022.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2039, Steve, ‘Building an SRAM Substitute,’ Sep. 4, 2015.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2040, Data Centers, ‘How Long Des NAND Flash Memory Last,’ Dec. 19, 2008.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2041, Lucas Mearian, ‘Micron boosts NAND flash endurance six-fold,’ Computer World, Oct. 19, 2009.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2043, Micron, ‘2, 4, 8Gb: x8/x16 Multiplexed NAND Flash Memory Features,’ @2004 Micron Technology, Inc.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2044, STMicroelectronics, M48Z512A M48Z512AY, M48Z512AV, 4 Mbit (512 Kbit x 8) Zeropower® SRAM, Jun. 2011.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2045, Everspin Technologies, MR4A08B, 2Mx8RAM Memory, Mar. 2018.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2046, ‘3DS Technology,’ @ 2023 Impress Corporation.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2047, Neil H. E. Weste et al., ‘CMOS VLSI Design—A Circuits and Systems Perspective,’ 4th Ed. © 2011, 2005, 1993, 1985 Pearson Education, Inc., publishing as Addison-Wesley.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2048, Ji-Young Kim et al. ‘An Energy-Efficient Design of TSV I/O for HBM With a Data Rate up to 10 Gb/s,’ IEEE Journal of Solid-State Circuits, 2023.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2049, NangateOpenCellLibrary_typical, Feb. 17, 2011.

US 12,308,087 B2

Page 10

(56)

References Cited

OTHER PUBLICATIONS

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2050, "Principles of VLSI Design—Capacitance and Resistance Model", CMPE 413, UMBC, undated.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2051, Kyomin Sohn et al., "A 1.2 V 20 nm 307 GB/s HBM DRAM With At-Speed Wafer-Level IO Test Scheme and Adaptive Refresh Considering Temperature Distribution," IEEE 2016.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2052, Jim Cooke, EDN, "Flash memory 101: An introduction to NAND flash," Mar. 20, 2006.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2053, ECE UNM, "Hardware Design with VHDL Design Example: SRAM," Dec. 3, 2008.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2054, Donghyuk Lee, et al. "Simultaneous Multi-Layer Access: Improving 3D-Stacked Memory Bandwidth at Low Cost," ACM Transactions on Architecture and Code Optimization, vol. 12, No. 4, Article 63, Publication date: Dec. 2015.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2055, Tech Blog, "High Bandwidth Memory is Catching on for Data-Intensive Computer," May 21, 2018.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2100, FIG. 3 of US Patent Publication No. 20110103156.

IPR2022-01427 of U.S. Pat. No. 9,318,160, Exhibit 2101, Jin Hee Cho, et al., "A 1.2V 64Gb 341GB/s HBM2 Stacked DRAM with Spiral Point-to-Point TSV Structure and Improved Bank Group Data Control," 2018 IEEE International Solid-State Circuits Conference.

Inter Partes Review of U.S. Pat. No. 10,949,339, Case No. IPR2022-00639, Petitioner's Authorized Reply to Patent Owner's Preliminary Response, filed Aug. 26, 2022.

Inter Partes Review of U.S. Pat. No. 10,949,339, Case No. IPR2022-00639, Petitioner's Reply to Patent Owner's Response, filed May 1, 2023.

Inter Partes Review of U.S. Pat. No. 10,949,339, Case No. IPR2022-00639, EX2015—Subramanian Supplemental Declaration IPR2022-00063, filed Jan. 26, 2023.

Inter Partes Review of U.S. Pat. No. 8,787,060, Case No. IPR2022-01428, Brogioli Deposition, Sep. 27, 2023.

Inter Partes Review of U.S. Pat. No. 8,787,060, Case No. IPR2022-01428, EX1063—Dictionary Deposition for "electrical", filed Oct. 19, 2023.

Inter Partes Review of U.S. Pat. No. 8,787,060, Case No. IPR2022-01428, EX1064—Dictionary Deposition for "communication", filed Oct. 19, 2023.

Inter Partes Review of U.S. Pat. No. 8,787,060, Case No. IPR2022-01428, Petitioner's Reply to Patent Owner Response, filed Oct. 19, 2023.

Inter Partes Review of U.S. Pat. No. 8,787,060, Case No. IPR2022-01428, Patent Owner's Sur Reply, filed Nov. 21, 2023.

USDC for Eastern District of Texas, Marshall Division, Civil Action No. 2:22-CV-00203-JRG-RSP, Claim Construction Order, issued Oct. 21, 2023.

Inter Partes Review of U.S. Pat. No. 8,787,060, Case No. IPR2022-01428, Exhibit 1053—Netlist's Technology Tutorial, filed Oct. 19, 2023.

Inter Partes Review of U.S. Pat. No. 8,787,060, Case No. IPR2022-01428, Hearing Transcript, Feb. 5, 2024.

Inter Partes Review of U.S. Pat. No. 8,787,060, Case No. IPR2022-01428, Patent Owner's Corrected Objections to Petitioner's Demonstratives, filed Jan. 9, 2024.

Inter Partes Review of U.S. Pat. No. 8,787,060, Case No. IPR2022-01428, Petitioner's Demonstratives for Oral Arguments, filed Jan. 8, 2024.

Inter Partes Review of U.S. Pat. No. 8,787,060, Case No. IPR2022-01428, Petitioner's Objections to Patent Owner's Demonstratives, filed Jan. 8, 2024.

Inter Partes Review of U.S. Pat. No. 8,787,060, Case No. IPR2022-01428, Netlist's Demonstratives, filed Jan. 8, 2024.

USDC Eastern District of Texas Marshall Division, Civil Action No. 2:22-cv-203-JRG-RSP, *Netlist, Inc. v. Micron Technology, Inc et al.*, Report & Recommendation, filed Jan. 11, 2024.

Yiran Li, et al., Exploiting Three-Dimensional (3D) Memory Stacking to Improve Image Data Access Efficiency for Motion Estimation Accelerators, Preprint submitted to Signal Processing: Image Communication, Mar. 31, 2010.

Li Jiang, et al., Modeling TSV Open Defects in 3D-Stacked DRAM, IEEE International Test Conference, Paper 6.1, 2010.

Xiaoxia Wu, Design Exploration for Three-dimensional Integrated Circuits (3D ICs), Doctor of Philosophy Dissertation, Oct. 25, 2010, pp. 30-32.

Paul Falkenstern, Electronic Design Automation Challenges in Three Dimensional Integrated Chips (3D ICs), Master of Science Master Thesis, Nov. 11, 2008, pp. 1-2.

R. Colin Johnson, Multidimensional Efforts Push Toward 3-D Chips, Electronic Engineering Times, Oct. 10, 2011, pp. 26-32.

Rakesh Anigundi, et al., Architecture Design Exploration of Three-Dimensional (3D) Integrated DRAM, 10th International Symposium on Quality Electronic Design, 2009.

Katsuyuki Sakuma, et al., 3D Chip-Stacking Technology With Through-Silicon Vias And Low-Volume Lead-free Interconnections, IBM Journal of Research and Development vol. 52 No. 6, Nov. 2008, pp. 611-622.

Qi Wu, et al., Impacts of Though-DRAM Vias in 3D Processor-DRAM Integrated Systems, 2009.

Wangyuan Zhang, et al., Exploring Phase Change Memory and 3D Die-Stacking for Power/Thermal Friendly, Fast and Durable Memory Architectures, 2009 18th International Conference on Parallel Architectures and Compilation Techniques, 2009.

Christian Weis, et al., Design Space Exploration for 3D-stacked DRAMs, 2011 Design, Automation & Test in Europe, 2011.

Ang-Chih Hsieh, et al., TSV Redundancy: Architecture and Design Issues in 3D IC, IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 2011.

John U. Knickerbocker, et al., Three-Dimensional Silicon Integration, IBM Journal of Research and Development vol. 52 No. 6, Nov. 2008, pp. 553-569.

Yung-Fa Chou, et al., Yield Enhancement by Bad-Die Recycling and Stacking With Though-Silicon Vias, IEEE Transactions On Very Large Scale Integration (VLSI) Systems, vol. 19, No. 8, Aug. 2011, pp. 1346-1356.

Igor Loi, et al., An Efficient Distributed Memory Interface For Many-Core Platform With 3D Stacked DRAM, 2010 Design, Automation & Test in Europe Conference & Exhibition, 2010.

Jung-Sik Kim, et al., A 1.2 v 12.8 GB/s 2 Gb mobile wide-I/O DRAM with 4 × 128 I/Os using TSV Based Stacking, IEEE Journal of Solid-State Circuits vol. 47 No. 1, Jan. 2012, pp. 107-116.

USDC Easter District of Texas, Marshall Division, *Netlist, Inc. v. Micron Technology, Inc et al.*, Civil Action No. 2:22-CV-00203-JRG-RSP, Claim Construction Order, Filed Oct. 22, 2023, 45 pages.

Inter Partes Review of U.S. Pat. No. 8,787,060, Case No. IPR2022-01428, Final Written Decision, Apr. 1, 2024.

Marco Facchini et al., "System-level power/performance evaluation of 3D stacked DRAMs for mobile applications," 2009 Design, Automation & Test in Europe Conference & Exhibition, Apr. 24, 2009.

X. Zhao et al., "Low-Power Clock Tree Design for Pre-Bond Testing of 3-D Stacked ICs," IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, vol. 30, Issue 5, Apr. 19, 2011.

Q. Wu et al., "Impacts of though-DRAM vias in 3D processor-DRAM integrated systems," 2009 IEEE International Conference on 3D System Integration, Sep. 30, 2009.

W. Ahmad et al., "Peak-to-Peak Ground Noise on a Power Distribution TSV Pair as a Function of Rise Time in 3-D Stack of Dies Interconnected Through TSVs," IEEE Transactions on Components, Packaging and Manufacturing Technology (vol. 1, Issue: 2, Feb. 2011).

R. Anigundi et al. "Architecture design exploration of three dimensional (3D) integrated DRAM," 2009 10th International Symposium on Quality Electronic Design, Mar. 18, 2009.

US 12,308,087 B2

Page 11

(56)

References Cited

OTHER PUBLICATIONS

Tak-Yung Kim et al., "Clock tree synthesis with pre-bond testability for 3D stacked IC Designs," IEEE Design Automation Conference, Jun. 18, 2010.

Uksong Kang et al., "8 Gb 3-D DDR3 DRAM Using Through-Silicon-Via Technology," IEEE Journal of Solid-State Circuits, vol. 45, No. 1, Jan. 2010.

Yi Zhao et al., "Cost-Effective TSV Grouping for Yield Improvement of 3D-ICs," retrieved from the Internet at: <https://typeset.io/pdf/cost-effective-tsv-grouping-for-yield-improvement-of-3d-ics-5gm5xcx0iq.pdf>, undated.

Ka Yao Tong et al., "A System Level Implementation of Rijndael on a Memory-slot based FPGA Card," retrieved from https://phwl.org/assets/papers/aes_fpt02.pdf, Jan. 2003.

Christian Plessl et al., "TKDM—A Reconfigurable Co-processor in a PC's Memory Slot," retrieved from from: <https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=615235db2bed5bb4823f9c3a9d383a92e1f602c7>, Dec. 17, 2003.

Maxim Integrated, "DC-DC Converter Tutorial," retrieved from https://pdfserv.maximintegrated.com/en/an/DCDC_Converter_Tutorial.pdf, Nov. 29, 2001.

Renesas Datasheet, ISL6537, "ACPI Regulator/Controller for Dual Channel DDR Memory Systems," Jul. 18, 2007.

Renesas Datasheet, ISL6534, "Dual PWM with Linear," Nov. 18, 2005.

Intersil Datasheet, ISL6532B, "ACPI Regulator/Controller for Dual Channel DDR Memory Systems," Jul. 18, 2004.

* cited by examiner




FIG. 1A




FIG. 1B




FIG. 2




FIG. 3





FIG. 4

U.S. Patent

May 20, 2025

Sheet 5 of 8

US 12,308,087 B2






FIG. 6A




FIG. 6B




FIG. 7




FIG. 8

US 12,308,087 B2

1

MEMORY PACKAGE HAVING STACKED ARRAY DIES AND REDUCED DRIVER LOAD

RELATED APPLICATION

This application is a continuation application of U.S. patent application Ser. No. 17/157,903, filed Jan. 25, 2021, which is a continuation application of U.S. patent application Ser. No. 16/412,308, filed May 14, 2019, now U.S. Pat. No. 10,902,886, which is a continuation application of U.S. patent application Ser. No. 15/602,099, filed May 22, 2017, now U.S. Pat. No. 10,602,099, which is a continuation application of U.S. patent application Ser. No. 15/095,288, filed Apr. 11, 2016, now U.S. Pat. No. 9,659,601 on May 23, 2017, which is a continuation of U.S. patent application Ser. No. 14/337,168, filed Jul. 21, 2014, now U.S. Pat. No. 9,318,160, which is a continuation of U.S. patent application Ser. No. 13/288,850, filed Nov. 3, 2011, now U.S. Pat. No. 8,787,060, which claims the benefit of priority under 35 U.S.C. § 119(e) of U.S. Provisional Patent Application No. 61/409,893, filed 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

Technical Field

The present disclosure relates to memory devices and memory modules. More specifically, the present disclosure relates to a memory package having stacked array dies and methods for reducing driver load in the memory package.

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

2

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 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 DRAM package operable to communicate with a system memory controller via a command/address bus and a data bus. The DRAM package comprises stacked array dies including at least a first array die and a second array die stacked over at least the first array die, and terminals including control terminals, via which the DRAM package receives command/address signals, and data terminals, via which the DRAM package receives or outputs data signals. The DRAM package further comprises control interconnects between respective ones of the control terminals and one or more array dies in the stacked array dies, first data interconnects between respective ones of the data terminals and the first array die, and second data interconnects distinct from the first data interconnects and between respective ones of the data terminals and the second array die. The DRAM package further comprises drivers coupled to the data terminals and configured to drive first data signals to the first array die via the first data interconnects and second data signals to the second array die via the second data interconnects, a respective signal of the first data signals being driven by one or more drivers having a first driver size, a respective signal of the second data signals being driven by one or more drivers having a second driver size different from the first driver size.

In certain embodiments, a method is provided for optimizing load in a memory package. The DRAM package comprises stacked array dies including at least a first array die and a second array die stacked over at least the first array die, and terminals including control terminals, via which the DRAM package receives command/address signals, and data terminals, via which the DRAM package receives or

US 12,308,087 B2

3

outputs data signals. The DRAM package further comprises control interconnects between respective ones of the control terminals and one or more array dies in the stacked array dies, first data interconnects between respective ones of the data terminals and the first array die, and second data interconnects distinct from the first data interconnects and between respective ones of the data terminals and the second array die. The method comprises selecting a first driver size to drive a first signal received at one of the data terminals to the first array die via one of the first data interconnects and selecting a second driver size to drive a second signal received at the one of the data terminals to the second array die via one of the second data interconnects. The second driver size is different from the first driver size.

In some embodiments, a respective one of the first data interconnects has a first load and a respective one of the second data interconnect has a second load larger than the first load. The second driver size is larger than the first driver size.

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.

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

4

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 web site, <http://www.intel.com/idf/index.htm>, which includes presentations and papers from the IDF2011 Intel Developer Forum including the keynote address given by Justin Rattner on Sep. 15, 2011. Additional examples of an HMC compatible with certain embodiments described herein are described by the Hybrid Memory Cube Consortium website, <http://www.hybridmemorycube.org>. The memory package **200** can include any type of memory package. For example, the memory package **200** may be a DRAM package, a SDRAM package, a flash memory package, or a DDR SDRAM package (e.g., DDR3, DDR4), to name a few. A memory module (not shown) may include one or more memory packages in accordance with the memory package **200**. Further, the memory package **200** may include input/output terminals (not shown) that are configured to be placed in electrical communication with circuitry of the memory module to transmit signals between the memory package **200** and a Memory Control Hub (MCH) (not shown).

The memory package **200** may include a plurality of array dies **210** (e.g., array dies **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

US 12,308,087 B2

5

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 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., 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

6

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

US 12,308,087 B2

7

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

8

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 signals 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

US 12,308,087 B2

9

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

10

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.

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

US 12,308,087 B2

11

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 any one of the conduits 332a to 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 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

12

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 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

US 12,308,087 B2

13

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).

$$CL=AD+S/M \quad (1)$$

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 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.

14

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

US 12,308,087 B2

15

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

16

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 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

US 12,308,087 B2

17

calculated using formula (1) as previously described above. In some embodiments, the first subset of array dies and the second subset of array dies may be selected to balance a load on a first driver and a load on a second driver based, at least in part, on the loads of the array dies **310** without considering the loads of the die interconnect segments from the first die interconnect and the second die interconnect.

The process **500** further comprises forming electrical connections between the first die interconnect and the first subset of array dies in an operational block **504**. The process **500** further comprises forming electrical connections between the second die interconnect and the second subset of array dies in an operational block **506**. Forming the electrical connections places the die interconnects in electrical communication with the respective subsets of array dies. In some embodiments, forming the electrical connections can comprise forming electrical connections between the die interconnects and at least one port from each array die of the respective subsets of array dies.

The process **500** further comprises selecting a driver size for a first driver at block **508** and selecting a driver size for a second driver at block **510**. Selecting the driver size can be based, at least in part, on the calculated load on the driver. Generally, the greater the load on the driver, the larger the driver is selected to drive a signal along, for example, a die interconnect. The size of the driver may be adjusted by the selection of the transistor size and/or number of transistors included in the driver. A larger driver often consumes more power than a smaller driver. Thus, in certain embodiments, balancing the loads on the drivers to reduce the load on each driver can reduce the power consumption of a memory package.

In some embodiments, the size of the first driver and the size of the second driver are both less than the size sufficient for a driver to drive a signal along a die interconnect to each of the array dies **310** (e.g., with less than a predetermined or threshold signal degradation). The threshold signal degradation can be based on any one or more characteristics of a signal. For example, the threshold signal degradation can be based on the amplitude of the signal, the frequency of the signal, the noise distortion included in or introduced into the signal, or the shape of the signal, to name a few.

The process **500** further comprises forming electrical connections between the first die interconnect and the first driver in an operational block **512**. Similarly, the process **500** further comprises forming electrical connections between the second die interconnect and the second driver in an operational block **514**. Forming the electrical connections places the die interconnects in electrical communication with the respective drivers. In some embodiments, forming the electrical connections can comprise forming electrical connections between the die interconnects and a data conduit that includes the respective driver.

Advantageously, certain embodiments of the present disclosure reduce the load on each conduit and on each corresponding driver. In certain embodiments, reducing the load on a driver may increase the speed of data transfer between the array dies and other components in a computer system (e.g. a MCH or a processor). Further, reducing the load on a driver may result in reduced power consumption by the driver and consequently, by the memory package. In addition, certain embodiments of the present disclosure may minimize current switching noise.

Operational Modes

A proposed three-dimensional stacking (3DS) standard for dual in-line memory modules (DIMMs) being considered by the Joint Electron Devices Engineering Council

18

(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; and 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., 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

US 12,308,087 B2

19

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; and 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

20

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 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

US 12,308,087 B2

21

814. The command/address bus **814**, although illustrated as a single line, may include as many lines or signal conduits as the number of bits of the command/address signal.

Further, the register device **812** may generate data path control signals, which can be provided to a control die **822**, or isolation device, of a memory package **820** via one or more data path control lines **816**. In certain embodiments, the control dies **822** can include some or all of the embodiments described above with respect to the control dies **230** and **330**. Moreover, in certain embodiments, the memory packages **820** can include some or all of the embodiments described above with respect to the memory package **200** and **300**. In general, the memory module **812** may include one or more memory packages **820**.

In certain embodiments, each control die **822**, or isolation device, may be capable of address pass-through. Address pass-through, in some cases, enables the control die **822** to provide an address signal to one or more array dies **824** without decoding the address signal. This is possible, in some implementations, because the address signal received by the control die **822** is not encoded.

Some implementations of the control dies **822** include a plurality of command/address buffers (not shown). These buffers may comprise latches. In certain embodiments, the buffers are configured to hold command/address signals to control the timing of command/address signals. In some cases, controlling the timing of the command/address signals may reduce or slow signal degradation. In some implementations, the control dies **822** include a plurality of data buffers, which may control the timing of data signals to reduce or slow signal degradation. Further, the control dies **822** may include a data path control circuit (not shown) that is configured to control command/address time slots and data bus time slots. Controlling the command/address time slots and the data bus time slots enables the control dies **822** to reduce or prevent signal collisions caused by multiple memory packages **820** sharing the data path control lines **816** and the data bus **818**. In some implementations, the data path control circuit may be separate from the control die **822**.

Each of the control dies **822** may be configured to receive data signals from the memory control hub **802** via the data bus **818**. Further, the control dies **822** may provide data signals to the memory control hub **802** via the data bus **818**. The control dies **822** may also receive data path control signals and/or command/address signals from the register device **812** via the data path control lines **816**.

As illustrated in FIG. 8, each memory package **820** may include one or more array dies **824** operatively coupled to the control die **822**. The array dies **824** may be configured to receive data signals from the control die **822**. As with the array dies **210** and **310**, the array dies **824** may include any type of Random Access Memory (RAM) die. For example, the array dies **824** may include DDR DRAM, SDRAM, flash memory, or SRAM, to name a few. Further, if the memory module **810** is utilized for graphics, the array dies **824** may include GDDR SDRAM or any other type of graphics memory. In addition, the array dies **824** may be configured in a stack arrangement as illustrated in FIG. 8. Alternatively, the array dies **824** may be arranged in a planar configuration or a hybrid configuration utilizing both a planar and a stack arrangement.

In some embodiments, the memory module **810** is selectively configurable into two operational modes. In the first operational mode, the register device **812** generates data path control signals provided to the control dies **822** via the data path control lines **816**. The control dies **822** may decode

22

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 embodiments of, and examples for, the invention are described above for illustrative purposes, various equivalent

US 12,308,087 B2

23

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 dynamic random access memory (DRAM) package, comprising:

stacked DRAM dies including at least a first plurality of DRAM dies and a second plurality of DRAM dies, each DRAM die of the stacked DRAM dies including C/A ports, data ports and DRAM memory cells, wherein the each DRAM die is configurable to transfer data between the data ports and the DRAM memory cells; terminals including command and/or address (C/A) terminals and data terminals, wherein the DRAM package is configured to receive C/A signals via the C/A terminals and is further configured to receive or output data signals via the data terminals in response to the (C/A) signals, wherein the DRAM package is configured to output first data signals in response to a first set of C/A signals associated with a memory read operation and to receive second data signals in response to a second set of C/A signals associated with a memory write operation;

die interconnects including C/A interconnects and data interconnects, the C/A interconnects including at least first C/A interconnects and second C/A interconnects, the first C/A interconnects configured to conduct the first set of C/A signals and the second set of C/A signals, the data interconnects including at least first

24

data interconnects and second data interconnects, the first data interconnects configured to conduct the first data signals and the second data signals, each of the die interconnects including one or more through silicon vias (TSVs) in one or more DRAM dies in the stacked DRAM dies and configured to conduct signals to and/or from the one or more DRAM dies in the stacked DRAM dies through the one or more TSVs;

a control die coupled between the terminals and the stacked DRAM dies, the control die including conduits, the conduits including C/A conduits and data conduits, the C/A conduits including at least first C/A conduits coupled to respective ones of the first C/A interconnects and second C/A conduits coupled to respective ones of the second C/A interconnects, the data conduits including at least first data conduits coupled to respective ones of the first data interconnects and second data conduits coupled to respective ones of the second data interconnects;

wherein a first C/A interconnect of the first C/A interconnects is in electrical communication with corresponding C/A ports on the first plurality of DRAM dies and not in electrical communication with any C/A port on any of the second plurality of DRAM dies;

wherein a second C/A interconnect of the second C/A interconnects is in electrical communication with corresponding C/A ports on the second plurality of DRAM dies and not in electrical communication with any C/A port on any of the first plurality of DRAM dies;

wherein a first data interconnect of the first data interconnects is in electrical communication with corresponding data ports on the first plurality of DRAM dies and not in electrical communication with any data port on any of the second plurality of DRAM dies, each of the first data interconnects including a first respective set of TSVs, the first respective set of TSVs including a TSV in each DRAM die of the first plurality of DRAM dies and at least one TSV in at least one DRAM die of the second plurality of DRAM dies, wherein the TSV in the each DRAM die of the first plurality of DRAM dies is in electrical communication with a corresponding data port on the each DRAM die, and wherein the at least one TSV in the at least one DRAM die of the second plurality of DRAM dies is not in electrical communication with any data port on the at least one DRAM die;

wherein a second data interconnect of the second data interconnects is in electrical communication with corresponding data ports on the second plurality of DRAM dies and not in electrical communication with any data port on any of the first plurality of DRAM dies;

wherein a first conduit of the first data conduits is coupled between the first data interconnect and a first data terminal of the data terminals, and a second conduit of the second data conduits is coupled between the second data interconnect and the first data terminal;

wherein the control die further includes control logic configurable to control respective states of the first and second conduits in response to one or more C/A signals received via one or more of the C/A terminals, wherein the one or more C/A signals do not include any chip select signal;

wherein the die interconnects further include first unidirectional interconnects configured to conduct signals from one or more DRAM dies of the stacked DRAM dies to the control die and not configured to conduct any signal from the control die to any of the stacked DRAM dies;

US 12,308,087 B2

25

wherein the die interconnects further include second unidirectional interconnects configured to conduct signals from the control die to one or more DRAM dies of the stacked DRAM dies and not configured conduct any signal from any of the stacked DRAM dies to the control die;

wherein the control die is configured to receive signals from one or more DRAM dies of the stacked DRAM dies via the first unidirectional interconnects and is not configured to drive any signal to any of the stacked DRAM dies via any of the first unidirectional interconnects;

wherein the control die is configured to drive signals to one or more DRAM dies of the stacked DRAM dies via the second unidirectional interconnects and is not configured to receive any signal from any of the stacked DRAM dies via any of the second unidirectional interconnects; and

wherein the control die is configured to, in response to the first set of C/A signals, receive first signals associated with the memory read operation from a DRAM die of the stacked DRAM dies via the first unidirectional interconnects, and in response to the second set of C/A signals, drive second signals associated with the memory write operation to one or more DRAM dies of the stacked DRAM dies via the second unidirectional interconnects.

2. The DRAM package of claim 1, wherein the stacked DRAM dies include at least 8 stacked DRAM dies.

3. The DRAM package of claim 2, wherein the stacked DRAM dies include 16 stacked DRAM dies.

4. The DRAM package of claim 1, wherein:

the first conduit includes a first driver configured to have a first driver size and the second conduit includes a second driver configured to have a second driver size; and the second driver size is different from the first driver size.

5. The DRAM package of claim 1, wherein the first data signals include 32 or 64 signals conducted via the first data interconnects in parallel and subsequently output by the DRAM package in parallel, and wherein the second data signals include 32 or 64 signals received by the DRAM package in parallel and subsequently conducted via the first data interconnects in parallel.

6. The DRAM package of claim 1, wherein:

the first set of C/A signals include one or more signals to select one DRAM die among the first plurality of DRAM dies to output the first data signals; and

the second set of C/A signals include one or more signals to select one DRAM die among the first plurality of DRAM dies to receive the second data signals;

wherein the control die does not include emulation logic that performs emulation causing two or more DRAM dies of the stacked DRAM dies to emulate a single DRAM die having a larger capacity than each of the two or more DRAM dies.

7. The DRAM package of claim 1, wherein the control die is further configured to control timing of command/address signals driven to the stacked DRAM dies and/or timing of data signals communicated between the stacked DRAM dies and the data terminals to prevent signal collisions.

8. The DRAM package of claim 1, wherein the at least one DRAM die of the second plurality of DRAM dies has at least one electrical connection leading from the at least one TSV, the at least one electrical connection is not electrically connected to any data port on the at least one DRAM die.

26

9. A method, comprising:

at a dynamic random access memory (DRAM) package including stacked DRAM dies, the stacked DRAM dies including a first plurality of DRAM dies and a second plurality of DRAM dies, each DRAM die of the stacked DRAM dies including command and/or address (C/A) ports, data ports and DRAM memory cells, wherein the each DRAM die is configurable to transfer data between the data ports and the DRAM memory cells in response to C/A signals received at the C/A ports, the DRAM package further including an interface, via which the DRAM package receives C/A signals, and via which the DRAM package receives or outputs data signals in response to C/A signals received via the interface, the interface including terminals, the DRAM package further including die interconnects, each of the die interconnects including one or more through silicon vias (TSVs) in one or more DRAM dies in the stacked DRAM dies, the die interconnects including C/A interconnects, data interconnects, first unidirectional interconnects and second unidirectional interconnects, the C/A interconnects including first C/A interconnects and second C/A interconnects, the data interconnects including first data interconnects and second data interconnects, receiving via the interface a first set of C/A signals for a memory read operation and a second set of C/A signals for a memory write operation, the first set of C/A signals including one or more first signals to select a first DRAM die in the first plurality of DRAM dies to respond to the first set of C/A signals, and the second set of C/A signals include one or more second signals to select a second DRAM die in the first plurality of stacked DRAM dies to respond to the second set of C/A signals;

in response to the first set of C/A signals, driving the first set of C/A signals to C/A ports of the first plurality of DRAM dies via the first C/A interconnects, receiving from the first DRAM die first data signals associated with the first set of C/A signals via the first data interconnects, driving the first data signals to the interface, and receiving from the first DRAM die one or more signals associated with the first set of C/A signals via one or more of the first unidirectional interconnects;

in response to the second set of C/A signals, driving the second set of C/A signals to C/A ports of the first plurality of DRAM dies via the first C/A interconnects, receiving second signals associated with the second set of C/A signals from the interface, driving the second data signals to the first plurality of DRAM dies via the first data interconnects and not to any of the second plurality of DRAM dies via any of the data interconnects, and driving one or more signals associated with the second set of C/A signals to the first plurality of DRAM dies via one or more of the second unidirectional interconnects;

wherein each of the first data interconnects is in electrical communication with corresponding data ports on a first plurality of DRAM dies of the stacked DRAM dies and not in electrical communication with any data port on any other DRAM die of the stacked DRAM dies, each of the first data interconnects including a respective set of TSVs, the respective set of TSVs including a TSV in each DRAM die of the first plurality of DRAM dies and at least one additional TSV in at least one DRAM die of the second plurality of DRAM dies, wherein the TSV in the each DRAM die of the first plurality of DRAM dies is in electrical communication with a corresponding data port on the each DRAM die, and

US 12,308,087 B2

27

wherein the at least one additional TSV in the at least one DRAM die of the second plurality of DRAM dies is not in electrical communication with any data port on the at least one DRAM die;

wherein each of the second data interconnects is in electrical communication with corresponding data ports on the second plurality of DRAM dies and not in electrical communication with any data port on any of the first plurality of DRAM dies;

wherein a first C/A interconnect of the first C/A interconnects is in electrical communication with corresponding C/A ports on the first plurality of DRAM dies and not in electrical communication with any C/A port on any of the second plurality of DRAM dies;

wherein a second C/A interconnect of the second C/A interconnects is in electrical communication with corresponding C/A ports on the second plurality of DRAM dies and not in electrical communication with any C/A port on any of the first plurality of DRAM dies;

wherein the first unidirectional interconnects are not configured to conduct any signals to any of the stacked DRAM dies and the second unidirectional interconnects are not configured to conduct any signals from any of the stacked DRAM dies;

wherein the first data signals include at least 32 signals that are received via the first data interconnects in parallel and subsequently driven to the interface in parallel;

and wherein the second data signals include at least 32 signals that are received from the interface in parallel and subsequently driven to the first plurality of DRAM dies in parallel.

10. The method of claim 9, wherein the stacked DRAM dies include at least 8 stacked DRAM dies.

11. The method of claim 9, wherein the first data signals include 32 or 64 signals conducted via the first data interconnects in parallel and subsequently output by the DRAM package via the interface in parallel, and wherein the second data signals include 32 or 64 signals received by the DRAM package via the interface in parallel and subsequently conducted via the first data interconnects in parallel.

12. The method of claim 9, wherein the first set of C/A signals and the second set of C/A signals are consecutive sets of C/A signals, the method further comprising:

- enabling first drivers to receive and drive the first data signals in response to the first set of C/A signals;
- enabling second drivers to receive and drive the second data signals in response to the second set of C/A signals; and
- maintaining each of the first drivers and the second drivers in a non-driving state after driving of the first data signals by the first drivers and before receiving of the second data signals by the second drivers;

wherein any driver of the first and second drivers in a non-driving state is not configured to receive or drive any signal.

13. A dynamic random access memory (DRAM) package, comprising:

- a plurality of stacked DRAM dies, each DRAM die of the plurality of stacked DRAM dies including command/address (C/A) ports, data ports and DRAM memory cells, wherein the each DRAM die is configurable to transfer data between the data ports and the DRAM memory cells in response to C/A signals received via the C/A ports, each of the data ports being a bi-directional data port;

28

an interface, via which the DRAM package receives command/address (C/A) signals, and via which the DRAM package receives or outputs data signals, the interface including terminals, the terminals including C/A terminals and data terminals;

die interconnects, each of the die interconnects including one or more through silicon vias (TSVs) in one or more DRAM dies in the plurality of stacked DRAM dies;

and a control die coupled between the interface and the plurality of stacked DRAM dies, the control die including signal conduits coupled, respectively, to the die interconnects;

wherein:

- the die interconnects include a first set of C/A interconnects, a second set of C/A interconnects, first data interconnects, second data interconnects, first unidirectional interconnects, and second unidirectional interconnects;
- each of the first data interconnects is a bi-directional interconnect configured to be in electrical communication with a corresponding data port on each of a first subset of the plurality of stacked DRAM dies and is not configured to be in electrical communication with any data port on any of a second subset of the plurality of stacked DRAM dies;
- each of the second data interconnects is a bi-directional interconnect configured to be in electrical communication with a corresponding data port of each of the second subset of the plurality of stacked DRAM dies and is not configured to be in electrical communication with any data port of the first subset of the plurality of stacked DRAM dies, each subset of the first subset and the second subset of the plurality of stacked DRAM dies including two or more DRAM dies;
- the control die is configurable to drive a first set of C/A signals received by the DRAM package via the interface to the first subset of the plurality of stacked DRAM dies via the first C/A interconnects, and is further configurable to drive a second set of C/A signals received by the DRAM package via the interface to the second subset of the plurality of stacked DRAM dies via the second C/A interconnects;
- a first C/A interconnect of the first C/A interconnects is in electrical communication with corresponding C/A ports on the first subset of the plurality of stacked DRAM dies and not in electrical communication with any C/A port on any of the second subset of the plurality of stacked DRAM dies;
- wherein a second C/A interconnect of the second C/A interconnects is in electrical communication with corresponding C/A ports on the second subset of the plurality of stacked DRAM dies and not in electrical communication with any C/A port on any of the first subset of the plurality of stacked DRAM dies;
- the first set of C/A signals include one or more first signals to select a first DRAM die in the first subset of the plurality of stacked DRAM dies to respond to the first set of C/A signals, and the second set of C/A signals include one or more second signals to select a second DRAM die in the second subset of the plurality of stacked DRAM dies to respond to the second set of C/A signals;
- the control die is further configurable to drive first data signals associated with the first set of C/A signals between the first data interconnects and the interface and to drive second data signals associated with the

US 12,308,087 B2

29

second set of C/A signals between the second data interconnects and the interface;

the control die is further configurable to, in response to the first set of C/A signals being for a memory read operation, receive one or more first signals associated with the memory read operation from one or more of the first unidirectional interconnects, and in response to the second set of C/A signals being for a memory write operation, drive one or more second signals associated with the memory write operation to one or more of the second unidirectional interconnects;

the one or more of the first unidirectional interconnects are configured to conduct the one or more first signals from the first DRAM die to the control die and not configured to conduct any signal from the control die to any of the plurality of stacked DRAM dies;

the one or more of the second unidirectional interconnects are configured to conduct the one or more second signals from the control die to the second DRAM die and not configured to conduct any signal from any of the plurality of stacked DRAM dies to the control die;

the control die is configurable to receive the one or more first signals from the first DRAM die via the one or more of the first unidirectional interconnects and is not configurable to drive any signal to any of the plurality of stacked DRAM dies via any of the first unidirectional interconnects;

the control die is configurable to drive the one or more second signals to the second DRAM die one or more DRAM dies of the plurality of stacked DRAM dies via the one or more of the second unidirectional interconnects and is not configurable to receive any signal from any of the plurality of stacked DRAM dies via any of the second unidirectional interconnects;

each respective data interconnect of the first data interconnects includes a respective set of through silicon vias (TSVs), the respective set of TSVs including a first subset of TSVs in respective DRAM dies in the first subset of DRAM dies and at least one second TSV in at least one second DRAM die of the second subset of plurality of stacked DRAM dies;

a respective TSV of the first subset of TSVs in a respective DRAM die of the first subset of DRAM dies is in electrical communication with a corresponding data port on the respective DRAM die in the first subset of DRAM dies; and

each of the at least one second TSV in a corresponding DRAM die of the at least one second DRAM die is not in electrical communication with and is not connected to any data port on the corresponding DRAM die.

14. The DRAM package of claim **13**, wherein the plurality of stacked DRAM dies include at least 8 stacked DRAM dies.

15. The DRAM package of claim **14**, wherein the plurality of stacked DRAM dies include 16 stacked DRAM dies.

16. The DRAM package of claim **13**, wherein the each of the at least one second TSV in the corresponding DRAM die of the at least one second DRAM die is connected to one or more electrical connections leading from the each of the at least one second TSV, the one or more electrical connections are not configured to provide any electrical communication with memory cells on the corresponding DRAM die.

17. The DRAM package of claim **13**, wherein the signal conduits in the control die include first data conduits coupled to respective ones of the first data interconnects and second data conduits coupled to respective ones of the second data interconnects, wherein the first set of C/A signals and the

30

second set of C/A signals are consecutive sets of C/A signals, wherein the control die further includes logic configured to control states of the data conduits in response to C/A signals received via the interface, whereby:

the first data conduits are enabled to receive and drive the first data signals in response to the first set of C/A signals;

the second conduits are enabled to receive and drive the second data signals in response to the second set of C/A signals;

each of the first data conduits are placed in a non-driving state after driving the first data signals;

each of the second data conduits are placed in a non-driving state after driving the second data signals; and

data conduit in the first and second data conduits in a non-driving state are not configured to receive or drive any signals.

18. The DRAM package of claim **13**, wherein the terminals include a first terminal coupled to a first die interconnect of the first data interconnects via a first conduit of the signal conduits and to a second die interconnect of the second data interconnects via a second conduit of the signal conduits, wherein the control die includes logic configurable to control respective states of the first conduit and the second conduit in response to one or more C/A signals received via one or more second terminals of the terminals, wherein the one or more C/A signals do not include any chip select signal.

19. The DRAM package of claim **13**, wherein the first data signals include 32 or 64 signals received or output by the DRAM package in parallel and conducted via the first data interconnects in parallel, and wherein the second data signals include 32 or 64 signals received or output by the DRAM package in parallel and conducted via the second data interconnects in parallel.

20. A dynamic random access memory (DRAM) package, comprising:

stacked DRAM dies, each DRAM die of the stacked DRAM dies including data ports and DRAM memory cells, wherein the each DRAM die is configurable to transfer data between the data ports and the DRAM memory cells;

terminals including command and/or address (C/A) terminals and data terminals, wherein the DRAM package is configured to receive C/A signals via the C/A terminals and receives or outputs data signals via the data terminals in response to the (C/A) signals, wherein the DRAM package is configured to output a first n-bit-wide data signal via the data terminals in response to a set of C/A signals associated with a memory read operation and to receive a second n-bit-wide data signal via the data terminals in response to a second set of C/A signals associated with a memory write operation, n being 32 or 64;

die interconnects including C/A interconnects and data interconnects, the C/A interconnects including a first set of C/A interconnects configured to conduct the first set of C/A signals and the second set of C/A signals, the data interconnects including a first set of data interconnects configured to conduct the first n-bit-wide data signal and the second n-bit-wide data signal, wherein each die interconnect of the die interconnects includes one or more TSVs in one or more DRAM dies in the stacked DRAM dies and is configured to conduct signals to and/or from the one or more DRAM dies through the one or more TSVs; and

US 12,308,087 B2

31

a control die coupled between the terminals and the stacked DRAM dies, the control die including conduits coupled to respective ones of the die interconnects; wherein the data interconnects include at least a first data interconnect and a second data interconnect, the first data interconnect configured to be in electrical communication with a corresponding data port of each of a first plurality of DRAM dies of the stacked DRAM dies and not in electrical communication with any data port of any of a second plurality of DRAM dies of the stacked DRAM dies, the second data interconnect configured to be in electrical communication with a corresponding data port of each of a second plurality of DRAM dies and not in electrical communication with any data port of any of the first plurality of DRAM dies; wherein the conduits include at least a first data conduit and a second data conduit, the first data conduits being coupled between the first data interconnect and a first data terminal of the data terminals, and the second data conduit being coupled between the second data interconnect and the first data terminal;

wherein the control die further includes control logic configurable to control respective states of the first data conduit and the second data conduit in response to one or more C/A signals received via one or more of the C/A terminals, wherein the one or more C/A signals do not include any chip select signal;

wherein the die interconnects further include first unidirectional interconnects and second unidirectional interconnects;

wherein the control die is configured to receive signals from the first unidirectional interconnects and to drive signals to the second unidirectional interconnects;

wherein the control die is not configured to drive any signal to any of the first unidirectional interconnects or to receive any signal from any of the second unidirectional interconnects; and

wherein, in response to the first set of C/A signals, the control die is configured to receive one or more first signals associated with the memory read operation from one of the one or more DRAM dies via one or more of the first unidirectional interconnects, and wherein, in response to the second set of C/A signals, the control die is configured to drive one or more second signals associated with the memory write operation to the one or more DRAM dies via one or more of the second unidirectional interconnects;

wherein the control die does not include emulation logic that performs emulation causing two or more DRAM dies of the stacked DRAM dies to emulate a single DRAM die having a larger capacity than each of the two or more DRAM dies.

21. The DRAM package of claim 20, wherein the stacked DRAM dies include at least 8 stacked DRAM dies.

22. The DRAM package of claim 20, wherein: the first set of C/A signals include one or more signals to select one DRAM die among the first plurality of DRAM dies to output the first n-bit-wide data; and the second set of C/A signals include one or more signals to select one DRAM die among the first plurality of DRAM dies to receive the second n-bit-wide data.

23. A dynamic random access memory (DRAM) package, comprising:

stacked DRAM dies, each DRAM die of the stacked DRAM dies including data ports and DRAM memory

32

cells, wherein the each DRAM die is configurable to transfer data between the data ports and the DRAM memory cells;

terminals including command and/or address (C/A) terminals, data terminals, and power supply terminals, wherein the DRAM package is configured to receive C/A signals via the C/A terminals and is further configured to receive or output data signals via the data terminals in response to the (C/A) signals, wherein the DRAM package is configured to output a first n-bit-wide data signal via the data terminals in response to a first set of C/A signals associated with a memory read operation and to receive a second n-bit-wide data signal via the data terminals in response to a second set of C/A signals associated with a memory write operation, n being 32 or 64;

die interconnects including C/A interconnects and data interconnects, the C/A interconnects including a first set of C/A interconnects configured to conduct the first set of C/A signals and the second set of C/A signals, the data interconnects including a first set of data interconnects configured to conduct the first n-bit-wide data signal and the second n-bit-wide data signal, wherein each die interconnect of the die interconnects includes one or more TSVs in one or more DRAM dies in the stacked DRAM dies and is configured to conduct signals to and/or from the one or more DRAM dies through the one or more TSVs; and

a control die coupled between the terminals and the stacked DRAM dies, the control die including conduits coupled to respective ones of the die interconnects;

wherein the data interconnects include at least a first data interconnect and a second data interconnect, the first data interconnect configured to be in electrical communication with a corresponding data port of each of a first plurality of DRAM dies of the stacked DRAM dies and not in electrical communication with any data port of any of a second plurality of DRAM dies of the stacked DRAM dies, the second data interconnect configured to be in electrical communication with a corresponding data port of each of a second plurality of DRAM dies and not in electrical communication with any data port of any of the first plurality of DRAM dies;

wherein the conduits include at least a first data conduit and a second data conduit, the first data conduits being coupled between the first data interconnect and a first data terminal of the data terminals, and the second data conduit being coupled between the second data interconnect and the first data terminal;

wherein the control die further includes control logic configurable to control respective states of the first data conduit and the second data conduit in response to one or more C/A signals received via one or more of the C/A terminals, wherein the one or more C/A signals do not include any chip select signal;

wherein the die interconnects further include first unidirectional interconnects and second unidirectional interconnects;

wherein the control die is configured to receive signals from the first unidirectional interconnects and to drive signals to the second unidirectional interconnects;

wherein the control die is not configured to drive any signal to any of the first unidirectional interconnects or to receive any signal from any of the second unidirectional interconnects;

wherein, in response to the first set of C/A signals, the control die is configured to receive one or more first

US 12,308,087 B2

33

signals associated with the memory read operation from one of the one or more DRAM dies via one or more of the first unidirectional interconnects, and wherein, in response to the second set of C/A signals, the control die is configured to drive one or more second signals associated with the memory write operation to the one or more DRAM dies via one or more of the second unidirectional interconnects;

wherein the power supply interconnects include a first supply voltage interconnect and a second supply voltage interconnect, wherein the first supply voltage interconnect is configured to conduct a supply voltage received at a power supply terminal to each of the first plurality of DRAM dies and not to any of the second plurality of DRAM dies, and wherein the second supply voltage interconnect is configured to conduct the supply voltage received at the power supply terminal to each of the second plurality of DRAM dies and not to any of the first plurality of DRAM dies; and

wherein the control die does not include emulation logic that performs emulation causing two or more DRAM dies of the stacked DRAM dies to emulate a single DRAM die having a larger capacity than each of the two or more DRAM dies.

24. A dynamic random access memory (DRAM) package, comprising:

stacked DRAM dies, wherein each DRAM die of the stacked DRAM dies includes signal ports and DRAM memory cells, and is configurable to receive C/A signals via a first portion of the signal ports and to transfer data between the DRAM memory cells and a second portion of the signal ports;

an interface via which the DRAM package receives C/A signals and receives or outputs data signals;

a control die coupled between the interface and the stacked DRAM dies, the control die including signal conduits and logic configured to control at least some of the signal conduits; and

die interconnects, each die interconnect of the die interconnects coupled between at least one corresponding conduit in the control die and one corresponding signal port in each of a subset of the stacked DRAM dies, the each die interconnect including through silicon vias (TSVs) in at least some of the stacked DRAM dies;

wherein:

the signal conduits include first conduits, second conduits, third conduits and fourth conduits;

the die interconnects include first die interconnects coupled to respective ones of the first conduits, second die interconnects coupled to respective ones of the second conduits, third die interconnects coupled to respective ones of the third conduits, and fourth die interconnects coupled to respective ones of the fourth conduits;

each of the first conduits being a unidirectional conduit configurable to drive a signal to a corresponding first die interconnect and not configurable to receive any signal from the corresponding first die interconnect;

each of the second conduits being a bidirectional conduit configurable to drive a signal to a corresponding one of the second die interconnects or to receive a signal from the corresponding second die interconnect;

each of the third conduits being a unidirectional conduit configurable to drive a signal received from a corresponding third die interconnect to the interface and not configurable to drive any signal to the corresponding third die interconnect;

34

each of the fourth conduits being a unidirectional conduit configurable to drive a signal received from the interface to a corresponding fourth die interconnect and not configurable to receive any signal from the corresponding fourth die interconnect;

the DRAM package is configurable to receive via the interface a first set of C/A signals for a memory read operation, a second set of C/A signals for a memory write operation, a third set of C/A signals for a memory read operation and a fourth set of C/A signals for a memory write operation;

in response to the first set of C/A signals, a first subset of the first conduits is configured to drive the first set of C/A signals to a first subset of the first die interconnects, a first subset of the second conduits is configured to receive first signals associated with the first set of C/A signals via a first subset of the second die interconnects and to drive the first signals to the interface; and a first subset of the third conduits is configured to receive one or more second signals associated with the first set of C/A signals from a first subset of the third die interconnect and drive the one or more second signals to the interface;

in response to the second set of C/A signals, the first subset of the first conduits is configured to drive the second set of C/A signals to the first subset of the first die interconnects, the first subset of the second conduits is configured to receive third signals associated with the second set of C/A signals from the interface and to drive the third signals to the first subset of the second die interconnects; and a first subset of the fourth conduits is configured to receive one or more fourth signals associated with the second set of C/A signals from the interface and to drive the one or more fourth signals to a first subset of the fourth die interconnect;

in response to the third set of C/A signals, a second subset of the first conduits is configured to drive the third set of C/A signals to a second subset of the first die interconnects, a second subset of the second conduits is configured to receive fifth signals associated with the second set of C/A signals via a second subset of the second die interconnects and to drive the fifth signals to the interface; and a second subset of the third conduits is configured to receive one or more sixth signals associated with the second set of C/A signals from a second subset of the third die interconnect and drive the one or more sixth signals to the interface;

in response to the fourth set of C/A signals, the second subset of the first conduits is configured to drive the fourth set of C/A signals to the second subset of the first die interconnects, the second subset of the second conduits is configured to receive seventh signals associated with the fourth set of C/A signals from the interface and to drive the seventh signals to the second subset of the second die interconnects; and a second subset of the fourth conduits is configured to receive one or more eighth signals associated with the second set of C/A signals from the interface and to drive the one or more eighth signals to a second subset of the fourth die interconnect;

each conduit in the first subset of the first, second, third or fourth conduits is distinct from any conduit in the second subset of the first, second, third or fourth conduits;

each die interconnect in the first subset of the first, second, third or fourth die interconnects is distinct from any die

US 12,308,087 B2

35

interconnect in the second subset of the first, second, third or fourth die interconnects;
the first signals driven by the first subset of the second conduits in response to first set of C/A signals include at least 32 signals driven in parallel;
the third signals driven by the first subset of the second conduits in response to second set of C/A signals include at least 32 signals driven in parallel;
the fifth signals driven by the second subset of the second conduits in response to third set of C/A signals include at least 32 signals driven in parallel; and
the sixth signals driven by the second subset of the second conduits in response to fourth set of C/A signals include at least 32 signals driven in parallel.

25. The DRAM package of claim 24, wherein the stacked DRAM dies include a first subset of DRAM dies and a second subset of DRAM dies, wherein each of the first subset of the second die interconnects is configured to be in electrical communication with a corresponding signal port of each of a first subset of DRAM dies and not in electrical communication with any signal port of any of the second subset of DRAM dies, wherein each of the second subset of the second die interconnects is configured to be in electrical communication with a corresponding signal port of each of the second subset of DRAM dies and not in electrical communication with any signal port of any of the first subset of DRAM dies, and wherein:

the first set of C/A signals include one or more signals selecting a DRAM die in the first subset of DRAM dies to output the first signals and the one or more second signals;
the second set of C/A signals include one or more signals selecting a DRAM die in the first subset of DRAM dies to receive the third signals and the one or more fourth signals;
the third set of C/A signals include one or more signals selecting a DRAM die in the second subset of DRAM dies to output the fifth signals and the one or more sixth signals; the fourth set of C/A signals include one or

36

more signals selecting a DRAM die in the second subset of DRAM dies to receive the seventh signals and the one or more eighth signals;
the interface includes terminals via which the DRAM package receives and/or outputs signals, the terminals including a first terminal coupled to a first die interconnect of the first subset of the second die interconnects via a first conduit of the signal conduits and to a second die interconnect of the second subset of the second die interconnects via a second conduit of the signal conduits; and

the logic is configurable to control respective states of the first and second conduits in response to one or more C/A signals received via one or more second terminals in the interface, wherein the one or more C/A signals do not include any chip select signal.

26. The DRAM package of claim 25, wherein the first set of C/A signals and the second set of C/A signals are consecutive sets of C/A signals, wherein the control die further includes logic configured to control states of the data conduits in response to C/A signals received via the interface, whereby:

the first subset of the second conduits is enabled to receive and drive the first signals in response to the first set of C/A signals;

the first subset of the second conduits is enabled to receive and drive the third signals in response to the second set of C/A signals; and

the first subset of the second conduits is maintained at a non-driving state after driving the first signals and before receiving the third signals;

a conduit in the first subset of the second conduits or the second subset of the second conduits in a non-driving state is not configured to receive or drive any signal.

27. The DRAM package of claim 25, wherein the stacked DRAM dies include at least 8 stacked DRAM dies.

28. The DRAM package of claim 27, wherein the stacked DRAM dies include 16 stacked DRAM dies.

* * * * *

EXHIBIT 2

(12) **United States Patent**
Lee et al.

(10) **Patent No.:** **US 10,025,731 B1**
(45) **Date of Patent:** **Jul. 17, 2018**

(54) **MEMORY MODULE AND CIRCUIT PROVIDING LOAD ISOLATION AND NOISE REDUCTION**

(56) **References Cited**

U.S. PATENT DOCUMENTS

(71) Applicant: **Netlist, Inc.**, Irvine, CA (US)

4,739,473 A 4/1988 Ng
5,313,624 A 5/1994 Harriman
(Continued)

(72) Inventors: **Hyun Lee**, Ladera Ranch, CA (US);
Jayesh R. Bhakta, Cerritos, CA (US);
Jeffrey C. Solomon, Irvine, CA (US);
Mario Jesus Martinez, Laguna Niguel, CA (US); **Chi-She Chen**, Walnut, CA (US)

FOREIGN PATENT DOCUMENTS

WO WO99/30240 6/1999

OTHER PUBLICATIONS

(73) Assignee: **NETLIST, INC.**, Seocho-Gu (KR)

Non-Final Office Action, U.S. Appl. No. 13/970,606, filed Aug. 20, 2013, dated Nov. 23, 2015.

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 108 days.

(Continued)

Primary Examiner — Gurtej Bansal

(21) Appl. No.: **14/715,491**

(74) *Attorney, Agent, or Firm* — Morgan, Lewis & Bockius LLP

(22) Filed: **May 18, 2015**

Related U.S. Application Data

(57) **ABSTRACT**

(63) Continuation of application No. 14/324,990, filed on Jul. 7, 2014, now Pat. No. 9,037,809, which is a (Continued)


Certain embodiments described herein include a memory module having a printed circuit board including at least one connector configured to be operatively coupled to a memory controller of a computer system. The memory module further includes a plurality of memory devices on the printed circuit board and a circuit including a first set of ports operatively coupled to at least one memory device. The circuit further includes a second set of ports operatively coupled to the at least one connector. The circuit includes a switching circuit configured to selectively operatively couple one or more ports of the second set of ports to one or more ports of the first set of ports. Each port of the first set and the second set comprises a correction circuit which reduces noise in one or more signals transmitted between the first set of ports and the second set of ports.

(51) **Int. Cl.**
G06F 12/00 (2006.01)
G06F 13/16 (2006.01)
(Continued)

(52) **U.S. Cl.**
CPC **G06F 13/1684** (2013.01); **G06F 13/4027** (2013.01); **G11C 5/02** (2013.01); **G11C 5/06** (2013.01); **G11C 8/18** (2013.01)

(58) **Field of Classification Search**
CPC ... G06F 12/1684; G06F 13/4027; G11C 5/02; G11C 5/06; G11C 8/18
See application file for complete search history.

18 Claims, 9 Drawing Sheets



US 10,025,731 B1

Page 2

Related U.S. Application Data

continuation of application No. 13/412,243, filed on Mar. 5, 2012, now Pat. No. 8,782,350, which is a continuation of application No. 12/422,853, filed on Apr. 13, 2009, now Pat. No. 8,154,901.

- (60) Provisional application No. 61/044,839, filed on Apr. 14, 2008, provisional application No. 61/044,825, filed on Apr. 14, 2008, provisional application No. 61/044,801, filed on Apr. 14, 2008.

(51) Int. Cl.

G11C 5/02 (2006.01)
G11C 5/06 (2006.01)
G11C 8/18 (2006.01)
G06F 13/40 (2006.01)

(56) References Cited

U.S. PATENT DOCUMENTS

5,630,096	A	5/1997	Zuravleff
5,905,401	A	5/1999	Sher
6,061,754	A	5/2000	Cepulis
6,141,245	A	10/2000	Bertin
6,233,650	B1	5/2001	Johnson
6,262,938	B1	7/2001	Lee
6,381,140	B1	4/2002	Liao
6,546,476	B1	4/2003	Gillingham
6,636,446	B2	10/2003	Lee
6,704,910	B2	3/2004	Hong
7,024,518	B2	4/2006	Halbert et al.
7,191,302	B2	3/2007	Usami
7,225,303	B2	5/2007	Choi
7,254,036	B2	8/2007	Pauley
7,730,254	B2	6/2010	Risse
7,990,746	B2	8/2011	Rajan
8,244,971	B2	8/2012	Rajan
8,335,894	B1	12/2012	Rajan
8,856,464	B2	10/2014	Karamcheti
2001/0008006	A1	7/2001	Klein
2002/0038405	A1	3/2002	Leddige et al.
2003/0070052	A1	4/2003	Lai
2008/0025137	A1	1/2008	Rajan et al.
2008/0162790	A1	7/2008	Im
2009/0248969	A1	10/2009	Wu
2009/0296503	A1	12/2009	Chu et al.
2010/0125681	A1	5/2010	Patel
2011/0085406	A1	4/2011	Solomon et al.

OTHER PUBLICATIONS

Response to Non-Final Office Action, U.S. Appl. No. No. 13/970,606, filed Aug. 20, 2013, dated Mar. 23, 2016.

Inter Partes Review Case No. 2014-00882, Patent Owner Response, filed May 8, 2015.

Inter Partes Review Case No. 2014-00882, Petitioner Diablo Technologies, Inc.'s Reply to Patent Owner Netlist, Inc.'s Response, filed Jun. 19, 2015.

Inter Partes Review Case No. 2014-00882, Petitioner's Request for Oral Argument, filed Jul. 1, 2015.

Inter Partes Review Case No. 2014-00882, Patent Owner's Request for Oral Argument, filed Jul. 3, 2015.

Inter Partes Review Case No. 2014-00882, Order—Request for Oral Argument—37 C.F.R. 42.70, filed Jul. 10, 2015.

Inter Partes Review Case No. 2014-00882, Record of Oral Hearing, filed Nov. 16, 2015.

Inter Partes Review Case No. 2014-00882, Final Written Decision 35 USC 318 and 37 CFR 42.73, filed Dec. 14, 2015.

Inter Partes Review Case No. 2014-00882, Netlist Inc.'s Notice of Appeal, filed Feb. 10, 2016.

Inter Partes Review Case No. 2014-00882, Exhibit 1022, 'Professor Carl Sechen Deposition Transcript,' filed Jun. 19, 2015.

Inter Partes Review Case No. 2014-00882, Exhibit 1023, 'Dr. Srinivasan Jagannathan Supplemental Declaration,' filed Jun. 19, 2015.

Inter Partes Review Case No. 2014-00882, Exhibit 2002, 'Declaration of Professor Carl Sechen,' filed May 8, 2015.

Inter Partes Review Case No. 2014-00882, Exhibit 2003, 'Transcript of Videotaped Deposition of Srinivasan Jagannathan, Ph. D.,' filed May 8, 2015.

Microsoft Computer Dictionary, Fifth Edition, Microsoft Press, p. 334, 2002.

Jacob, Bruce, et al. "Memory Systems—Cache, DRAM, Disk," Elsevier Inc. 2008, Chapters 7 and 10.

U.S. District Court Northern District of California, Case No. 4:13-cn-05889-YGR, *Netlist Inc. v. Smart Storage Systems, Inc., and Diablo Technologies, Inc.*, Joint Claim Construction and Prehearing Statement Pursuant to Patent L. T. 4-3, filed Aug. 27, 2014, in 28 pages.

Excerpts from the Authoritative Dictionary of IEEE Standard Terms, Seventh Edition, Standards Information Network, IEEE Press (2000), pp. 349-411.

Excerpts from the Authoritative Dictionary of IEEE Standard Terms, Seventh Edition, Standards Information Network, IEEE Press (2000), pp. 133-265.

Selected Entries in Oxford English Dictionary, Oxford University Press 2014, <http://www.oed.com/view/Entry/175039?redirectedFrom=selectively#eid>.

Web search results from the Oxford English Dictionary, Oxford University Press 2014, http://www.oed.com/search?searchType=dictionary&q=select&_searchBin=Search.

Inter Partes Review Case No. 2014-00883, Patent Owner Response, filed May 8, 2015.

Inter Partes Review Case No. 2014-00883, Petitioner Diablo Technologies, Inc.'s Reply to Patent Owner Netlist, Inc.'s Response, filed Jun. 19, 2015.

Inter Partes Review Case No. 2014-00883, Petitioner's Request for Oral Argument, filed Jul. 1, 2015.

Inter Partes Review Case No. 2014-00883, Patent Owner's Request for Oral Hearing, filed Jul. 3, 2015.

Inter Partes Review Case No. 2014-00883, Order—Request for Oral Argument—37 C.F.R. 42.70, filed Jul. 10, 2015.

Inter Partes Review Case No. 2014-00883, Record of Oral Hearing, filed Nov. 16, 2015.

Inter Partes Review Case No. 2014-00883, Final Written Decision—35 USC 318 and 37 CFR 42.73, filed Dec. 14, 2015.

Inter Partes Review Case No. 2014-00883, Netlist Inc.'s Notice of Appeal, filed Feb. 10, 2016.

Inter Partes Review Case No. 2014-00883, Exhibit 1025, 'Professor Carl Sechen Deposition Transcript,' filed Jun. 19, 2015.

Inter Partes Review Case No. 2014-00883, Exhibit 1026, 'Dr. Srinivasan Jagannathan Supplemental Declaration,' filed Jun. 19, 2015.

Inter Partes Review Case No. 2014-01011, Patent Owner Response, filed May 8, 2015.

Inter Partes Review Case No. 2014-01011, Petitioner Diablo Technologies, Inc.'s Reply to Patent Owner Netlist, Inc.'s Response, filed Jun. 19, 2015.

Inter Partes Review Case No. 2014-01011, Petitioner's Request for Oral Argument, filed Jul. 1, 2015.

Inter Partes Review Case No. 2014-01011, Patent Owner's Request for Oral Hearing, filed Jul. 3, 2015.

Inter Partes Review Case No. 2014-01011, Order—Request for Oral Argument—37 C.F.R. 42.70, filed Jul. 10, 2015.

Inter Partes Review Case No. 2014-01011, Record of Oral Hearing, filed Nov. 16, 2015.

Inter Partes Review Case No. 2014-01011, Final Written Decision—35 USC 318 and 37 CFR 42.73, filed Dec. 14, 2015.

Inter Partes Review Case No. 2014-01011, Netlist Inc.'s Notice of Appeal, filed Feb. 10, 2016.

Inter Partes Review Case No. 2014-01011, Exhibit 1027, 'Professor Carl Sechen Deposition Transcript,' filed Jun. 19, 2015.

Inter Partes Review Case No. 2014-01011, Exhibit 1028, 'Dr. Srinivasan Jagannathan Supplemental Declaration,' filed Jun. 19, 2015.

US 10,025,731 B1

Page 3

(56) **References Cited**

OTHER PUBLICATIONS

Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Record of Oral Hearing, dated Jul. 26, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Final Written Decision, issued Sep. 28, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Exhibit 3001 to Final Decision, issued Sep. 28, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Exhibit 3001-2 to Final Decision, issued Sep. 28, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Exhibit 3002 to Final Decision, issued Sep. 28, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Exhibit 3003 to Final Decision, issued Sep. 28, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Exhibit 3004 to Final Decision, issued Sep. 28, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Final Decision, issued Sep. 28, 2016.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Petition for Inter Partes Review, filed Dec. 30, 2016.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibit 1003, "Declaration of Harold S. Stone," filed Dec. 30, 2016.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Petition for Inter Partes Review, filed Jan. 5, 2017.
Inter Partes Review of U.S. Pat. No. 8,516,185, Case No. IPR2017-00577, Exhibit 1003, "Declaration of Harold S. Stone," filed Jan. 5, 2017.
Inter Partes Review of U.S. Pat. No. 8,756,364, Case No. IPR2017-00549, Exhibit 1008, Declaration of John J. Kelly Regarding Records of Joint Electron Device Engineering Council (JEDEC), filed Dec. 30, 2016.
JEDEC Standard Double Data Rate (DDR) SDRAM Specification, JESD79 (Jun. 2000).
Texas Instruments SN74LS245 octal bus transceivers with 3-state outputs datasheet, 2002.
Stone, H.S. Microcomputer Interfacing. Reading, MA: Addison Wesley, 1982.
128Mbit SDRAM, Samsung datasheet K4S281632C, Rev. 00. Mar. 2000.
G. Moore, "Cramming more components onto integrated circuits," Electronics, vol. 38, No. 8, Apr. 19, 1965.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Related Matters, filed Apr. 28, 2015.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Patent Owner Response to Petition, filed Feb. 22, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Petitioner's Objections to Evidence, filed Feb. 29, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Petitioner's Reply, filed May 19, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Petitioner Request for Oral Argument, filed Jun. 2, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Patent Owner Motion for Observations, filed Jun. 2, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Patent Owner Request for Oral Argument, filed Jun. 2, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Order Trial Hearing, filed Jun. 8, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Patent Owner's Submission on Propriety of Petitioner Reply, filed Jun. 9, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Petitioner's Response to Motion for Observations, filed Jun. 10, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Petitioner's Response to Netlist's Submission, filed Jun. 13, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Patent Owner's Objections to Petitioner's Demonstrative Exhibits, filed Jun. 23, 2016.

Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Patent Owner's Demonstrative Exhibits, filed Jun. 24, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Petitioner's Updated Exhibit List, filed Jun. 24, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Supplemental Declaration of Dr. Jagannathan, filed May 19, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, May 10, 2016 Deposition of Carl Sechen, filed May 19, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Petitioner's Demonstrative Exhibits, filed Jun. 24, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Exhibit 2001, "Declaration of Professor Carl Sechen," filed Feb. 22, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Exhibit 2002, "Videotaped Deposition of Srinivasan Jagannathan, Ph.D on Feb. 12, 2016," filed Feb. 22, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Exhibit 2003, "Videotaped Deposition of Srinivasan Jagannathan, Ph.D on Apr. 20, 2015," filed Feb. 22, 2016.
Inter Partes Review of U.S. Pat. No. 7,881,150, Case No. IPR2015-01020, Exhibit 2007, "Deposition of Dr. Srinivasan Jagannathan on May 25, 2016," filed Jun. 2, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Patent Owner Response to Petition, filed Feb. 22, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Petitioner's Objections to Evidence, filed Feb. 29, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Petitioner's Reply, filed May 19, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Petitioner Request for Oral Argument, filed Jun. 2, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Patent Owner Motion for Observations, filed Jun. 2, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Patent Owner Request for Oral Argument, filed Jun. 2, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Patent Owner's Submission on Propriety of Petitioner Reply, filed Jun. 9, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Petitioner's Response to Motion for Observations, filed Jun. 10, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Petitioner's Response to Patent Owner's Submission, filed Jun. 13, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Patent Owner's Objections to Petitioner's Demonstrative Exhibits, filed Jun. 23, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Patent Owner's Demonstrative Exhibits, filed Jun. 24, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Petitioner's Updated Exhibit List, filed Jun. 24, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Supplemental Declaration of Dr. Jagannathan, filed May 19, 2016.
Inter Partes Review of U.S. Pat. No. 8,081,536, Case No. IPR2015-01021, Exhibit 2001, "Declaration of Professor Carl Sechen," filed Feb. 22, 2016.
International Search Report and Written Opinion, International Application No. PCT/US14/48517, filed Jul. 28, 2014, dated Oct. 27, 2014.
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.
Inter Partes Review Case No. IPR2014-01369, Decision Denying Institution of Inter Partes Review 37 C.F.R. § 42.108, issued Mar. 9, 2014.
Examination Report, European Patent Application No. 10730021.2, dated Apr. 14, 2014.
Response to Examination Report, European Patent Application No. 10730021.2, dated Jun. 4, 2014.

US 10,025,731 B1Page 4

(56)

References Cited

OTHER PUBLICATIONS

Examination Report, European Patent Application No. 10730021.2, dated Apr. 29, 2015.

Response to Examination Report, European Patent Application No. 10730021.2, dated Apr. 29, 2015, filed Nov. 4, 2015.

English Translation of the Notice of Grounds for Rejection, Korean Patent Application No. 2012-7004038, dated May 11, 2016.

Notice of Allowance, U.S. Appl. No. 13/970,606, filed Aug. 20, 2013, dated Jun. 27, 2016.




FIG. 1





FIG. 2






FIG. 4




FIG. 5





FIG. 6






FIG. 8




FIG. 9




FIG. 10

US 10,025,731 B1

1

MEMORY MODULE AND CIRCUIT PROVIDING LOAD ISOLATION AND NOISE REDUCTION

CROSS-REFERENCE TO RELATED APPLICATIONS

The present application is a continuation of U.S. patent application Ser. No. 14/324,990, filed Jul. 7, 2014, which is a continuation of U.S. patent application Ser. No. 13/412,243, filed Mar. 5, 2012, now U.S. Pat. No. 8,782,350, which is a continuation of U.S. patent application Ser. No. 12/422,853, filed Apr. 13, 2009, now U.S. Pat. No. 8,154,901, which claims the benefit of priority from U.S. Provisional Appl. No. 61/044,839, filed Apr. 14, 2008, from U.S. Provisional Appl. No. 61/044,825, filed Apr. 14, 2008, and from U.S. Provisional Appl. No. 61/044,801, filed Apr. 14, 2008, each of which are incorporated in its entirety by reference herein. This application is also related to U.S. patent application Ser. No. 12/422,912, filed on Apr. 13, 2009, now abandoned, and U.S. patent application Ser. No. 12/422,925, filed on Apr. 13, 2009, now U.S. Pat. No. 8,001,434, both of which are incorporated in their entirety by reference herein. This application is further related to U.S. patent application Ser. No. 13/183,253, filed Jul. 14, 2011, now U.S. Pat. No. 8,359,501, U.S. patent application Ser. No. 13/745,790, filed Jan. 19, 2013, now U.S. Pat. No. 8,689,064, and U.S. patent application Ser. No. 14/229,844, filed Mar. 29, 2014. This application is further related to U.S. patent application Ser. No. 12/422,912, filed on Apr. 13, 2009, now abandoned, U.S. patent application Ser. No. 12/422,925, filed on Apr. 13, 2009, now U.S. Pat. No. 8,001,434, U.S. patent application Ser. No. 13/183,253, filed Jul. 14, 2011, now U.S. Pat. No. 8,359,501, U.S. patent application Ser. No. 13/745,790, filed Jan. 19, 2013, now U.S. Pat. No. 8,689,064, U.S. patent application Ser. No. 14/299,844, filed Mar. 29, 2014, U.S. patent application Ser. No. 12/504,131, filed Jul. 16, 2009, now U.S. Pat. No. 8,417,870, U.S. patent application Ser. No. 12/761,179, filed Apr. 15, 2010, now U.S. Pat. No. 8,516,185, U.S. patent application Ser. No. 13/970,606, filed Aug. 20, 2013, U.S. patent application Ser. No. 13/287,042, filed Nov. 1, 2011, now, U.S. Pat. No. 8,756,364, U.S. patent application Ser. No. 13/971,231, filed Aug. 20, 2013, and U.S. application Ser. No. 13/288,850, filed Nov. 3, 2011.

BACKGROUND

The present application relates generally to the field of electrical signal integrity and more specifically to circuits for improving electrical signal integrity. Description of the Related Art

As the demand on the performance of electronic systems increase, the systems include higher speed processors and higher memory densities with relatively heavy fan-outs. However, higher speed processors and higher memory densities translate to higher power dissipation and increased memory access time, impairing the performance of the system.

Traditional systems sometimes incorporate multiplexers and/or demultiplexers to help quiet inactive signal paths, minimizing the overall dynamic power dissipation, and isolating unused subsystems. FIG. 1 schematically illustrates an example of such a system 10 which incorporates a multiplexer-demultiplexer 12 to isolate unused subsystems of a plurality of subsystems 16, 18. This arrangement can reduce the load on each driver 14 when the original signal path includes multiple destinations by inserting a multi-

2

plexer on each signal path between the driver 14, the receiver 20 and the multiple destinations. However, the multiplexer and/or demultiplexer 12 can present transmission line discontinuity regions in electronic systems (e.g., in electronic systems with signal paths with electrical lengths longer than $\frac{1}{4}$ of the wavelength of the operational frequency or the signal transition rate). In such systems, signal reflections may occur, changing wave characteristics and degrading system performance. These signal reflections may cause signal distortions and signal integrity issues which can contribute to reduced system performance and eventual failure of the system, thereby limiting the effectiveness of the multiplexer and/or demultiplexer in addressing the signal integrity and power dissipation issues.

In memory applications, some systems control “chip select” or device select signals, along with the memory address and control signals, based on the device select signals, in order to disable memory devices when they are not being accessed. FIG. 2 schematically illustrates an example of such a system 20 which isolates unused subsystems of a plurality of subsystems 16, 18 using device select signals 22, 24 (e.g., chip select signals). The devices are disabled by not activating (asserting) the address, control, and/or device select signals, thus logically isolating the memory devices which are not accessed. This arrangement generally requires a device select signal for each physical memory rank and the address and device select signals are generally controllable based on whether the corresponding memory device or group of memory devices is being accessed. Generating controllable address and device select signals can introduce timing delays on the device select, address, and control paths. In addition, this arrangement does not address the power dissipation and may degrade performance degradation by, for example, requiring that secondary control signals be generated for the memory devices.

SUMMARY

Certain embodiments described herein include a memory module comprising a printed circuit board comprising at least one connector configured to be operatively coupled to a memory controller of a computer system. The memory module comprises a plurality of memory devices on the printed circuit board and a circuit. The circuit comprises a first set of ports comprising a plurality of bi-directional ports, each port of the first set of ports operatively coupled to at least one memory device of the plurality of memory devices. The memory module further comprises a second set of ports comprising one or more bi-directional ports, each port of the second set of ports operatively coupled to the at least one connector. The memory module of certain embodiments comprises a switching circuit configured to selectively operatively couple one or more ports of the second set of ports to one or more ports of the first set of ports. The one or more ports of the first set of ports and the one or more ports of the second set of ports each comprise a correction circuit in certain embodiments. The correction circuit reduces noise in one or more signals transmitted between the one or more ports of the first set of ports and the one or more ports of the second set of ports.

A method of using a memory module with a computer system is described in certain embodiments. The method includes providing a memory module comprising a printed circuit board having at least one connector configured to be operatively coupled to a memory controller of a computer system and to a plurality of memory devices. The ‘memory

US 10,025,731 B1

3

module further including a circuit which comprises a first set of ports comprising a plurality of bi-directional ports, each port of the first set of ports operatively coupled to at least one memory device of the plurality of memory devices. The circuit further comprises a second set of ports comprising one or more bi-directional ports. Each port of the second set of ports is operatively coupled to the at least one connector. The circuit is configured to selectively operatively couple one or more ports of the second set of ports to one or more ports of the first set of ports in some embodiments. The one or more ports of the first set of ports and the one or more ports of the second set of ports each comprise a correction circuit. The correction circuit reduces noise in one or more signals transmitted between the one or more ports of the first set of ports and the one or more ports of the second set of ports. The method includes activating the circuit to selectively operatively couple at least one of the one or more ports of the second set of ports to at least one of the one or more ports of the first set of ports. The method also includes activating the circuit to reduce noise in the one or more signals.

Certain embodiments are described herein provide a circuit comprising a first set of ports comprising a plurality of bi-directional ports. Each port of the first set of ports may be configured to be operatively coupled to at least one memory device of a plurality of memory devices on a printed circuit board. The circuit further comprises a second set of ports comprising one or more bi-directional ports, each port of the second set of ports configured to be operatively coupled to at least one connector on the printed circuit board.

The at least one connector is configured to be operatively coupled to a memory controller of a computer system. The circuit further comprises a switching sub-circuit configured to selectively operatively couple one or more ports of the second set of ports to one or more ports of the first set of ports. The one or more ports of the first set of ports and the one or more ports of the second set of ports each comprise a correction circuit in certain embodiments. The correction circuit can reduce noise in one or more signals transmitted between the one or more ports of the first set of ports and the one or more ports of the second set of ports.

In certain embodiments, a circuit comprises a first set of ports comprising a plurality of bi-directional ports. Each port of the first set of ports can be configured to be operatively coupled to at least one subsystem of a plurality of subsystems. The circuit may comprise a second set of ports comprising one or more bi-directional ports, each port of the second set of ports configured to be operatively coupled to at least one subsystem of one or more subsystems. In certain embodiments, the circuit comprises a switching sub-circuit configured to selectively operatively couple one or more ports of the second set of ports to one or more ports of the first set of ports. The one or more ports of the first set of ports and the one or more ports of the second set of ports each comprise a correction circuit in certain embodiments. The correction circuit can reduce noise in one or more signals transmitted between the one or more ports of the first set of ports and the one or more ports of the second set of ports.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 schematically illustrates an example system which isolates unused subsystems and incorporates a multiplexer-demultiplexer.

FIG. 2 schematically illustrates an example system which isolates unused subsystems using device select signals.

4

FIG. 3 schematically illustrates an example memory system incorporating a memory module in accordance with certain embodiments described herein.

FIG. 4 schematically illustrates an example correction circuit in accordance with certain embodiments described herein.

FIG. 5 schematically illustrates an example programmable impedance matching circuit in accordance with certain embodiments described herein.

FIG. 6 schematically illustrates an example self-adjusting damper circuit in accordance with certain embodiments described herein.

FIG. 7 schematically illustrates example configuration including an example filter element connected to an example repeater/amplifier in accordance with certain embodiments described herein.

FIG. 8 schematically illustrates an example system having an example circuit connecting the system with one or more other systems in accordance with certain embodiments described herein.

FIG. 9 schematically illustrates an example system including an example circuit in accordance with certain embodiments described herein.

FIG. 10 is a flowchart of an example method 400 of using a memory module with a computer system in accordance with certain embodiments described herein.

DETAILED DESCRIPTION

Electronic systems incorporating devices and methods in accordance with embodiments described herein may include memory subsystems that contain multiple memory boards or modules that are driven by a single memory controller which synchronizes memory access and manages memory coherency. However, having a single memory controller may gradually degrade system performance and increase power consumption as the physical and the logical size of the memory subsystem increases. Increases in memory speed may further increase the power consumption. Additionally, impedance mismatches between the system board, the memory boards, and the memory devices may also substantially degrade the performance and power efficiency of traditional memory systems. Such memory systems may be used with a power supply which helps to stabilize the voltage level, and/or a thermo-mechanical component to moderate memory subsystem temperature, and/or software for throttling the processor to lower the memory access rate.

Embodiments described herein address these speed and power issues in memory systems. For example, a circuit of certain embodiments can comprise an isolation switch and can logically isolate inactive subsystems and reduce signal reflection and/or other noise at, for example, the interface between the system board and the memory board and at the interface between the memory board and the memory devices.

Certain embodiments described herein present a circuit for isolating inactive subsystems and maintaining and improving signal integrity in a variety of systems. Certain embodiments described herein are used to increase the performance of electronic systems without increasing the power dissipation or the memory access time. In certain embodiments, a plurality of circuits are configured so as to improve the interconnection (e.g., timing, power consumption, and/or signal integrity) between each of the plurality of circuits and a plurality of corresponding devices (e.g., memory devices) of a system or subsystem (e.g., a memory module of a computer system). For example, the data width

US 10,025,731 B1

5

of each of the plurality of circuits may be selected to optimize the interconnection (e.g., timing, power consumption, and/or signal integrity). In various embodiments, the data width of each of the plurality of devices may be 4, 8, or 16 bits, for example.

Certain embodiments described herein describe a circuit for use in high performance electronic systems which may contain transmission lines and relatively heavy fan-outs. A circuit is described herein which, in certain embodiments, includes at least some of the following components: a signal router, a programmable coarse impedance matching circuit, a self-adjusting damper circuit, a filter and/or a repeater/amplifier.

FIG. 3 schematically illustrates an example system incorporating a memory module 100 in accordance with certain embodiments described herein. The memory module 100 comprises a printed circuit board (PCB) 102. The PCB 102 includes at least one connector 104 configured to be operatively coupled to a memory controller 106 of a computer system 108. The memory module 100 includes a plurality of memory devices 110 on the PCB 102.

The memory module 100 of certain embodiments further includes a circuit 112 comprising a first set of ports 122 comprising a plurality of bi-directional ports 124, each port 124 of the first set of ports 122 operatively coupled to at least one memory device 110 of the plurality of memory devices 110. The circuit 112 further comprises a second set of ports 126 comprising one or more bi-directional ports 128, each port 128 of the second set of ports 126 operatively coupled to the at least one connector 104. The circuit 112 further comprises a switching sub-circuit 140 configured to selectively operatively couple one or more ports 128 of the second set of ports 126 to one or more ports 124 of the first set of ports 122. The one or more ports 124 of the first set of ports 122 and the one or more ports 128 of the second set of ports 126 each comprises a correction circuit 150 (see, e.g., FIG. 4) which reduces noise in one or more signals transmitted between the one or more ports 124 of the first set of ports 122 and the one or more ports 128 of the second set of ports 126.

In certain embodiments, the memory module 100 has a memory capacity of 512-MB, 1-GB, 2-GB, 4-GB, 8-GB, or higher or lower. Other memory capacities are also compatible with certain embodiments described herein. In addition, memory modules 100 having widths of 4 bits, 8 bits, 16 bits, 32 bits, as well as other widths (in bytes or in bits), are compatible with embodiments described herein. In certain embodiments, the PCB 102 has an industry-standard form factor. For example, the PCB 102 can have a low profile (LP) form factor with a height of 30 millimeters and a width of 133.35 millimeters. In certain other embodiments, the PCB 102 has a very high profile (VHP) form factor with a height of 50 millimeters or more. In certain other embodiments, the PCB 102 has a very low profile (VLP) form factor with a height of 18.3 millimeters. Other form factors including, but not limited to, small-outline (SO-DIMM), unbuffered (UDIMM), registered (RDIMM), fully-buffered (FB-DIMM), mini-DIMM, mini-RDIMM, VLP mini-DIMM, micro-DIMM, and SRAM DIMM are also compatible with certain embodiments described herein. For example, in other embodiments, certain non-DIMM form factors are possible such as, for example, single inline memory module (SIMM), multi-media card (MMC), and small computer system interface (SCSI).

In certain embodiments, the plurality of memory devices 110 of the memory module 100 may be arranged as one or more ranks, each rank of memory generally having a bit

6

width. In certain embodiments, each rank may comprise an independent set of memory devices 111 of the plurality of memory devices 110 that can be accessed by the memory controller 106 to access the full bit-width of the memory bus of the memory module 100. For example, a memory module 100 in which each rank of the memory module is 64 bits wide is described as having an "x 64" organization. Similarly, a memory module 100 having 72-bit-wide ranks is described as having an "x 72" organization. The number of memory devices 111 of a memory module 100 can be increased by increasing the number of memory devices 111 per rank or by increasing the number of ranks of the memory module 100. For example, a memory module with four ranks with each rank having N 512-MB memory devices 111 has double the memory capacity of a memory module with two ranks with each rank having N 512-MB memory devices 111 and four times the memory capacity of a memory module with one rank with each rank having N 512-MB memory devices.

During operation, one or more ranks of a memory module 100 may be selected or activated from the plurality of ranks by control signals that are received from a component of the system (e.g., a system memory controller 106 or a local memory controller of the memory module 100). Examples of such control signals include, but are not limited to, rank-select signals, also called chip-select signals. In certain other embodiments, the memory module 100 comprises only one rank of memory devices 111.

As discussed, the PCB 102 includes at least one connector 104 configured to be operatively coupled to a memory controller 106 of a computer system 108, such as a host computer system 108. For example, the connector 104 is electrically coupled, logically coupled, or both, with the memory controller 106 in some other manner. In certain embodiments, the connector 104 comprises a wireless radio frequency (RF) interface or a portion thereof and is wirelessly connected (e.g., via one or more RF antennas) to the memory controller 106. In some embodiments, some other connection mechanism is used besides the connector 104. Examples of host computer systems 108 include, but are not limited to, blade servers, 1U servers, personal computers (PCs), data storage systems and other applications in which space is constrained or limited. The memory controller 106 may comprise a disk controller of the computer system 108, for example. The memory controller 106 of certain embodiments may be mounted on a system board 109 of the host computer 108. The connector 104 can comprise a plurality of edge connections distributed along one or more edges of the PCB 102, which fit into a corresponding slot connector of the host system. The connector 104 of certain embodiments provides a conduit for power voltage as well as data, address, and control signals between the memory module 100 and the host system 108. For example, the connector 104 can comprise a standard 240-pin DDR2 edge connector. Additionally, in certain embodiments, more than one memory module 100 is coupled to the host system 108.

The plurality of memory devices 110 on the PCB 102 may include one or more volatile memory components. For example, the plurality of memory devices 110 of certain embodiments comprises two or more dynamic random-access memory (DRAM) elements 111. For example, in the block diagram of FIG. 3, the plurality of memory devices 110 comprises eight DRAM devices 111. Types of DRAM devices 111 compatible with certain embodiments described herein include, but are not limited to, DDR, DDR2, DDR3, and synchronous DRAM (SDRAM). The memory devices 111 may comprise other types of memory elements such as

US 10,025,731 B1

7

static random-access memory (SRAM). In addition, volatile memory devices **111** having bit widths of 4, 8, 16, 32, as well as other bit widths, are compatible with certain embodiments described herein. Memory devices **111** compatible with certain embodiments described herein have packaging which include, but are not limited to, thin small-outline package (TSOP), ball-grid-array (BGA), fine-pitch BGA (FBGA), micro, BGA (μ BGA), mini-BGA (mBGA), and chip-scale packaging (CSP).

The first set of ports **122** of the circuit **112** comprises a plurality of bidirectional ports **124** where each port of the first set of ports **122** is operatively coupled to at least one memory device **111** of the plurality of memory devices **110**. In addition, the second set of ports **126** of the connector **104** comprises one or more bi-directional ports **128** where each port of the second set of ports **126** is operatively coupled to at least one connector **104**. The switching sub-circuit **140** is configured to selectively operatively couple one or more ports **128** of the second set of ports **126** to one or more ports **124** of the first set of ports **122**. The operative coupling of the first set of ports **122** to the memory devices **110**, the second set of ports **126** to the connector **104**, and the selective operative coupling between the first and second set of ports **122**, **126** may comprise electrical coupling, logical coupling, or both. In other embodiments, optical coupling or some other manner of coupling may be employed.

In certain embodiments, each bi-directional port **128** of the second set of ports **126** may correspond to a signal of a memory bus **130** which operatively couples the memory controller **106** and the memory module **100** via the connector **104**. For example, the memory controller **106** may generally read data from and/or write data to the plurality of memory devices **110** on the memory module **100** through the memory bus **130** and through the bi-directional ports **124**, **128**. In certain embodiments, the number of bi-directional ports **128** of the second set of ports **126** may correspond to the number of data bits in the memory bus **130** (i.e., the bit-width of the memory bus **130**).

In certain embodiments, each of the bi-directional ports **124** in the first set of ports **122** may correspond to and be operatively coupled to a signal port of one of the memory devices **111** of the plurality of memory devices **110**. The bi-directional ports **124** correspond to the data signal ports, for example, of the plurality of memory devices **110**. Because each memory device **111** includes associated data signal ports, the number of bidirectional ports **124** in the first set of ports of certain embodiments may be larger than the number of bi-directional ports **128** in the second set of ports **122** (e.g., the number of bits in the memory bus **130**). In general, the overall number of ports of the plurality of memory devices **110** and corresponding bi-directional ports **124** in the first set of ports **126** may depend on the organization of the memory module. For example, the number of ports may depend on the number of memory devices and/or the number of ranks in the memory module. In certain embodiments, the memory module **100** includes a plurality of circuits **112**. The circuits **112** may be configured so as to improve the interconnection (e.g., timing, power consumption, and/or signal integrity) between each of the plurality of circuits **112** and a plurality of corresponding memory devices **111** of memory module **100**. For example, the data width of each of the plurality of circuits **112** may be selected to optimize the interconnection (e.g., timing, power consumption, and/or signal integrity). In various embodiments, the data width of each of the plurality of devices may be 4, 8, or 16 bits, for example. Other data widths are possible in other embodiments. In some embodiments, the data width of

8

each circuit **112** corresponds to the data width of each of the memory devices **111** such that each circuit **112** is operatively coupled to one of the plurality of memory devices **110**. In other embodiments, the data width of each circuit **112** corresponds to a multiple of the data width of each memory device **111** such that each circuit **112** may be operatively coupled to the data bits of more than one memory device **111**. In various embodiments, the data width of each circuit **112** corresponds to a segment of the data width of each memory device **111** or to segments of multiple memory devices **111**. As such, in some embodiments, each circuit **112** may be operatively coupled to a segment of the data bits of one memory device **111** or to segments of more than one memory device **111**. Other configurations are possible which optimize the interconnection (e.g., timing, power consumption, and/or signal integrity). In addition, in certain embodiments, the location of each of the plurality of circuits **112** on the PCB **102** is selected to optimize the interconnection, timing, power consumption, and signal integrity.

In other embodiments, the bi-directional ports **124** correspond to one or more other bits of the memory bus **130** instead of, or in addition to, the data bits. For example, one or more of the ports **124** may correspond to address, chip select, control (e.g., read/write) signals. Moreover, in some embodiments, one or more ports of the first set of ports **122** is not bi-directional. For example, in one embodiment, signals corresponding to address, chip select and/or control signals from the first set of ports **122** to the plurality of memory devices **110** and not vice versa. For example, one or more of the address, chip select, and control signals may be used to control the routing of the bi-directional ports **124** which correspond to data bits.

As discussed, the switching sub-circuit **140** is configured to selectively operatively couple one or more ports of the second set of ports **126** to one or more ports of the first set of ports **122**. In certain embodiments, the switching sub-circuit **112** may be capable of interconnecting any number of the second set of ports **126** to any number of the first set of ports **122** simultaneously. In certain embodiments, the switching sub-circuit **140** can pass signals (e.g., information such as the data signals described above) from one port **128** of the second set of ports **126** to an operatively coupled port **124** of the first set of ports **122** and vice versa. Additionally, the switching sub-circuit **140** of some embodiments can pass signals from one port **128** of the second set of ports **126** to multiple operatively coupled ports **124** of the first set of ports **122** and vice versa. In certain embodiments, the switching sub-circuit **140** isolates and provides a known static value to all un-driven ports, thereby reducing any dynamic power consumption associated with floating or toggling electrical nodes.

The switching sub-circuit **140** of certain embodiments can function as an isolation switch which isolates and/or inactivates individual subsystems not actively transferring information. For example, in an example embodiment, the plurality of memory devices **110** comprises multiple independently accessible ranks. When the memory controller accesses one of the ranks of the memory module **100**, the switching sub-circuit **140** isolates and/or inactivates at least the one or more other, un-accessed ranks. For example, the switching sub-circuit **140** does not operatively couple ports from the second of set of ports **126** to the first set of ports **122** which are coupled to the memory devices **111** of the unaccessed one or more ranks of memory devices **110**. In certain embodiments, the switching sub-circuit **140** may function as or be characterized as a scalable isolation switch router which routes the information from one subsystem to

US 10,025,731 B1

9

another (e.g., host computer **108** to the memory module **100** via the memory controller **106**), establishing a one-to-one connection for transferring data from point-to-point, and/or establishing a one-to-many connection for broadcasting data. The switching sub-circuit **140** comprises or is characterized as a bidirectional multiplexer-demultiplexer in certain embodiments. In other embodiments, the switching sub-circuit **140** comprises or is characterized as a router.

In certain embodiments, the memory module **100** further includes circuitry (not shown) for controlling the operation of the circuit **112** and the components therein (e.g., the switching sub-circuit **140** and the correction circuit **150**, as described more fully below). For example, in one embodiment, the memory module **100** comprises a local memory controller which is operatively coupled to one or more control inputs (not shown) of the circuit **112** to control the selective coupling of the second set of ports **126** to the first set of ports **122**. In certain other embodiments, some other component of the memory module **100** controls the operation of the circuit **112** or the circuit **112** comprises internal control circuitry. In various embodiments, the circuitry for controlling the circuit **112** includes discrete logic, one or more microprocessors, a field-programmable gate array (FPGA), or a computer programmable logic device (CPLD). In yet other configurations, an external component controls the operation of the circuit **112**.

The circuit **112** of certain embodiments comprises one physically integral component (e.g., one integrated circuit package mounted on the memory module **100**). Various alternative configurations for the physical organization of the circuit **112** are possible, however. In certain other embodiments, the circuit **112** comprises multiple, physically separate components. For example, in one embodiment, the circuit **112** comprises one physically separate component per memory device **111** (e.g., integrate circuit packages mounted on the memory module **100**). In another embodiment, the circuit **112** comprises one physically separate component per memory rank. In some embodiments, the circuit **112** comprises one physically separate component corresponding to multiple memory devices **111**.

FIG. 4 schematically illustrates an example correction circuit **150** in accordance with certain embodiments described herein. As shown, in certain embodiments, the correction circuit **150** comprises one or more correction elements (e.g., at least one coarse correction element **152** to provide coarse correction of the noise and at least one fine correction fine correction element **154** to provide fine correction of the noise). The correction circuit **150** of certain embodiments further comprises at least one filter element **156**. In certain embodiments, the correction circuit **150** further comprises at least one repeater/amplifier element **158**, which improves the signal quality of the one or more signals.

Referring to FIGS. 3 and 4, the correction circuit **150** reduces noise in one or more signals transmitted between the one or more ports of the first set of ports **122** and the one or more ports of the second set of ports **126**. For example, the correction circuit **150** may reduce noise due to impedance mismatches (e.g., due to discontinuity regions) at nodes between the memory controller **106** and the plurality of memory devices **110**. In certain embodiments, the noise may comprise signal reflections. For example, in certain embodiments, the signal paths between the host computer **108** and the plurality of memory devices **110** or portions thereof may be characterized as transmission lines if they are greater than or about equal to a particular fraction of the wavelength of

10

the signal, such as, for example, greater than or equal to about $\frac{1}{4}$ of the wavelength of the signal. In other configurations, the signal paths may be characterized as transmission lines when the signal path is greater than or equal to some other fraction of the wavelength of the signal such as, for example, $\frac{1}{2}$, $\frac{1}{8}$, or $\frac{1}{10}$ of the wavelength. In certain embodiments, the system (e.g., memory module **100**) incorporating the circuit **112** has signal delays which are related to signal characteristics such as the signal transition rate and/or the rise/fall time of the signals. For example, the system incorporating the circuit **112** may include signal delays which are more than some fraction (e.g., $\frac{1}{2}$) of the signal transition rate and/or the rise/fall time of the signals. In other embodiments, the signal delays are related in some other way to the signal transition rate and/or the rise/fall time of the signals (e.g., are more than $\frac{1}{4}$, $\frac{1}{8}$ of those values) or are related to some other signal property.

The reflections may be caused by one or more discontinuity regions along the signal paths. For example, discontinuity regions comprise points or regions at which the characteristic impedance of the corresponding signal path changes. For example, the interface between the system board **109** of the host computer **108** and the memory module **100** (e.g., the connector **104**) may comprise a discontinuity region along each of the signal paths. The interfaces between the connector **104** of the memory module **100** and each of the memory devices **111** mounted on the PCB **102** may comprise one or more discontinuity regions along each of the signal paths.

The noise may comprise cross-coupling between the signal paths in certain embodiments. For example, a cross-coupling impedance may be present between one or more of the signal paths or portions thereof. In certain embodiments, the cross-coupling impedance may include cross-coupling due to cross-coupling capacitance, inductance (e.g., mutual inductance), transconductance, or a combination thereof. Additionally, a variety of other noise are possible such as, for example, electromagnetic interference or other types of noise.

Referring to FIG. 4, in certain embodiments, the correction circuit **150** includes one or more elements to provide for coarse correction of signal noise due to impedance mismatches (e.g., discontinuity regions) in the system. The at least one coarse correction element **152** comprises at least one programmable impedance matching circuit in certain embodiments. FIG. 5 schematically illustrates an example programmable impedance matching circuit **160** in accordance with certain embodiments described herein. For example, the programmable impedance matching circuit **160** may comprise at least one programmable resistor **162**. The resistance of the programmable resistor **162** may be controllable by a control signal **164**. In certain embodiments, the coarse correction element **152** may be integrated with the switching sub-circuit **140**, the fine correction element **154**, the filter element **156**, and/or the repeater/amplifier element **158**. Alternatively, in other embodiments, the coarse correction element **152** may be implemented with discrete components.

In other embodiments, the coarse correction element **152** comprises a network of resistors **162** which can be configured to provide programmable impedance matching. In some embodiments, the coarse correction element **152** may comprise a set of programmable coarse impedance matching circuits **160** that are connected to each of the ports **124**, **128** of the first set of ports **122** and the ports **128** of the second set of ports **126**. Advantageously, the impedance values of the coarse correction element **152** may be dynamically

US 10,025,731 B1

11

controllable, configurable, and/or programmable, so as to match the characteristic impedances of the transmission lines (e.g., signal paths) in the system. The values of these impedances may differ (e.g., may be programmed differently) based on which particular signal path is activated (e.g., by the switching sub-circuit 140). Where the coarse correction element 152 comprises at least one impedance matching circuit, the granularities of the configurable impedance values may be configured based on the application of the particular system. In some embodiments, the impedance matching circuit of the coarse correction element 152 may include a real time controllable impedance control input for preprogramming and/or dynamically controlling the impedance matching circuit. In this manner, the impedance matching circuit 152 may, in certain embodiments, provide dynamically controllable variable impedance which can be pre-set for matching the impedance of the subsystem that drives the circuit 112.

In certain embodiments, the at least one fine correction element 154 of the correction circuit 150 further reduces signal reflections due to impedance mismatches and/or reduces other noise. In certain embodiments, for example, the coarse correction element 152 may be characterized as a first refining stage and the fine correction element 154 may be referred to as a second refining stage. The fine correction element 154 may smooth out the signal after it has undergone coarse correction (e.g., impedance matching) and reduces or removes transmission line characteristics from the received signal. Referring again to FIG. 4, the fine correction element 154 can be connected, for example, in parallel with the programmable coarse correction element 152. In some embodiments, the fine correction element 154 operates to reduce the impedance mismatches by introducing a current source that damps to the slew rate of the input signal, reducing signal reflection. The fine correction element 154 actively counteracts the noise in certain embodiments. In addition, the fine correction element 154 of certain embodiments is dynamically responsive to one or more changes in the noise, such as, for example, changes in the frequency of the noise component.

The fine correction element 154 may include at least one damper circuit. The damper circuit includes at least one self-adjusting damper circuit 170 in certain embodiments, such as the example self-adjusting damper circuit 170 schematically illustrated by FIG. 6. In certain embodiments, the self-adjusting damper circuit 170 comprises an input terminal 172 adapted to receive an input signal 174 comprising an input noise component and an input signal component. The damper circuit 170 further comprises an output terminal 176. In certain embodiments, the damper circuit 170 includes a first sub-circuit 184 adapted to generate a first signal 187 responsive to at least a portion of the noise component. For example, the first sub-circuit 184 may generate the first signal 187 in response to the portion of the noise component not reduced by the coarse correction element 152 in certain embodiments. The damper circuit 170 of certain embodiments further comprises a second sub-circuit 188 adapted to generate a second signal 190 corresponding to a delayed version of at least a portion of the input signal. The damper circuit 170 combines the first signal 187 with the second signal 190 to generate an output signal 108 on the output terminal 176 wherein the output signal comprises an output noise component smaller than the input noise component. For example, the first signal 187 and the second signal 190 may be combined at the node 182 of the damper circuit 170. The output signal 176 of certain embodiments is substantially representative of the signal

12

component of the input signal 174. The first subcircuit 184 of certain embodiments comprises one or more current sources 185, 186. For example, the current sources 185, 186 may each comprise a controlled current source such as a voltage-controlled current source 185, 186 in certain embodiments. In some embodiments, the first sub-circuit 184 may comprise or be characterized as a negative feedback current source. In various embodiments, the self-adjusting damper circuit 170 includes an enable control input (not shown) for enabling the self-adjusting damper circuit 170. Examples of self-adjusting damper circuits 170 in accordance with certain embodiments described herein are described in U.S. Provisional Application No. 61/044, 825, filed Apr. 14, 2008, and copending U.S. application Ser. No. 12/422,912, filed by the assignee of the present application on Apr. 13, 2009. Both applications are incorporated in their entirety by reference herein.

In certain embodiments, the at least one filter element 156 and the at least one repeater/amplifier element 158 may also reduce noise from other sources such as, for example, signal coupling (e.g., cross-coupling). The filter element 156 may reduce noise from the one or more signals where the noise includes a predetermined range of frequencies. For example, the filter element 156 may include at least one band-pass filter and may reduce any low and/or high frequency noise components. In other embodiments, the filter element 156 comprises at least one other type of a filter such as a notch filter, a low-pass filter, or a high-pass filter.

One or more of the filter element 156 and the repeater/amplifier element 158 of certain embodiments enhance signal slew rate to reduce the transient power dissipation at target devices or components (e.g., one or more of the memory devices 110). FIG. 7 schematically illustrates example configuration including an example filter element 156 connected to an example repeater/amplifier element 158 in accordance with certain embodiments described herein. The repeater/amplifier element 158 includes an output enable control input 159 for controlling the bi-directional signal path. In various embodiments one or more of the filter element 156 and repeater/amplifier element 158 are programmable and/or dynamically controllable. The repeater/amplifier element may further include a reference input 157. In certain embodiments, the repeater/amplifier element 158 may include one or more operational amplifiers, for example. Referring again to FIG. 4, in certain embodiments, only filter element 156, only the repeater, only the amplifier, or any combination of the three elements may be included in the circuit.

The filter element 156 and/or the repeater/amplifier element 158 of certain embodiments provide wave reshaping functions and generally reduces signal noise. For example, the filter element 156 and/or the repeater/amplifier element 158 may reduce signal ringing due to reflection (e.g., reflections at one or more interfaces along the signal paths such as interfaces between the circuit 112, the connector 104, the memory devices 110 and/or the system board 109 of FIG. 3). In certain embodiments, the repeater/amplifier element 158 reshapes the signal, which may be deteriorated due to transmission over lossy transmission lines having discontinuity regions. The repeater/amplifier element 158 may be configured to repeat the input signal, generating an output signal with earlier transition trigger points and faster slew rates in order to regain speed and to reduce the transient power.

In certain embodiments, the circuit 112 or portions thereof (e.g., the switching sub-circuit 140, the correction circuit 150 or elements thereof, etc.) is selected from a group

US 10,025,731 B1

13

consisting of: a programmable-logic device (PLD), an application-specific integrated circuit (ASIC), a field-programmable gate array (FPGA), a custom-designed semiconductor device, and a complex programmable-logic device (CPLD). For example, in certain embodiments, the circuit **112** or portions thereof is implemented in one or more of a PLD, ASIC, FPGA, or CPLD. In other embodiments, the circuit **112** or portions thereof is implemented using some other technology, such as, for example, using discrete components.

In various embodiments, the elements of the circuit **112** such as the switching sub-circuit **140**, the coarse correction element **152**, the fine correction element **154**, the filter element **156**, and the repeater/amplifier element **158**, comprise or may be characterized as independent, fully controllable functional blocks. Additionally, in certain embodiments, any subset of one or more of these functional blocks may be implemented in order to enhance performance, reduce power dissipation, or provide other advantages.

FIG. **8** schematically illustrates an example system **200** having an example circuit **112** connecting a system **200** with one or more other systems **220** in accordance with certain embodiments described herein. The circuit **112** may be any of the circuits **112** discussed herein, such as the circuit **112** discussed above with respect to FIG. **3**, for example. The system **200** of FIG. **8** includes a driver **202**, a receiver **201**, and two subsystems **203**, **204**. For illustration purposes, the system **200** of FIG. **8** is described as an open system that communicates with other systems **220** via the driver **202** and the receiver **201**. Furthermore, in some embodiments, a connector (not shown) connects the system **200** with other systems by means of transmission lines. For example, the system **200** may be a memory system and the receiver **201** and the driver **202** of FIG. **8** may represent memory signals (e.g., data, address and/or control signals) transmitted between the system **220** (e.g., a memory controller and the subsystems **203**, **204** (e.g., two sets of independently addressable memory devices such as DRAM devices). In certain embodiments, for example, the circuit **112** may be disposed on a memory board that is populated with sets of memory devices **203**, **204**.

In various embodiments, the receiver **201** may or may not contain any active circuits. For example, the receiver **201** may be a connector such as a socket or DIMM interface with a system board (e.g., the connector **104** described above in regards to FIG. **3**). The output signal from the receiver **201** may exhibit the characteristics of the signal on a transmission line such as, for example, wave dispersion, attenuation, and reflection. In certain embodiments, the output signal from the receiver **201** goes through a number of refining, filtering and/or amplification stages in circuit **112** as described herein to improve the performance of the overall system. For example, the signal quality is enhanced by introducing at least one circuit **112** including at least one correction circuit **150** comprising one or more coarse correction elements **152** (e.g., impedance matching circuits), fine correction elements **154** (e.g., self-adjusting dampers), filter elements **156** (e.g., bandpass filters) and/or repeater/amplifier elements **158** as described herein on the signal paths (e.g., from receiver **101** to subsystem **103**). For example, the circuit **112** may reduce power consumption by selectively isolating the inactive (or un-accessed) subsystems and may enhance the access rate of a targeted subsystem. For example, in one embodiment, the system **200** comprises a memory module **100** such as one of the memory modules described herein, the subsystems **203** and **204** comprise memory devices such as the memory devices **111**

14

described herein, and the other system **220** comprises a host computer **108** including a memory controller **106** as described herein.

In certain embodiments, the performance is enhanced and power dissipation is reduced by removing the load associated with other (e.g., unused or inactive) subsystems (e.g., the subsystem **204**) and reshaping the signal waveform using a filter element **156** and/or the repeater/amplifier element **158**. In some embodiments, there are two refining stages for reducing signal reflections due to impedance mismatches. In some embodiments, for example, the first refining stage is the coarse correction element **152** (e.g., an impedance matching circuit) that matches the impedance of the transmission line (e.g., wire) prior to a receiver **201** with the characteristic impedance of the transmission line (e.g., wire) in the system **100**. The second stage may be the fine correction element **154** (e.g., a self adjusting damper for further reducing signal reflections due to impedance mismatches and/or reducing other noise).

FIG. **9** schematically illustrates an example system **300** including an example circuit **112** in accordance with certain embodiments described herein. The circuit **112** may be any of the circuits **112** discussed herein, such as the circuit **112** discussed above with respect to FIG. **3**, for example. The circuit **112** comprises a first set of ports **122** comprising a plurality of bi-directional ports, each port of the first set of ports **122** configured to be operatively coupled to at least one subsystem **311** of a plurality of subsystems **310**. The circuit **112** further comprises a second set of ports **126** comprising one or more bi-directional ports, each port of the second set of ports **126** configured to be operatively coupled to at least one subsystem **321** of one or more subsystems **320**. In certain embodiments, the circuit **112** is configured to selectively operatively couple one or more ports of the second set of ports **126** to one or more ports of the first set of ports **122**. In certain embodiments, the one or more ports of the first set of ports **122** and the one or more ports of the second set of ports **126** each comprises a correction circuit **150** which reduces noise in one or more signals transmitted between the one or more ports of the first set of ports **122** and the one or more ports of the second set of ports **126**. The correction circuit **150** of certain embodiments may comprise one of the correction circuits **150** described herein such as the correction circuits **150** described above with respect to FIGS. **4-7** and may include/at least one coarse correction element **152**, at least one fine correction element **154**, at least one filter element **156**, and at least one repeater/amplifier element **158**, or any combination or sub-combination thereof. In other embodiments, the fine correction element **154** may be placed after the switching sub-circuit **140** in order to achieve better results. For example, the fine correction element **154** may be connected after the filter element **156** (e.g., between the filter **156** and the repeater/amplifier element **158**) in certain embodiments.

In certain embodiments, the system **300** includes a first set of interfaces **330** positioned between (e.g., logically and/or physically between) and operatively coupled to the first set of ports **122** and to the plurality of subsystems **310**. In addition, the system **300** of certain embodiments may include a second set of interfaces **340** positioned between (e.g., logically and/or physically between) and operatively coupled to the second set of ports **126** and the one or more subsystems **320**. For example, the first and/or second sets of interfaces **330**, **340** may comprise port connections.

In an example configuration, the plurality of subsystems **310** include a plurality of memory devices **311** (e.g., DRAM devices) mounted on a memory module and the first set of

US 10,025,731 B1

15

interfaces 330 comprise interfaces between the circuit 112 and the plurality of memory devices 311. For example, each of the first set of ports 122 may correspond to a memory signal (e.g., a data, address, and/or control signal) and each of the first set of interfaces 330 is operatively coupled to a corresponding memory signal port of a corresponding memory device 311. The one or more subsystems 320 of the example configuration comprise a host computer system and the second set of interfaces 340 comprises memory board (e.g., a PCB of a memory module) interfaces to the host system. For example, the second set of interfaces 340 form part of a connector. Each of the second set of ports 126 and corresponding interfaces of the second set of interfaces 340 corresponds to a memory signal (e.g., a data, address, and/or control signal) of at least one memory bus between the host system and the memory module.

The coarse correction elements 152 (e.g., impedance matching circuits) and the fine correction elements (e.g., self-adjusting dampers) of the second set of ports 126 of the example configuration are configured to reduce the signal reflection due to the impedance mismatching between the system board and the memory board. The coarse correction elements 152 (e.g., impedance matching circuits) and the fine correction elements (e.g., self-adjusting dampers) of the second set of ports 122 of the example configuration are configured to reduce the signal reflection due to the impedance mismatching between the circuit 112, the memory board and/or the plurality of memory devices. The filter elements 156 (e.g., a bandpass filter) reduce or filter out cross-coupling between signals, for example. For example, the filter elements 158 reduce or filter out cross-coupling between data, address, and/or control signals (e.g., wires).

FIG. 10 is a flowchart of an example method 400 of using a memory module 100 with a computer system in accordance with certain embodiments described herein. While the method 400 is described herein by reference to the memory module 100 schematically illustrated by FIG. 3, other memory modules, electronic systems or subsystems, and/or circuits are also compatible with the embodiments described herein. In certain embodiments, at operational block 402, the method 400 comprises providing a memory module 100 comprising a printed circuit board 102 having at least one connector 104 configured to be operatively coupled to a memory controller 106 of a computer system 108 and to a plurality of memory devices 110. The memory module 100 further comprises a circuit 112 including a first set of ports 122 comprising a plurality of bi-directional ports 124. In certain embodiments, each port 124 of the first set of ports 122 is operatively coupled to at least one memory device 111 of the plurality of memory devices 110. The circuit of certain embodiments further includes a second set of ports 126 comprising one or more bidirectional ports 128' where each port 128 of the second set of ports 126 is operatively coupled to the at least one connector 104.

The circuit 112 of certain embodiments is configured to selectively operatively couple one or more ports 128 of the second set of ports 126 to one or more ports 124 of the first set of ports 122. The one or more ports 124 of the first set of ports 122 and the one or more ports 128 of the second set of ports 126 each comprises at least one correction circuit 150 which reduces noise in one or more signals transmitted between the one or more ports 124 of the first set of ports 122 and the one or more ports 128 of the second set of ports 126.

The method 400 of certain embodiments further includes activating the circuit 112 at operational block 404 to selectively operatively couple at least one of the one or more

16

ports 128 of the second set of ports 126 to at least one of the one or more ports 124 of the first set of ports 122. For example, in certain embodiments, activating the circuit 112 to perform the selective operative coupling may include sending the appropriate control signal or combination of control signals to a switching sub-circuit 140 of the circuit 112. In certain embodiments, the method 400 further includes activating the circuit 112 to reduce noise in the one or more signals at operational block 406. For example, in one embodiment, activating the circuit 112 to reduce noise comprises activating one or more of the elements of the correction circuit 150 (e.g., at least one coarse correction element 152, at least one fine correction element 154, at least one filter element 156, and at least one repeater/amplifier element 158) of ports 124, 128 of the first and second sets of ports 122, 126 corresponding to active signal paths between the memory controller 106 and the memory devices 110. Activating one or more of the elements 152, 154, 156, 158 of the correction circuit 150 may include sending control signals (e.g., enable signals) to the appropriate elements 152, 154, 156, 158 and/or receiving the control signals (e.g., enable signals) at the appropriate elements 152, 154, 156, 158.

The operational blocks 402, 404, 406 of the method 400 may be performed in a different order than provided or one or more of the operational blocks may be performed substantially at the same time or during overlapping time periods. For example, in certain embodiments, at least the operational block 404 and the operational block 406 are performed substantially simultaneously. In addition, in some embodiments, one or more of the operational blocks 402, 404, 406 may not be included, an additional operational block may be included, or both. The devices and methods described herein may be incorporated in any system (e.g. a board level system) or device that contains one or more connectors or wires.

Although certain preferred embodiments and examples are discussed above, it is understood that the inventive subject matter extends beyond the specifically disclosed embodiments to other alternative embodiments and/or uses of the invention and obvious modifications and equivalents thereof. It is intended that the scope of the inventions disclosed herein should not be limited by the particular disclosed embodiments. Thus, for example, in any method or process disclosed herein, the acts or operations making up the method/process may be performed in any suitable sequence and are not necessarily limited to any particular disclosed sequence. Various aspects and advantages of the embodiments have been described where appropriate. It is to be understood that not necessarily all such aspects or advantages may be achieved in accordance with any particular embodiment. Thus, for example, it should be recognized that the various embodiments may be carried out in a manner that achieves or optimizes one advantage or group of advantages as taught herein without necessarily achieving other aspects or advantages as may be taught or suggested herein.

What is claimed is:

1. A memory module operable to communicate data with a system memory controller via a memory bus in response to address and control signals from the system memory controller, comprising:

a printed circuit board comprising at least one connector configured to be operatively coupled to the memory bus, the at least one connector including a plurality of edge connections distributed along one or more edges of the printed circuit board;

US 10,025,731 B1

17

a plurality of memory devices on the printed circuit board, the plurality of memory devices being arranged in multiple ranks, each rank comprising an independent set of memory devices that can be accessed by the system memory controller to access a full bit-width of the memory bus;

at least one circuit coupled between the at least one connector and the plurality of memory devices, the at least one circuit comprising a first set of port connections coupled to respective ones of the plurality of edge connections and a second set of port connections coupled to the plurality of memory devices, each circuit of the at least one circuit further comprising a set of correction circuits, each correction circuit of the set of correction circuits being configured to make corrections in one or more signals transmitted between a respective port connection of the first set of port connections and a corresponding port connection in the second set of port connections, the each correction circuit of the set of correction circuits including at least one programmable impedance matching circuit; and control circuitry configured to receive the address and control signals from the system memory controller and to control the plurality of memory devices in response to the address and control signals, wherein the control circuitry is further configured to dynamically control the at least one programmable impedance matching circuit in the each correction circuit of the set of correction circuits in the each circuit of the at least one circuit based on which of the multiple ranks is selected to communicate data with the system memory controller in response to the address and control signals.

2. The memory module of claim 1, wherein the at least one impedance matching circuit in the each correction circuit of the set of correction circuits in the each circuit of the at least one circuit has dynamically programmable impedances whose values may differ depending on which of the multiple ranks is selected to communicate data with the system memory controller in response to the address and control signals.

3. The memory module of claim 1, wherein the at least one impedance matching circuit in the each correction circuit of the set of correction circuits in the each circuit of the at least one circuit comprises at least one network of programmable resistors that is dynamically controllable.

4. The memory module of claim 1, wherein the each correction circuit of the set of correction circuits in the each circuit of the at least one circuit further comprises one or more self-adjusting damper circuits.

5. The memory module of claim 1, wherein the each correction circuit of the set of correction circuits in the each circuit of the at least one circuit further comprises first one or more dynamically programmable repeater/amplifier elements.

6. The memory module of claim 1, wherein the each correction circuit of the set of correction circuits in the each

18

circuit of the at least one circuit further includes a delay circuit that is configured to generate delayed versions of the signals transmitted between the respective port connection of the first set of port connections and the corresponding port connection in the second set of port connections.

7. The memory module of claim 1, wherein the number of port connections in the first set of port connections corresponds to a bit width of the memory bus.

8. The memory module of claim 1, wherein the at least one circuit comprises multiple physically separate integrated circuit packages each having a data width corresponding to a data width of one of the plurality of memory devices.

9. The memory module of claim 1, wherein the at least one circuit comprises multiple physically separate integrated circuit packages each having a data width corresponding to a data width of multiple memory devices of the plurality of memory devices.

10. The memory module of claim 1, wherein the control circuitry is configured to activate signal paths in the each circuit of the at least one circuit in response to the address and control signals.

11. The memory module of claim 1, wherein the control circuitry is configured to send at least one control signal to the at least one circuit in response to the address and control signals.

12. The memory module of claim 1, wherein the each correction circuit of the set of correction circuits in the each circuit of the at least one circuit includes one or more circuit elements, and wherein the control circuitry is configured to activate the one or more circuit elements in response to the address and control signals.

13. The memory module of claim 1, wherein each circuit in the at least one circuit has a data width of 4, 8 or 16 bits.

14. The memory module of claim 1, wherein the at least one circuit is configured to improve signal integrity in signals transmitted between the at least one connector and the plurality of memory devices.

15. The memory module of claim 1, wherein the at least one circuit is configured to improve timing of signals transmitted between the at least one connector and the plurality of memory devices.

16. The memory module of claim 1, the at least one circuit is configured to improve power consumption by signals transmitted between the at least one connector and the plurality of memory devices.

17. The memory module of claim 1, wherein the at least one circuit includes signal paths, wherein the control circuitry is configured to activate a subset of the signal paths in response to the address and control signals.

18. The memory module of claim 17, wherein the at least one circuit is further configured to provide a known static value to each port connection in the second set of port connections that is not in an active signal path.

* * * * *

EXHIBIT 3

(12) **United States Patent**
Chen et al.

(10) **Patent No.:** **US 12,373,366 B2**
(45) **Date of Patent:** **Jul. 29, 2025**

(54) **MEMORY WITH ON-MODULE POWER MANAGEMENT**

(56) **References Cited**

(71) Applicant: **Netlist, Inc.**, Irvine, CA (US)

U.S. PATENT DOCUMENTS
5,272,677 A 12/1993 Yamamura
5,563,838 A 10/1996 Mart et al.
(Continued)

(72) Inventors: **Chi-She Chen**, Walnut, CA (US);
Jeffrey C. Solomon, Irvine, CA (US);
Scott H. Milton, Irvine, CA (US);
Jayesh Bhakta, Cerritos, CA (US)

FOREIGN PATENT DOCUMENTS

(73) Assignee: **NETLIST, INC.**, Irvine, CA (US)

EP 0 744 748 B1 10/2000
EP 0 908 892 B1 7/2005
(Continued)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

OTHER PUBLICATIONS

(21) Appl. No.: **17/582,797**

Jedec Standard, Fbdimm, Advanced Memory Buffer (AMB), JES82-20 Jedec Solid State Technology Association, Mar. 2007, 190 pages.*

(22) Filed: **Jan. 24, 2022**

(Continued)

(65) **Prior Publication Data**

US 2022/0222191 A1 Jul. 14, 2022

Primary Examiner — Hashem Farrokh

(74) *Attorney, Agent, or Firm* — Vorys, Sater, Seymour and Pease LLP; Khaled Shami

Related U.S. Application Data

(63) Continuation of application No. 17/328,019, filed on May 24, 2021, now Pat. No. 11,232,054, which is a (Continued)

(57)

ABSTRACT


In certain embodiments, a memory module includes a printed circuit board (PCB) having an interface that couples it to a host system for provision of power, data, address and control signals. First, second, and third buck converters receive a pre-regulated input voltage and produce first, second and third regulated voltages. A converter circuit reduces the pre-regulated input voltage to provide a fourth regulated voltage. Synchronous dynamic random access memory (SDRAM) devices are coupled to one or more regulated voltages of the first, second, third and fourth regulated voltages, and a voltage monitor circuit monitors an input voltage and produces a signal in response to the input voltage having a voltage amplitude that is greater than a threshold voltage.

(51) **Int. Cl.**
G06F 13/28 (2006.01)
G06F 1/18 (2006.01)
(Continued)

(52) **U.S. Cl.**
CPC **G06F 13/28** (2013.01); **G06F 1/185** (2013.01); **G06F 3/0613** (2013.01);
(Continued)

(58) **Field of Classification Search**
CPC .. G06F 1/185; G06F 12/0246; G06F 12/0638;
G06F 13/1694; G06F 13/28;
(Continued)

38 Claims, 22 Drawing Sheets



US 12,373,366 B2

Page 2

Related U.S. Application Data

continuation of application No. 17/138,766, filed on Dec. 30, 2020, now Pat. No. 11,016,918, which is a continuation of application No. 15/934,416, filed on Mar. 23, 2018, now abandoned, which is a continuation of application No. 14/840,865, filed on Aug. 31, 2015, now Pat. No. 9,928,186, which is a continuation of application No. 14/489,269, filed on Sep. 17, 2014, now Pat. No. 9,158,684, which is a continuation of application No. 13/559,476, filed on Jul. 26, 2012, now Pat. No. 8,874,831, which is a continuation-in-part of application No. 12/240,916, filed on Sep. 29, 2008, now Pat. No. 8,301,833, which is a continuation of application No. 12/131,873, filed on Jun. 2, 2008, now abandoned.

- (60) Provisional application No. 61/512,871, filed on Jul. 28, 2011, provisional application No. 60/941,586, filed on Jun. 1, 2007.

(51) Int. Cl.

G06F 3/06 (2006.01)
G06F 12/02 (2006.01)
G06F 12/06 (2006.01)
G06F 13/16 (2006.01)
G06F 13/40 (2006.01)
G06F 13/42 (2006.01)
G11C 7/10 (2006.01)
G11C 14/00 (2006.01)

(52) U.S. Cl.

CPC **G06F 3/0659** (2013.01); **G06F 3/0685** (2013.01); **G06F 12/0246** (2013.01); **G06F 12/0638** (2013.01); **G06F 13/1694** (2013.01); **G06F 13/4027** (2013.01); **G06F 13/4243** (2013.01); **G11C 7/1072** (2013.01); **G11C 14/0018** (2013.01); **G06F 2212/205** (2013.01); **G06F 2212/7208** (2013.01)

(58) Field of Classification Search

CPC **G06F 13/4027**; **G06F 13/4243**; **G06F 2212/205**; **G06F 2212/7208**; **G06F 3/0613**; **G06F 3/0659**; **G06F 3/0685**; **G11C 11/005**; **G11C 11/409**; **G11C 14/0018**; **G11C 16/10**; **G11C 7/1072**; **Y02D 10/00**

USPC 711/110
 See application file for complete search history.

(56) References Cited

U.S. PATENT DOCUMENTS

5,563,839 A 10/1996 Herdt et al.
 5,630,096 A 5/1997 Zuravleff et al.
 5,675,467 A 10/1997 Nishimura et al.
 5,831,421 A 11/1998 Taira et al.
 6,031,762 A 2/2000 Saitoh
 6,067,241 A 5/2000 Lu
 6,351,827 B1 2/2002 Co et al.
 6,418,065 B2 7/2002 Sato et al.
 6,434,044 B1 8/2002 Gongwer et al.
 6,493,282 B2 12/2002 Iida et al.
 6,556,475 B2 4/2003 Yamazaki et al.
 6,693,840 B2 2/2004 Shimada et al.
 6,798,709 B2 9/2004 Sim et al.
 6,856,556 B1 * 2/2005 Hajack G11C 16/30 365/189.11
 6,889,335 B2 5/2005 Hargis et al.
 7,016,249 B2 3/2006 Ross
 7,023,187 B2 4/2006 Shearon et al.

7,034,569 B1 4/2006 Balasubramanian et al.
 7,061,214 B2 6/2006 Mayega et al.
 7,061,821 B2 6/2006 Coteus et al.
 7,117,334 B2 10/2006 Almeida et al.
 7,127,622 B2 10/2006 Schnepfer
 7,161,820 B2 1/2007 Funaba et al.
 7,166,934 B2 1/2007 Dewey
 7,221,617 B2 5/2007 Flach et al.
 7,224,595 B2 5/2007 Dreps et al.
 7,269,764 B2 9/2007 Dart et al.
 7,280,417 B2 10/2007 Choi et al.
 7,426,649 B2 9/2008 Brittain et al.
 7,495,943 B2 2/2009 Takemura et al.
 7,587,559 B2 9/2009 Brittain et al.
 7,613,880 B2 11/2009 Miura et al.
 7,619,912 B2 11/2009 Bhakta et al.
 7,644,216 B2 1/2010 Fahr et al.
 7,656,735 B2 2/2010 Kagan et al.
 7,679,967 B2 3/2010 Chung et al.
 7,692,938 B2 4/2010 Petter
 7,702,928 B2 4/2010 Schumacher et al.
 7,719,866 B2 5/2010 Boldo
 7,721,130 B2 5/2010 Prete et al.
 7,724,604 B2 * 5/2010 Amidi G11C 7/22 711/115
 7,779,285 B2 8/2010 Doblar et al.
 8,089,795 B2 1/2012 Rajan et al.
 8,301,833 B1 10/2012 Chen et al.
 8,874,831 B2 10/2014 Lee et al.
 9,158,684 B2 10/2015 Lee et al.
 9,858,218 B1 1/2018 Lee
 9,928,186 B2 3/2018 Lee et al.
 10,217,523 B1 2/2019 Lee et al.
 10,474,595 B2 11/2019 Lee
 10,672,458 B1 6/2020 Shaeffer et al.
 10,860,506 B2 12/2020 Lee et al.
 10,949,339 B2 3/2021 Lee et al.
 11,016,918 B2 5/2021 Chen et al.
 11,232,054 B2 1/2022 Chen et al.
 2006/0080515 A1 * 4/2006 Spiers G06F 11/1456 714/14
 2006/0136765 A1 6/2006 Poisner et al.
 2006/0174140 A1 * 8/2006 Harris G11C 5/04 713/300
 2006/0294437 A1 12/2006 Washburn et al.
 2007/0136523 A1 * 6/2007 Bonella G06F 12/0866 711/E12.019
 2008/0101147 A1 5/2008 Amidi
 2008/0126624 A1 * 5/2008 Prete G11C 7/1006 710/53
 2008/0238536 A1 10/2008 Hayashi et al.
 2008/0307240 A1 12/2008 Dahan et al.
 2009/0020608 A1 1/2009 Bennett et al.
 2009/0027844 A1 1/2009 Chen et al.
 2009/0034354 A1 2/2009 Resnick
 2010/0110748 A1 5/2010 Best
 2010/0205470 A1 8/2010 Moshayedi et al.
 2010/0257304 A1 10/2010 Rajan et al.
 2012/0271990 A1 * 10/2012 Chen G11C 29/52 711/E12.008
 2019/0004985 A1 1/2019 Lee et al.

FOREIGN PATENT DOCUMENTS

JP 9-34806 A 2/1997
 JP 2001-318741 A 11/2001
 JP 2001-357669 A 12/2001
 JP 3707888 B2 10/2005
 JP 2006-268419 A 10/2006
 KR 10-2005-0073902 A 7/2005
 WO 01/37090 A1 5/2001
 WO 2005/066968 A1 7/2005

OTHER PUBLICATIONS

Jedec Standard, Fbdimm Specification, DDR2 Sdram Fully Buffered Dimm (FBDIMM) Design Specification., JESD205, Jedec Solid State Technology Association, Mar. 2007, 129 pages.*

US 12,373,366 B2

Page 3

(56)

References Cited

OTHER PUBLICATIONS

Exhibits to Defendant's Second Amended Answer, Document 58-8, Case 1:21-cv-01453, filed Nov. 1, 2022.

IPR2022-000063, Decision Granting Inter Partes Review of U.S. Pat. No. 10,217,523, entered May 5, 2022.

Micron's Preliminary Identification of Proposed Claim Constructions in Civil Action Nos. 6:21-cv-00431-ADA and 6:21-cv-00430-ADA, IPR2022-00418 Exhibit 1013, filed in IPR on Jan. 14, 2022.

Petition for Inter Partes Review of U.S. Pat. No. 11,016,918, IPR No. IPR2023-00406, 151 pages, filed Jan. 6, 2023.

Petition for Inter Partes Review of U.S. Pat. No. 11,232,054, IPR No. IPR2023-00405, 147 pages, filed Jan. 6, 2023.

Decision Granting Institution of Inter Partes Review and Granting Motion for Joinder for U.S. Pat. No. 11,016,918, IPR No. IPR2023-00406, 7 pages, entered Jun. 23, 2023.

Decision Granting Institution of Inter Partes Review and Granting Motion for Joinder for U.S. Pat. No. 11,232,054, IPR No. IPR2023-00405, 7 pages, entered Jun. 23, 2023.

Decision of Final Written Decision Determining All Challenged Claims Unpatentable Dismissing Petitioner's Motion to Exclude for U.S. Pat. No. 11,016,918, IPR No. IPR2022-00996, 91 pages, entered Dec. 6, 2023.

Decision of Final Written Decision Determining All Challenged Claims Unpatentable Dismissing Petitioner's Motion to Exclude for U.S. Pat. No. 11,232,054, IPR No. IPR2022-00999, 80 pages, entered Dec. 5, 2023.

Machine Translation of KR 10-2005-0073902 A, published Jul. 18, 2005, 7 pages.

Renesas, ISL6537 ACPI Regulator/Controller for Dual Channel DDR Memory Systems Data Sheet, Rev 6.00, 16 pages, Jul. 18, 2007.

Intersil, ISL6532B ACPI Regulator/Controller for Dual Channel DDR Memory Systems Data Sheet, 16 pages, Jul. 2004.

Freescale Semiconductor, Hardware and Layout Design Considerations for DDR Memory Interfaces, Rev. 6, NXP, 48 pages, Apr. 2007.

Plessl and Platzner, TKDM—A Reconfigurable Co-processor in a PC's Memory Slot, Swiss Federal Institute of Technology, Jan. 15, 2004, 8 pages.

Renesas, ISL6534 Dual PWM with Linear Data Sheet, Rev 2.00, 29 pages, Nov. 18, 2005.

Maxim Integrated, DC-DC Converter Tutorial, APP 2031, 14 pages, Nov. 29, 2001.

Exhibit 2014, Qimonda, Internet Data Sheet, Rev. 1.25, Jun. 2007, 24 pages.

Exhibit 2015, Siemens, 3.3V 2M x 64/72-Bit 1 Bank SDRAM Module, HYS64/72V2200GU-8/-10, 3.3V 4M x 64/72-Bit 2 Bank SDRAM Module, HYS64/72V4200GU-8/-10, 168 pin unbuffered DIMM Modules, downloaded from www.datasheetcatalog.com; 1998, 18 pages.

Exhibit 2016, PCI Bus Variation—Technology Note, copyright 2006 by Euresys, 9 pages.

Exhibit 2017, Qimonda, Internet Data Sheet, Rev. 1.42, Sep. 2007, 27 pages.

Exhibit 2018, Transcend, Product Support, What is the difference between SDRAM, DDR1, DDR2, DDR3 and DDR4?, available online at <https://www.transcend-info.com/support/faq-296#:~:text=DDR3's prefetch buffer width is,according to the temperature variation,downloaded Aug. 24, 2022>, 2 pages.

Exhibit 2021, Texas Instruments: Application Report, Low Dropout Operation in a Buck Converter, SLUA928A—Dec. 2018—revised Mar. 2019, 18 pages.

Exhibit 2022, Electronic Design, Simple Soft-Start Circuitry Provides Long Startup Times, available online at <http://electronicdesign.com/power/simple-soft-start-circuitry-provides-long-startup-times>, published Jun. 22, 1998, 3 pages.

Exhibit 2024, SCHMID, Understanding Hard Drive Performance, available online at <https://www.tomshardware.com/reviews/>

understanding-hard-drive-performance, 1557-5.html, downloaded Aug. 15, 2022, published Mar. 5, 2007, 17 pages.

Exhibit 2025, Micron, SDR SDRAM, MT48LC64M4A2—16 Meg x 4 x 4 Banks, MT48LC32M8A2-8 Meg x 8 x 4 Banks, MT48LC16M16A2-4 Meg x 16 x 4 Banks, copyright 1999 by Micron Technology, Inc., 94 pages.

Exhibit 2026, Micron, Sdr SDRAM, MT48LC16M4A2-4 Meg x 4 x 4 Banks, MT48LC8M8A2-2 Meg x 8 x 4 Banks, MT48LC4M16A2-1 Meg x 16 x 4 Banks, copyright 1999 by Micron Technology, Inc., 84 pages.

Exhibit 2027, Transcend, DDR2 SO-DIMM, copyright 2012 by Transcend, 3 pages.

Exhibit 2028, Micron Technical Note: DDR3 Point-to-Point Design Support, copyright 2013 by Micron Technology, Inc., 27 pages.

Exhibit 2030, IPR 2022-00996 and IPR 2022-00999, Videoconference Deposition of Andrew Wolfe, Ph.D. regarding U.S. Pat. Nos. 11,016,918 and 11,232,054, 311 pages, dated Mar. 16, 2023.

Exhibit 2031, IPR 2022-00996, Declaration of Dr. Andrew Wolfe regarding U.S. Pat. No. 11,016,918, 452 pages, executed on May 16, 2022.

Exhibit 2032, Claim Construction Order, filed Dec. 14, 2022, 35 pages.

Exhibit 2033, Defendants' Objections to Claim Construction Memorandum Order, filed Dec. 29, 2022, 11 pages.

Exhibit 2043, Samsung Unveils New Power Management Solutions for DDR5 Modules, May 18, 2021, 6 pages.

Exhibit 2044, Micron, DDR2 SDRAM RDRAM, MT9HTF3272Y—256MB, MT9HTF6472PY—512MB, MT9HTF12872PY—1GB, copyright 2003 by Micron Technology, Inc., 23 pages.

Exhibit 2046, JEDEC Standard No. 21C, p. 4.20.11-1, PC2-6400/PC2-5300/PC2-4200/PC2-3200, DDR2 Unbuffered SO-DIMM, Reference Design Specification, Revision 2.5, Jul. 2008, 60 pages.

Exhibit 2048, Micron, DDR2 SDRAM FB-DIMM, MT18HTF12872FD—1GB, MT18HTF25672FD—2GB, copyright 2005 by Micron Technology, Inc., 39 pages.

Exhibit 2056, IPR 2022-00615, Remote Videotaped Deposition by Virtual Zoom of Andrew Wolfe, Ph.D. regarding U.S. Pat. No. 7,619,912, 297 pages, dated Jan. 4, 2023.

Exhibit 2057, Micron Technical Brief: Micron DDR5: Key Module Features, DDR4 vs. DDR5 Dimm, copyright 2020 by Micron Technology, Inc., 7 pages.

Exhibit 2063, Netlist, Demonstrative Exhibit for IPR 2022-00996 and IPR 2022-00999, filed Sep. 6, 2023, 180 pages.

Exhibit 2102, Lee et al., PCI Express Power and Mini Card Solutions—Design Note 346, Linear Technology—Design Notes, copyright 2004 by Linear Technology Corporation, 2 pages.

Paper 1, IPR 2022-00996, Petition for Inter Partes Review of U.S. Pat. No. 11,016,918, filed May 17, 2022, 151 pages.

Paper 1, IPR 2022-00999, Petition for Inter Partes Review of U.S. Pat. No. 11,232,054, filed May 17, 2022, 149 pages.

Paper 7, IPR 2022-00996, Patent Owner Preliminary Response, filed Sep. 9, 2022, 85 pages.

Paper 7, IPR 2022-00999, Patent Owner Preliminary Response, filed Sep. 8, 2022, 82 pages.

Paper 10, IPR 2022-00996, Decision - Granting Institution of Inter Partes Review, dated Dec. 7, 2022, 56 pages.

Paper 21, IPR 2022-00996, Patent Owner Response, filed Mar. 28, 2023, 91 pages.

Paper 22, IPR 2022-00999, Patent Owner Response, filed Mar. 28, 2023, 90 pages.

Paper 26, IPR 2022-00999, Petitioner's Reply to Patent Owner's Response, filed Jun. 20, 2023, 57 pages.

Paper 31, IPR 2022-00996, Patent Owner's Sur-Reply, filed Aug. 4, 2023, 48 pages.

Paper 33, IPR 2022-00999, Patent Owner's Sur-Reply, filed Aug. 4, 2023, 48 pages.

Transcript, Day 2 (redacted), Trial on the Merits, *Netlist, Inc. (plaintiff) vs. Samsung Electronics Co., Ltd., et al. (defendants)*, Apr. 17, 2023, 375 pages.

Transcript, Day 4 (redacted), Trial on the Merits, *Netlist, Inc. (plaintiff) vs. Samsung Electronics Co., Ltd., et al. (defendants)*, Apr. 19, 2023, 338 pages.

US 12,373,366 B2

Page 4

(56)

References Cited

OTHER PUBLICATIONS

Exhibit 1003, IPR 2022-00999, Declaration of Dr. Andrew Wolfe regarding U.S. Pat. No. 11,232,054, 431 pages, executed on May 16, 2022.

Exhibit 1015, IPR 2014-00994, Decision—Denying Institution of Inter Partes Review, issued Dec. 16, 2014, 16 pages.

Exhibit 1016, Corrected Petition for Inter Partes Review of U.S. Pat. No. 8,301,833, filed Sep. 22, 2014, 67 pages.

Exhibit 1017, IPR 2014-01370, Decision—Denying Institution of Inter Partes Review, issued Mar. 13, 2015, 19 pages.

Exhibit 1019, IPR 2017-00649, Decision—Denying Institution of Inter Partes Review, issued Jul. 24, 2017, 17 pages.

Exhibit 1020, Petition for Inter Partes Review of U.S. Pat. No. 8,874,831, filed Jan. 17, 2017, 78 pages.

Exhibit 1021, IPR 2017-00692, Final Written Decision, issued Jul. 5, 2018, 42 pages.

Exhibit 1022, Petition for Inter Partes Review of U.S. Pat. No. 8,301,833, filed Jan. 14, 2022, 76 pages.

Exhibit 1026, JEDEC Standard, DDR2 SDRAM Specification, JESD79-2B (Revision of JESD7-2A), Jedec Solid State Technology Association, Jan. 2005, 113 pages.

Exhibit 1027, JEDEC Standard, Fbdim: Advanced Memory Buffer (AMB), JESD82-20, Jedec Solid State Technology Association, Mar. 2007, 190 pages.

Exhibit 1028, Jedec Standard, Fbdim: Advanced Memory Buffer (FBDIMM) Design Specification, JESD205, JEDEC Solid State Technology Association, Mar. 2007, 129 pages.

Exhibit 1031, PCI Local Bus Specification, Revision 2.2, Dec. 18, 1998, 322 pages.

Exhibit 1032, Power Electronics: Converters, Applications, and Design, Second Edition, copyright 1995 by John Wiley & Sons, Inc., 37 pages.

Exhibit 1040, Texas Instruments, TPS51020, Dual, Voltage Mode, DDR Selectable, Synchronous, Step-Down Controller for Notebook System Power, SLUS564B, Jul. 2003 (revised Dec. 2003), 27 pages.

Exhibit 1041, Fairchild Semiconductor, FAN5026, Dual DDR/Dual-Output PWM Controller, Oct. 2005, 17 pages.

Exhibit 1042, Murata Manufacturing Co., Ltd, Power Supply Reference Guide for XILINX FPGAs, Sep. 2006, 33 pages.

Exhibit 1043, Murata Products, Power Supply Reference Guide for ALTERA FPGAs, Catalog No. DC-03-A, 2008, 40 pages.

Exhibit 1045, Jedec Standard, Double Data Rate (DDR) SDRAM Specification, JESD79, Jedec Solid State Technology Association, Jun. 2000, 77 pages.

Exhibit 1046, JEDEC Standard, DDR3 SDRAM Specification, JESD79-3A (Revision of JESD79-3), JEDEC Solid State Technology Association, Sep. 2007, 189 pages.

Exhibit 1048, Murata Manufacturing, DC-DC Converter Specification (DRAFT), MPD4S014S, Jan. 21, 2008, 11 pages.

Exhibit 1049, Micron, NAND Flash Memory, MT29F2G08AABWP/MT29F2G16AABWP, MT29F4G08BABWP/MT29F4G16BABWP, MT29F8G08FABWP, copyright 2004 by Micron Technology, Inc., 57 pages.

Exhibit 1058, LENK, Simplified Design of Switching Power Supplies, Linear Technology, EDN Series for Design Engineers, copyright 1995 by Butterworth-Heinemann, 9 pages.

Exhibit 1061, Analog Devices, Super Sequencer with Margining Control and Auxiliary ADC Inputs, ADM1066, copyright 2006 by Analog Devices, Inc., 36 pages.

Exhibit 1062, Analog Devices, Analog Dialogue, vol. 40, No. 2, 2006, 24 pages.

Exhibit 1063, National Semiconductor, LMC6953, PCI Local Bus Supervisor, Oct. 1996, 12 pages.

Exhibit 1067, Lattice Semiconductor Corporation, LatticeXP Family Data Sheet, Version 04.2, Mar. 2006, 127 pages.

Exhibit 1068, Complaint for Declaratory Judgment of Non-Infringement and Unenforceability; Breach of Contract, filed Oct. 15, 2021, 82 pages.

Exhibit 1069, First Amended Complaint for Declaratory Judgment of Non-Infringement and Unenforceability; Breach of Contract, filed Jan. 18, 2022, 122 pages.

Exhibit 1071, Complaint, filed Dec. 20, 2021, 49 pages.

Exhibit 1072, The Samsung Defendants' Answer, filed Apr. 12, 2022, 37 pages.

Exhibit 1073, First Amended Complaint, filed May 3, 2022, 71 pages.

Exhibit 1075, IPR 2022-00999, Deposition of William Mangione-Smith, Ph.D. regarding U.S. Pat. No. 11,232,054, 245 pages, dated Jun. 12, 2023.

Exhibit 1076, IPR 2022-00999, redline comparing Declarations—Declaration of Dr. William Henry Mangione-Smith regarding U.S. Pat. No. 11,232,054, 165 pages, executed on Mar. 29, 2022.

Exhibit 1077, Netlist, Technology Tutorial: In the Matter of Certain Memory Modules and Components Thereof and Products Containing Same, Inv No. 337-TA-1023, May 8, 2017, 11 pages.

Exhibit 1078, Analog Devices, Step-Down DC-to-DC Controller, ADP1821, copyright 2005-2007 by Analog Devices, Inc., 24 pages.

Exhibit 2001, IPR 2022-00999, Declaration of Dr. Sunil P. Khatri regarding U.S. Pat. No. 11,232,054, 68 pages, executed on Sep. 8, 2022.

Exhibit 2003, BELLONI et al., 24.6 A 4-Output Single-Inductor DC-DC Buck Converter with Self-Boosted Switch Drivers and 1.2A Total Output Current, ISSCC 2008/ Session 24/ Analog Power Techniques/ 24.6, 2008 IEEE International Solid-State Circuits Conference, 3 pages.

Exhibit 2004, MA et al., Single-Inductor Multiple-Output Switching Converters With Time-Multiplexing Control in Discontinuous Conduction Mode, IEEE Journal of Solid-State Circuits, vol. 38, No. 1, Jan. 2003, pp. 89-100 (12 pages).

Exhibit 2006, Micron, Technical Note: Power Solutions for DDR2 Notebook PCs, TN-47-05 DDR2 Power Solutions for Notebooks, copyright 2004 by Micron Technology, Inc., 16 pages.

Exhibit 2007, Texas Instruments, LP2996-N, LP2996A Ddr Termination Regulator, SNOSA40K, Nov. 2002 (revised Dec. 2016), 39 pages.

Exhibit 2008, Jameco Electronics, Jameco Part No. 843868, National Semiconductor: LP2996: DDR Termination Regulator, Jun. 2006, 19 pages.

Exhibit 2009, Jameco Electronics, Jameco Part No. 843930, National Semiconductor: LP2997: DDR-II Termination Regulator, Jun. 2006, 13 pages.

Exhibit 2010, National Semiconductor: LP2998: DDR-II and DDR-I Termination Regulator, Dec. 12, 2007, 18 pages.

Exhibit 2011, Bergveld, Battery Management Systems: Design by Modelling, copyright 2001 by Royal Philips Electronics N.V., 328 pages.

Exhibit 2012, Romo, DDR Memories: Comparison and Overview, NXP, pp. 70-73 (4 pages), Sep. 21, 2007.

Exhibit 2031, IPR 2022-00996, Declaration of Dr. William Henry Mangione-Smith regarding U.S. Pat. No. 11,016,918, 144 pages, executed on Mar. 28, 2023.

Paper 25, IPR 2022-00996, Petitioner's Reply to Patent Owner's Response, filed Jun. 20, 2023, 57 pages.

Exhibit 1014, Petition for Inter Partes Review of U.S. Pat. No. 8,301,833, filed Jun. 20, 2014, 69 pages.

Exhibit 1018, Petition for Inter Partes Review of U.S. Pat. No. 8,301,833, filed Jan. 13, 2017, 70 pages.

Exhibit 2013, JEDEC Standard No. 21-C, p. 4.20.3-1; PC133 SDRAM Unbuffered SO-DIMM, Reference Design Specification, Revision 1.02, 26 p. downloaded Aug. 16, 2022.

FBDIMM: Architecture and Protocol, JEDEC Standard JESD206.01, Jan. 2007, 128 pages.

ISL6537 "ACPI Regulator/Controller for Dual Channel DDR Memory Systems", Renesas, 16 pages, Jul. 18, 2007, taken from internet on Apr. 16, 2024 from: <https://www.renesas.com/us/en/document/dst/isl6537-datasheet>.

NXP "Hardware and Layout Design Considerations for DDR Memory Interfaces", Freescale Semiconductor, Inc., 48 pages, 2007, taken from internet on Apr. 16, 2024 from: <https://www.nxp.com/docs/en/application-note/AN2582.pdf>.

US 12,373,366 B2

Page 5

(56)

References Cited

OTHER PUBLICATIONS

TKDM “TKDM—A Reconfigurable Co-processor in a PC’s Memory Slot”, Plessi, 8 pages, Dec. 17, 2003, taken from internet on Apr. 17, 2024 from: <https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=615235db2bed5bb4823f9c3a9d383a92e1f602c7>.

ISL6534 “Dual PWM with Linear”, Renesas, 29 pages, Nov. 18, 2005, taken from internet on Apr. 16, 2024 from: <https://dir.heisener.com/DatasheetDownload/ISL6534CV-TR5229.pdf>.

Tong et al., “A System Level Implementation of Rijndael on a Memory-slot based FPGA Card”, Tong, 8 pages, Jan. 2003, taken from internet on Jun. 6, 2024 from: https://phwl.org/assets/papers/aes_fpt02.pdf.

National Semiconductor (Nov. 2006). LP3913 Power Management IC for Flash Memory Based Portable Media Players. Retrieved from <https://datasheetspdf.com/pdf-file/572488/NationalSemiconductor/LP3913/1>. See pp. 1-3.

Intersil (Dec. 2006). ISL6227 Dual Mobile-Friendly PWM Controller with DDR Option. Retrieved from http://pdfhtml.ic37.com/pdf_file_A/20200531/pdf_pdf/pdf5/INTERSIL/ISL6227_datasheet_647405/105904/ISL6227_datasheet.pdf. See pp. 2, 10, 12, 14 & 20.

Xilinx (Jul. 2005). OPB Synchronous DRAM (SDRAM) Controller (v1.00e). Retrieved from <https://datasheet.datasheetarchive.com/originals/library/Datasheet-080/DASF007621.pdf>. See pp. 9 & 11.

Phillips (2003). New multimarket products quarterly highlights. Retrieved from <https://datasheet.datasheetarchive.com/originals/library/Datasheets-DAV2/DSADA0022280.pdf>.

Maxim (Sep. 2005). Buck Converters Proliferate in Handhelds as Features and Processing Power Increase. Retrieved from <https://datasheet.datasheetarchive.com/originals/library/Datasheet-081/DASF0031684.pdf>.

Keller, S. J. (2006). The Physical Design of the ILLIAC 6 Supercomputing Platform (Doctoral dissertation, University of Illinois at Urbana-Champaign). See pp. 12, 13 & highlighted sections.

Interstill (2004). Power Management Application Guide for Altera FPGAs. Retrieved from <https://www.datasheetarchive.com/pdf/>

<download.php?id=c9facaf95737a8327c7217fbff086b497d4f48&type=P&term=buck%2520SDRAM>. See pp. 4 & 9.

Summit (2006). Application Note 59. Retrieved from <https://www.datasheetarchive.com/pdf/download.php?id=a7ccb0edb0c2b62d33ea2ca63eeb31b0884da8&type=P&term=buck%2520SDRAM>. See pp. 1, 3 & 4.

Analog Devices (2006). TigerSHARC Embedded Processor ADSP-TS201S.

Renesas (2005). ISL6531 Datasheet.

Texas Instruments (2001). Designing for Small-Size, High Frequency Applications Using TPS546xx DC/DC Converters.

Intersil (2002). Powering DDR Memory in Desktop Computers and Datacomm Systems. Retrieved from <https://datasheet.datasheetarchive.com/originals/library/Datasheet-02/DSA0031945.pdf>.

Intel (2002). Mobile Intel® Pentium® 4 Processor-M and Intel® 845MP/845MZ Chipset Platform Design Guide.

Asus (2001). A7V133 A7V133-C jumperfree PC133/VC133 200/266MHZ Fsb Agp Pro/4X Socket A Motherboard User’s Manual.

Trescases, O. (2007). Integrated Power-Supplies for Portable Applications.

Van Schaik, C. F. (2002). The Development of a Node for a Hardware Reconfigurable Parallel Processor (Master’s thesis, University of Cape Town).

McPartland, R. J., Loeper, D. J., Higgins, F. P., Singh, R., MacDonald, G., Komoriya, G., . . . & Leung, C. W. (2000, May). SRAM Embedded Memory with Low Cost, Flash EEPROM-Switch-Controlled Redundancy. In Proceedings of the IEEE 2000 Custom Integrated Circuits Conference (Cat. No. 00CH37044) (pp. 287-289). IEEE.

Crisp, R. (1997). Direct Rambus Technology: The New Main Memory Standard. IEEE Micro, 17(6), 18-28. PC2-4200/ PC2-3200 DDR2.

PC2-4200/PC2-3200 DDR2 Registered Mini-DIMM. (Dec. 2006). Design Spec Rev. 2.0. Retrievable from <https://www.jedec.org/standards-documents/docs/module-42014>.

* cited by examiner




FIG. 1
(PRIOR ART)





FIG. 3A


Spansion EcoRAM Configurations

256GB Spansion EcoRAM Solution – Single Accelerator



256GB Single Accelerator Spansion EcoRAM Solution

256GB Spansion EcoRAM Solution – Dual Accelerator



256GB Single Accelerator Spansion EcoRAM Solution

**FIG. 2
(PRIOR ART)**




FIG. 4B




FIG. 3B




FIG. 4A




FIG. 5A




FIG. 6




FIG. 5B




FIG. 7




FIG. 8A




FIG. 8B




FIG. 9




FIG. 10

<i>DRAM density (GB)</i>	<i># of blocks per bank</i>	<i>Flash wr-time to rd-time ratio</i>	<i>Avg block use time (sec)</i>	<i>Flash write time (sec)</i>	<i>Max allowed Closed Blk in queue to be written back to Flash</i>
1	250	55	1.00E-03	2.00E-02	0
1	250	55	1.00E-02	2.00E-02	2
1	250	55	2.00E-02	2.00E-02	5
1	250	55	5.00E-02	2.00E-02	11
2	500	55	1.00E-03	2.00E-02	0
2	500	55	1.00E-02	2.00E-02	5
2	500	55	2.00E-02	2.00E-02	9
2	500	55	5.00E-02	2.00E-02	23
4	1000	55	1.00E-03	2.00E-02	1
4	1000	55	1.00E-02	2.00E-02	9
4	1000	55	2.00E-02	2.00E-02	18
4	1000	55	5.00E-02	2.00E-02	45

FIG. 11




FIG. 13




FIG. 14




FIG. 15A




FIG. 15B




FIG. 15C




FIG. 16




FIG. 17




FIG. 18




FIG. 20




FIG. 21




FIG. 22

US 12,373,366 B2

1

**MEMORY WITH ON-MODULE POWER
MANAGEMENT****PRIORITY CLAIM**

This application is a continuation of U.S. patent application Ser. No. 17/328,019, filed May 24, 2021, titled "FLASH-DRAM HYBRID MEMORY MODULE", now U.S. Pat. No. 11,232,054, which is a continuation of U.S. patent application Ser. No. 17/138,766, filed Dec. 30, 2020, titled "Flash-Dram Hybrid Memory", now U.S. Pat. No. 11,016,918, which is a continuation of U.S. patent application Ser. No. 15/934,416, filed Mar. 23, 2018, titled "Flash-Dram Hybrid Memory Module," which is a continuation of U.S. patent application Ser. No. 14/840,865, filed Aug. 31, 2015, titled "Flash-Dram Hybrid Memory Module," now U.S. Pat. No. 9,928,186, which is a continuation of U.S. patent application Ser. No. 14/489,269, filed Sep. 17, 2014, titled "Flash-Dram Hybrid Memory Module," now U.S. Pat. No. 9,158,684, which is a continuation of U.S. patent application Ser. No. 13/559,476, filed Jul. 26, 2012, titled "Flash-Dram Hybrid Memory Module," now U.S. Pat. No. 8,874,831, which claims the benefit of U.S. Provisional Patent Application No. 61/512,871, filed Jul. 28, 2011, and is a continuation-in-part of U.S. patent application Ser. No. 12/240,916, filed Sep. 29, 2008, titled "Non-Volatile Memory Module," now U.S. Pat. No. 8,301,833, which is a continuation of U.S. patent application Ser. No. 12/131,873, filed Jun. 2, 2008, which claims the benefit of U.S. Provisional Patent Application No. 60/941,586, filed Jun. 1, 2007, the contents of all of which are incorporated herein by reference in their entirety.

This application may be considered related to U.S. patent application Ser. No. 14/173,242, titled "Isolation Switching For Backup Of Registered Memory," filed Feb. 5, 2014, which is a continuation of U.S. patent application Ser. No. 13/905,053, titled "Isolation Switching For Backup Of Registered Memory," filed May 29, 2013, now U.S. Pat. No. 8,677,060, issued Mar. 18, 2014, which is a continuation of U.S. patent application Ser. No. 13/536,173, titled "Data Transfer Scheme For Non-Volatile Memory Module," filed Jun. 28, 2012, now U.S. Pat. No. 8,516,187, issued Aug. 20, 2013, which is a divisional of U.S. patent application Ser. No. 12/240,916, titled "Non-Volatile Memory Module," filed Sep. 29, 2008, now U.S. Pat. No. 8,301,833, issued Oct. 30, 2012, which is a continuation of U.S. patent application Ser. No. 12/131,873, filed Jun. 2, 2008, now abandoned, which claims the benefit of U.S. Provisional Application No. 60/941,586, filed Jun. 1, 2007, the contents of which are incorporated by reference herein in their entirety.

This application may also be considered related to U.S. patent application Ser. No. 15/000,834, filed Jan. 19, 2016, which is a continuation of U.S. patent application Ser. No. 14/489,332, filed Sep. 17, 2014, now U.S. Pat. No. 9,269,437, which is a continuation of U.S. patent application Ser. No. 14/173,219, filed Feb. 5, 2014, now U.S. Pat. No. 8,904,099, which is a continuation of U.S. patent application Ser. No. 13/905,048, filed May 29, 2013, now U.S. Pat. No. 6,671,243, which is a continuation U.S. patent application Ser. No. 13/536,173 above.

This application may also be considered related to U.S. patent application Ser. No. 15/924,866, which is a continuation of U.S. patent application Ser. No. 14/489,281, filed Sep. 17, 2014, now U.S. Pat. No. 9,921,762, which is a continuation of U.S. patent application Ser. No. 13/625,563, filed Sep. 24, 2012, now U.S. Pat. No. 8,904,098, which

2

claims the benefit of U.S. Provisional Application No. 61/583,775, filed Sep. 23, 2011.

TECHNICAL FIELD

The present disclosure relates generally to computer memory devices, and more particularly, to devices that employ different types of memory devices such as combinations of Flash and random access memories.

BACKGROUND

As technology advances and the usage of portable computing devices, such as tablet notebook computers, increases, more data needs to be transferred among data centers and to/from end users. In many cases, data centers are built by clustering multiple servers that are networked to increase performance.

Although there are many types of networked servers that are specific to the types applications envisioned, the basic concept is generally to increase server performance by dynamically allocating computing and storage resources. In recent years, server technology has evolved to be specific to particular applications such as 'finance transactions' (for example, point-of-service, inter-bank transaction, stock market transaction), 'scientific computation' (for example, fluid dynamic for automobile and ship design, weather prediction, oil and gas expeditions), 'medical diagnostics' (for example, diagnostics based on the fuzzy logic, medical data processing), 'simple information sharing and searching' (for example, web search, retail store website, company home page), 'email' (information distribution and archive), 'security service', 'entertainment' (for example, video-on-demand), and so on. However, all of these applications suffer from the same information transfer bottleneck due to the inability of a high speed CPU (central processing unit) to efficiently transfer data in and out of relatively slower speed storage or memory subsystems, particularly since data transfers typically pass through the CPU input/output (I/O) channels.

The data transfer limitations by the CPU are exemplified by the arrangement shown in FIG. 1, and apply to data transfers between main storage (for example the hard disk (HD) or solid state drive (SSD) and the memory subsystems (for example DRAM DIMM (Dynamic Random Access Memory Dual In-line Memory Module) connected to the front side bus (FSB)). In arrangements such as that of FIG. 1, the SSD/HD and DRAM DIMM of a conventional memory arrangement are connected to the CPU via separate memory control ports (not shown). FIG. 1 specifically shows, through the double-headed arrow, the data flow path between the computer or server main storage (SSD/HD) to the DRAM DIMMs. Since the SSD/HD data I/O and the DRAM DIMM data I/O are controlled by the CPU, the CPU needs to allocate its process cycles to control these I/Os, which may include the IRQ (Interrupt Request) service which the CPU performs periodically. As will be appreciated, the more time a CPU allocates to controlling the data transfer traffic, the less time the CPU has to perform other tasks. Therefore, the overall performance of a server will deteriorate with the increased amount of time the CPU has to expend in performing data transfer.

There have been various approaches to increase the data transfer throughput rates from/to the main storage, such as SSD/HD, to local storage, such as DRAM DIMM. In one example as illustrated in FIG. 2, EcoRAM™ developed by Spansion provides a storage SSD based system that assumes

US 12,373,366 B2

3

a physical form factor of a DIMM. The EcoRAM™ is populated with Flash memories and a relatively small memory capacity using DRAMs which serve as a data buffer. This arrangement is capable of delivering higher throughput rate than a standard SSD based system since the EcoRAM™ is connected to the CPU (central processing unit) via a high speed interface, such as the HT (Hyper Transport) interface, while an SSD/HD is typically connected via SATA (serial AT attachment), USB (universal serial bus), or PCI Express (peripheral component interface express). For example, the read random access throughput rate of EcoRAM™ is near 3 GB/s compared with 400 MB/s for a NAND SSD memory subsystem using the standard PCI Express-based. This is a 7.5× performance improvement. However, the performance improvement for write random access throughput rate is less than 2× (197 MBs for the EcoRAM vs. 104 MBs for NAND SSD). This is mainly due to the fact that the write speed is cannot be faster than the NAND Flash write access time. FIG. 2 is an example of EcoRAM™ using SSD with the form factor of a standard DIMM such that it can be connected to the FSB (front side bus). However, due to the interface protocol difference between DRAM and Flash, an interface device, EcoRAM Accelerator™, which occupies one of the server's CPU sockets is used, and hence further reducing server's performance by reducing the number of available CPU sockets available, and in turn reducing the overall computation efficiency. The server's performance will further suffer due to the limited utilization of the CPU bus due to the large difference in the data transfer throughput rate between read and write operations.

The EcoRAM™ architecture enables the CPU to view the Flash DIMM controller chip as another processor with a large size of memory available for CPU access.

In general, the access speed of a Flash based system is limited by four items: the read/write speed of the Flash memory, the CPU's FSB bus speed and efficiency, the Flash DIMM controller's inherent latency, and the HT interconnect speed and efficiency which is dependent on the HT interface controller in the CPU and Flash DIMM controller chip.

The published results indicate that these shortcomings are evident in that the maximum throughput rate is 1.56 GBs for the read operation and 104 MBs for the write operation. These access rates are 25% of the DRAM read access speed, and 1.7% of the DRAM access speed at 400 MHz operation. The disparity in the access speed (15 to 1) between the read operation and write operation highlight a major disadvantage of this architecture. The discrepancy of the access speed between this type of architecture and JEDEC standard DRAM DIMM is expected to grow wider as the DRAM memory technology advances much faster than the Flash memory.

Certain types of memory modules comprise a plurality of dynamic random-access memory (DRAM) devices mounted on a printed circuit board (PCB). These memory modules are typically mounted in a memory slot or socket of a computer system (e.g., a server system or a personal computer) and are accessed by the computer system to provide volatile memory to the computer system.

Volatile memory generally maintains stored information only when it is powered. Batteries have been used to provide power to volatile memory during power failures or interruptions. However, batteries may require maintenance, may need to be replaced, are not environmentally friendly, and the status of batteries can be difficult to monitor.

4

Non-volatile memory can generally maintain stored information while power is not applied to the non-volatile memory. In certain circumstances, it can therefore be useful to backup volatile memory using non-volatile memory.

OVERVIEW

Described herein is a memory module couplable to a memory controller of a host system. The memory module includes a non-volatile memory subsystem, a data manager coupled to the non-volatile memory subsystem, a volatile memory subsystem coupled to the data manager and operable to exchange data with the non-volatile memory subsystem by way of the data manager, and a controller operable to receive commands from the memory controller and to direct (i) operation of the non-volatile memory subsystem, (ii) operation of the volatile memory subsystem, and (iii) transfer of data between any two or more of the memory controller, the volatile memory subsystem, and the non-volatile memory subsystem based on at least one received command from the memory controller.

Also described herein is a method for managing a memory module by a memory controller, the memory module including volatile and non-volatile memory subsystems. The method includes receiving control information from the memory controller, wherein the control information is received using a protocol of the volatile memory subsystem. The method further includes identifying a data path to be used for transferring data to or from the memory module using the received control information, and using a data manager and a controller of the memory module to transfer data between any two or more of the memory controller, the volatile memory subsystem, and the non-volatile memory subsystem based on at least one of the received control information and the identified data path.

Also described herein is a memory module wherein the data manager is operable to control one or more of data flow rate, data transfer size, data buffer size, data error monitoring, and data error correction in response to receiving at least one of a control signal and control information from the controller.

Also described herein is a memory module wherein the data manager controls data traffic between any two or more of the memory controller, the volatile memory subsystem, and the non-volatile memory subsystem based on instructions received from the controller.

Also described herein is a memory module wherein data traffic control relates to any one or more of data flow rate, data transfer size, data buffer size, data transfer bit width, formatting information, direction of data flow, and the starting time of data transfer.

Also described herein is a memory module wherein the controller configures at least one of a first memory address space of the volatile memory subsystem and a second memory address space of the non-volatile memory subsystem in response to at least one of a received command from the memory controller and memory address space initialization information of the memory module.

Also described herein is a memory module wherein the data manager is configured as a bi-directional data transfer fabric having two or more sets of data ports coupled to any one of the volatile and non-volatile memory subsystems.

Also described herein is a memory module wherein at least one of the volatile and non-volatile memory subsystems comprises one or more memory segments.

US 12,373,366 B2

5

Also described herein is a memory module wherein each memory segment comprises at least one memory circuit, memory device, or memory die.

Also described herein is a memory module wherein the volatile memory subsystem comprises DRAM memory.

Also described herein is a memory module wherein the non-volatile memory subsystem comprises flash memory.

Also described herein is a memory module wherein at least one set of data ports is operated by the data manager to independently and/or concurrently transfer data to or from one or more memory segments of the volatile or non-volatile memory subsystems.

Also described herein is a memory module wherein the data manager and controller are configured to effect data transfer between the memory controller and the non-volatile memory subsystem in response to memory access commands received by the controller from the memory controller.

Also described herein is a memory module wherein the volatile memory subsystem is operable as a buffer for the data transfer between the memory controller and non-volatile memory.

Also described herein is a memory module wherein the data manager further includes a data format module configured to format data to be transferred between any two or more of the memory controller, the volatile memory subsystem, and the non-volatile memory subsystem based on control information received from the controller.

Also described herein is a memory module wherein the data manager further includes a data buffer for buffering data delivered to or from the non-volatile memory subsystem.

Also described herein is a memory module wherein the controller is operable to perform one or more of memory address translation, memory address mapping, address domain conversion, memory access control, data error correction, and data width modulation between the volatile and non-volatile memory subsystems.

Also described herein is a memory module wherein the controller is configured to effect operation with the host system in accordance with a prescribed protocol.

Also described herein is a memory module wherein the prescribed protocol is selected from one or more of DDR, DDR2, DDR3, and DDR4 protocols.

Also described herein is a memory module wherein the controller is operable to configure memory space in the memory module based on at least one of a command received from the memory controller, a programmable value written into a register, a value corresponding to a first portion of the volatile memory subsystem, a value corresponding to a first portion of the non-volatile memory subsystem, and a timing value.

Also described herein is a memory module wherein the controller configures the memory space of the memory module using at least a first portion of the volatile memory subsystem and a first portion of the non-volatile memory subsystem, and the controller presents a unified memory space to the memory controller.

Also described herein is a memory module wherein the controller configures the memory space in the memory module using partitioning instructions that are application-specific.

Also described herein is a memory module wherein the controller is operable to copy booting information from the non-volatile to the volatile memory subsystem during power up.

Also described herein is a memory module wherein the controller includes a volatile memory control module, a

6

non-volatile memory control module, data manager control module, a command interpreter module, and a scheduler module.

Also described herein is a memory module wherein commands from the volatile memory control module to the volatile memory subsystem are subordinated to commands from the memory controller to the controller.

Also described herein is a memory module wherein the controller effects pre-fetching of data from the non-volatile to the volatile memory.

Also described herein is a memory module wherein the pre-fetching is initiated by the memory controller writing an address of requested data into a register of the controller.

Also described herein is a memory module wherein the controller is operable to initiate a copy operation of data of a closed block in the volatile memory subsystem to a target block in the non-volatile memory subsystem.

Also described herein is a memory module wherein, if the closed block is re-opened, the controller is operable to abort the copy operation and to erase the target block from the non-volatile memory subsystem.

Also described herein is a method for managing a memory module wherein the transfer of data includes a bidirectional transfer of data between the non-volatile and the volatile memory subsystems.

Also described herein is a method for managing a memory module further comprising operating the data manager to control one or more of data flow rate, data transfer size, data width size, data buffer size, data error monitoring, data error correction, and the starting time of the transfer of data.

Also described herein is a method for managing a memory module further comprising operating the data manager to control data traffic between the memory controller and at least one of the volatile and non-volatile memory subsystems.

Also described herein is a method for managing a memory module wherein data traffic control relates to any one or more of data transfer size, formatting information, direction of data flow, and the starting time of the transfer of data.

Also described herein is a method for managing a memory module wherein data traffic control by the data manager is based on instructions received from the controller.

Also described herein is a method for managing a memory module further comprising operating the data manager as a bi-directional data transfer fabric with two or more sets of data ports coupled to any one of the volatile and non-volatile memory subsystems.

Also described herein is a method for managing a memory module wherein at least one of the volatile and non-volatile memory subsystems comprises one or more memory segments.

Also described herein is a method for managing a memory module wherein each memory segment comprises at least one memory circuit, memory device, or memory die.

Also described herein is a method for managing a memory module wherein the volatile memory subsystem comprises DRAM memory.

Also described herein is a method for managing a memory module wherein the non-volatile memory subsystem comprises Flash memory.

Also described herein is a method for managing a memory module further comprising operating the data ports

US 12,373,366 B2

7

to independently and/or concurrently transfer data to or from one or more memory segments of the volatile or non-volatile memory subsystems.

Also described herein is a method for managing a memory module further comprising directing transfer of data bi-directionally between the volatile and non-volatile memory subsystems using the data manager and in response to memory access commands received by the controller from the memory controller.

Also described herein is a method for managing a memory module further comprising buffering the data transferred between the memory controller and non-volatile memory subsystem using the volatile memory subsystem.

Also described herein is a method for managing a memory module further comprising using the controller to perform one or more of memory address translation, memory address mapping, address domain conversion, memory access control, data error correction, and data width modulation between the volatile and non-volatile memory subsystems.

Also described herein is a method for managing a memory module further comprising using the controller to effect communication with a host system by the volatile memory subsystem in accordance with a prescribed protocol.

Also described herein is a method for managing a memory module wherein the prescribed protocol is selected from one or more of DDR, DDR2, DDR3, and DDR4 protocols.

Also described herein is a method for managing a memory module further comprising using the controller to configure memory space in the memory module based on at least one of a command received from the memory controller, a programmable value written into a register, a value corresponding to a first portion of the volatile memory subsystem, a value corresponding to a first portion of the non-volatile memory subsystem, and a timing value.

Also described herein is a method for managing a memory module wherein the controller configures the memory space of the memory module using at least a first portion of the volatile memory subsystem and a first portion of the non-volatile memory subsystem, and the controller presents a unified memory space to the memory controller.

Also described herein is a method for managing a memory module wherein the controller configures the memory space in the memory module using partitioning instructions that are application-specific.

Also described herein is a method for managing a memory module further comprising using the controller to copy booting information from the non-volatile to the volatile memory subsystem during power up.

Also described herein is a method for managing a memory module wherein the controller includes a volatile memory control module, the method further comprising generating commands by the volatile memory control module in response to commands from the memory controller, and transmitting the generated commands to the volatile memory subsystem.

Also described herein is a method for managing a memory module further comprising pre-fetching of data from the non-volatile memory subsystem to the volatile memory subsystem.

Also described herein is a method for managing a memory module wherein the pre-fetching is initiated by the memory controller writing an address of requested data into a register of the controller.

8

Also described herein is a method for managing a memory module further comprising initiating a copy operation of data of a closed block in the volatile memory subsystem to a target block in the non-volatile memory subsystem.

Also described herein is a method for managing a memory module further comprising aborting the copy operation when the closed block of the volatile memory subsystem is re-opened, and erasing the target block in the non-volatile memory subsystem.

Also described herein is a memory system having a volatile memory subsystem, a non-volatile memory subsystem, a controller coupled to the non-volatile memory subsystem, and a circuit coupled to the volatile memory subsystem, to the controller, and to a host system. In a first mode of operation, the circuit is operable to selectively isolate the controller from the volatile memory subsystem, and to selectively couple the volatile memory subsystem to the host system to allow data to be communicated between the volatile memory subsystem and the host system. In a second mode of operation, the circuit is operable to selectively couple the controller to the volatile memory subsystem to allow data to be communicated between the volatile memory subsystem and the nonvolatile memory subsystem using the controller, and the circuit is operable to selectively isolate the volatile memory subsystem from the host system.

Also described herein is a method for operating a memory system. The method includes coupling a circuit to a host system, a volatile memory subsystem, and a controller, wherein the controller is coupled to a non-volatile memory subsystem. In a first mode of operation that allows data to be communicated between the volatile memory subsystem and the host system, the circuit is used to (i) selectively isolate the controller from the volatile memory subsystem, and (ii) selectively couple the volatile memory subsystem to the host system. In a second mode of operation that allows data to be communicated between the volatile memory subsystem and the nonvolatile memory subsystem via the controller, the circuit is used to (i) selectively couple the controller to the volatile memory subsystem, and (ii) selectively isolate the volatile memory subsystem from the host system.

Also described herein is a nontransitory computer readable storage medium storing one or more programs configured to be executed by one or more computing devices. The programs, when executing on the one or more computing devices, cause a circuit that is coupled to a host system, to a volatile memory subsystem, and to a controller that is coupled to a nonvolatile memory subsystem, to perform a method in which, in a first mode of operation that allows data to be communicated between the volatile memory subsystem and the host system, operating the circuit to (i) selectively isolate the controller from the volatile memory subsystem, and (ii) selectively couple the volatile memory subsystem to the host system. In a second mode of operation that allows data to be communicated between the volatile memory subsystem and the nonvolatile memory subsystem via the controller, operating the circuit to (i) selectively couple the controller to the volatile memory subsystem, and (ii) selectively isolate the volatile memory subsystem from the host system.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated into and constitute a part of this specification, illustrate one or more examples of embodiments and, together with the

US 12,373,366 B2

9

description of example embodiments, serve to explain the principles and implementations of the embodiments.

In the drawings:

FIG. 1 is a block diagram illustrating the path of data transfer, via a CPU, of a conventional memory arrangement;

FIG. 2 is a block diagram of a known EcoRAM™ architecture;

FIGS. 3A and 3B are block diagrams of a non-volatile memory DIMM or NVDIMM;

FIGS. 4A and 4B are block diagrams of a Flash-DRAM hybrid DIMM or FDHDIMM;

FIG. 5A is a block diagram of a memory module 500 in accordance with certain embodiments described herein;

FIG. 5B is a block diagram showing some functionality of a memory module such as that shown in FIG. 5A;

FIG. 6 is a block diagram showing some details of the data manager (DMgr);

FIG. 7 is a functional block diagram of the on-module controller (CDC);

FIG. 8A is a block diagram showing more details of the prior art Flash-DRAM hybrid DIMM (FDHDIMM) of FIGS. 4A and 4B;

FIG. 8B is a block diagram of a Flash-DRAM hybrid DIMM (FDHDIMM) in accordance with certain embodiments disclosed herein;

FIG. 9 is a flow diagram directed to the transfer of data from Flash memory to DRAM memory and vice versa in an exemplary FDHDIMM;

FIG. 10 is a block diagram showing an example of mapping of DRAM address space to Flash memory address space; and

FIG. 11 is a table showing estimates of the maximum allowed closed blocks in a queue to be written back to Flash memory for different DRAM densities using various average block use time.

FIG. 12 is a block diagram of an example memory system compatible with certain embodiments described herein.

FIG. 13 is a block diagram of an example memory module with ECC (error-correcting code) having a volatile memory subsystem with nine volatile memory elements and a non-volatile memory subsystem with five non-volatile memory elements in accordance with certain embodiments described herein.

FIG. 14 is a block diagram of an example memory module having a microcontroller unit and logic element integrated into a single device in accordance with certain embodiments described herein.

FIGS. 15A-15C schematically illustrate example embodiments of memory systems having volatile memory subsystems comprising registered dual in-line memory modules in accordance with certain embodiments described herein.

FIG. 16 schematically illustrates an example power module of a memory system in accordance with certain embodiments described herein.

FIG. 17 is a flowchart of an example method of providing a first voltage and a second voltage to a memory system including volatile and non-volatile memory subsystems.

FIG. 18 is a flowchart of an example method of controlling a memory system operatively coupled to a host system and which includes at least 100 percent more storage capacity in non-volatile memory than in volatile memory.

FIG. 19 schematically illustrates an example clock distribution topology of a memory system in accordance with certain embodiments described herein.

FIG. 20 is a flowchart of an example method of controlling a memory system operatively coupled to a host system,

10

the method including operating a volatile memory subsystem at a reduced rate in a back-up mode.

FIG. 21 schematically illustrates an example topology of a connection to transfer data slices from two DRAM segments of a volatile memory subsystem of a memory system to a controller of the memory system.

FIG. 22 is a flowchart of an example method of controlling a memory system operatively coupled to a host system, the method including backing up and/or restoring a volatile memory subsystem in slices.

DESCRIPTION OF EXAMPLE EMBODIMENTS

Example embodiments are described herein in the context of a system of computers, servers, controllers, memory modules, hard disk drives and software. Those of ordinary skill in the art will realize that the following description is illustrative only and is not intended to be in any way limiting. Other embodiments will readily suggest themselves to such skilled persons having the benefit of this disclosure. Reference will now be made in detail to implementations of the example embodiments as illustrated in the accompanying drawings. The same reference indicators will be used to the extent possible throughout the drawings and the following description to refer to the same or like items.

In the interest of clarity, not all of the routine features of the implementations described herein are shown and described. It will, of course, be appreciated that in the development of any such actual implementation, numerous implementation-specific decisions must be made in order to achieve the developer's specific goals, such as compliance with application- and business-related constraints, and that these specific goals will vary from one implementation to another and from one developer to another. Moreover, it will be appreciated that such a development effort might be complex and time-consuming, but would nevertheless be a routine undertaking of engineering for those of ordinary skill in the art having the benefit of this disclosure.

In accordance with this disclosure, the components, process steps, and/or data structures described herein may be implemented using various types of operating systems, computing platforms, computer programs, and/or general purpose machines. In addition, those of ordinary skill in the art will recognize that devices of a less general purpose nature, such as hardwired devices, field programmable gate arrays (FPGAs), application specific integrated circuits (ASICs), or the like, may also be used without departing from the scope and spirit of the inventive concepts disclosed herein. Where a method comprising a series of process steps is implemented by a computer or a machine and those process steps can be stored as a series of instructions readable by the machine, they may be stored on a tangible medium such as a computer memory device (e.g., ROM (Read Only Memory), PROM (Programmable Read Only Memory), EEPROM (Electrically Erasable Programmable Read Only Memory), Flash memory, Jump Drive, and the like), magnetic storage medium (e.g., tape, magnetic disk drive, and the like), optical storage medium (e.g., CD-ROM, DVD-ROM, paper card, paper tape and the like) and other types of program memory.

The term "exemplary" where used herein is intended to mean "serving as an example, instance or illustration." Any embodiment described herein as "exemplary" is not necessarily to be construed as preferred or advantageous over other embodiments.

Disclosed herein are arrangements for improving memory access rates and addressing the high disparity (15 to 1 ratio)

US 12,373,366 B2

11

between the read and write data throughput rates. In one arrangement, a Flash-DRAM-hybrid DIMM (FDHDIMM) with integrated Flash and DRAM is used. Methods for controlling such an arrangement are described.

In certain embodiments, the actual memory density (size or capacity) of the DIMM and/or the ratio of DRAM memory to Flash memory are configurable for optimal use with a particular application (for example, POS, inter-bank transaction, stock market transaction, scientific computation such as fluid dynamics for automobile and ship design, weather prediction, oil and gas expeditions, medical diagnostics such as diagnostics based on the fuzzy logic, medical data processing, simple information sharing and searching such as web search, retail store website, company home page, email or information distribution and archive, security service, and entertainment such as video-on-demand).

In certain embodiments, the device contains a high density Flash memory with a low density DRAM, wherein the DRAM is used as a data buffer for read/write operation. The Flash serves as the main memory. Certain embodiments described herein overcome the needs of having a long separation period between an Activate command (may be referred to as RAS) and a corresponding read or write command (may be referred to as first CAS command).

In accordance with one embodiment, described with reference to FIGS. 3A and 3B, a memory system 300 includes a non-volatile (for example Flash) memory subsystem 302 and a volatile (for example DRAM) memory subsystem 304. The examples of FIGS. 3A and 3B are directed to architectures of a non-volatile DIMM (NVDIMM) NVDIMM system that may use a power subsystem (not shown) that can include a battery or a capacitor as a means for energy storage to copy DRAM memory data into Flash memory when power loss occurs, is detected, or is anticipated to occur during operation. When normal power is restored, a restore NVDIMM operation is initiated and the data stored in the Flash memory is properly restored to the DRAM memory. In this architecture, the density of the Flash is about the same as the DRAM memory size or within a few multiples, although in some applications it may be higher. This type of architecture may also be used to provide non-volatile storage that is connected to the FSB (front side bus) to support RAID (Redundant Array of Independent Disks) based systems or other type of operations. An NVDIMM controller 306 receives and interprets commands from the system memory controller hub (MCH). The NVDIMM controller 306 control the NVDIMM DRAM and Flash memory operations. In FIG. 3A, the DRAM 304 communicates data with the MCH, while an internal bus 308 is used for data transfer between the DRAM and Flash memory subsystems. In FIG. 3B, the NVDIMM controller 306' of NVDIMM 300' monitors events or commands and enables data transfer to occur in a first mode between the DRAM 304' and Flash 302' or in a second mode between the DRAM and the MCH.

In accordance with one embodiment, a general architecture for a Flash and DRAM hybrid DIMM (FDHDIMM) system 400 is shown in FIG. 4A. The FDHDIMM interfaces with an MCH (memory controller hub) to operate and behave as a high density DIMM, wherein the MCH interfaces with the non-volatile memory subsystem (for example Flash) 402 is controlled by an FDHDIMM controller 404. Although the MCH interfaces with the Flash via the FDHDIMM controller, the FDHDIMM overall performance is governed by the Flash access time. The volatile memory subsystem (for example DRAM) 406 is primarily used as a data buffer or a temporary storage location such that data from the Flash memory 402 is transferred to the DRAM 406

12

at the Flash access speed, and buffered or collected into the DRAM 406, which then transfers the buffered data to the MCH based on the access time of DRAM. Similarly, when the MCH transfers data to the DRAM 406, the FDHDIMM controller 404 manages the data transfer from the DRAM 406 to the Flash 402. Since the Flash memory access speed (both read and write) is relatively slower than DRAM, (e.g. for example a few hundred microseconds for read access), the average data throughput rate of FDHDIMM 400 is limited by the Flash access speed. The DRAM 406 serves as a data buffer stage that buffers the MCH read or write data. Thus, the DRAM 406 serves as a temporary storage for the data to be transferred from/to the Flash 402. Furthermore, in accordance with one embodiment, the MCH recognizes the physical density of an FDHDIMM operating as a high density DIMM as the density of Flash alone.

In accordance with one embodiment, a read operation can be performed by the MCH by sending an activate command (may be simply referred to as RAS, or row address strobe) to the FDHDIMM 400 to conduct a pre-fetch read data operation from the Flash 402 to the DRAM 406, with the pre-fetch data size being for example a page (1 KB or 2 KB, or may be programmable to any size). The MCH then sends a read command (may be simply referred to as CAS, or column address strobe) to read the data out input of the DRAM. In this embodiment, the data transfer from Flash to DRAM occurs at Flash access speed rates, while data transfer from DRAM to MCH occurs at DRAM access speed rates. In this example, data latency and throughput rates are the same as any DRAM operation as long as the read operations are executed onto the pages that were opened with the activate command previously sent to pre-fetch data from the Flash to DRAM. Thus, a longer separation time period between the RAS (e.g. Activate command) and the first CAS (column address strobe e.g. read or write command) is required to account for the time it takes to pre-fetch data from the Flash to DRAM.

An example of FDHDIMM operating as a DDR DIMM with SSD is shown in FIG. 4B, wherein the FDHDIMM 400' supports two different interface interpretations to the MCH. In the first interface interpretation, the MCH views the FDHDIMM 400' as a combination of DRAM DIMM and SSD (not illustrated). In this mode the MCH needs to manage two address spaces, one for the DRAMs 402' and one for the Flash 404'. The MCH is coupled to, and controls, both of the DRAM and Flash memory subsystems. One advantage of this mode is that the CPU does not need to be in the data path when data is moved from DRAM to Flash or from Flash to DRAM. In the second interface interpretation, the MCH views the FDHDIMM 400' as an on-DIMM Flash with the SSD in an extended memory space that is behind the DRAM space. Thus, in this mode, the MCH physically fetches data from the SSD to the DDR DRAM and then the DRAM sends the data to the MCH. Since all data movement occurs on the FDHDIMM, this mode will provide better performance than if the data were to be moved through or via the CPU.

In accordance with one embodiment and as shown in FIG. 4B, the FDHDIMM 400' receives control signals 408 from the MCH, where the control signals may include one or more control signals specifically for the DRAM 402' operation and one or more control signals specifically for the Flash 404' operation. In this embodiment, the MCH or CPU is coupled to the FDHDIMM via a single data bus interface 410 which couples the MCH to the DRAM.

FIGS. 5A and 5B are block diagrams of a memory module 500 that is couplable to a host system (not shown). The host

US 12,373,366 B2

13

system may be a server or any other system comprising a memory system controller or an MCH for providing and controlling the read/write access to one or more memory systems, wherein each memory system may include a plurality of memory subsystems, a plurality of memory devices, or at least one memory module. The term “read/write access” means the ability of the MCH to interface with a memory system or subsystem in order to write data into it or read data from it, depending on the particular requirement at a particular time.

In certain embodiments, memory module 500 is a Flash-DRAM hybrid memory subsystem which may be integrated with other components of a host system. In certain embodiments, memory module 500 is a Flash-DRAM hybrid memory module that has the DIMM (dual-inline memory module) form factor, and may be referred to as a FDHDIMM, although it is to be understood that in both structure and operation it may be different from the FDHDIMM discussed above and described with reference to FIGS. 4A and 4B. Memory module 500 includes two on-module intermediary components: a controller and a data manager. These on-module intermediary components may be physically separate components, circuits, or modules, or they may be integrated onto a single integrated circuit or device, or integrated with other memory devices, for example in a three dimensional stack, or in any one of several other possible expedients for integration known to those skilled in the art to achieve a specific design, application, or economic goal. In the case of a DIMM, these on-module intermediary components are an on-DIMM Controller (CDC) 502 and an on-DIMM data manager (DMgr) 504. While the DIMM form factor will predominate the discussion herein, it should be understood that this is for illustrative purposes only and memory systems using other form factors are contemplated as well. CDC 502 and data manager DMgr 504 are operative to manage the interface between a non-volatile memory subsystem such as a Flash 506, a volatile memory subsystem such as a DRAM 508, and a host system represented by MCH 510.

In certain embodiments, CDC 502 controls the read/write access to/from Flash memory 506 from/to DRAM memory 508, and to/from DRAM memory from/to MCH 510. Read/write access between DRAM 508, Flash 506 and MCH 510 may be referred to herein generally as communication, wherein control and address information C/A 560 is sent from MCH 510 to CDC 502, and possible data transfers follow as indicated by Data 550, Data 555, and/or Data 556. In certain embodiments, the CDC 502 performs specific functions for memory address transformation, such as address translation, mapping, or address domain conversion, Flash access control, data error correction, manipulation of data width or data formatting or data modulation between the Flash memory and DRAM, and so on. In certain embodiments, the CDC 502 ensures that memory module 500 provides transparent operation to the MCH in accordance with certain industry standards, such as DDR, DDR2, DDR3, DDR4 protocols. In the arrangement shown in FIGS. 5A and 5B, there is no direct access from the MCH 510 to the Flash 506 memory subsystem. Thus in accordance with certain embodiments, the Flash access speed has minimal impact on the overall FDHDIMM access speed. In the schematic illustration of FIG. 5B and in accordance with one embodiment, the CDC controller 502 receives standard DDR commands from the MCH, interprets, and produces commands and/or control signals to control the operation of the Data manager (DMgr), the Flash memory and the DRAM memory. The DMgr controls the data path routing

14

amongst DRAMs, Flash and MCH, as detailed below. The data path routing control signals are independently operated without any exclusivity.

An exemplary role of DMgr 504 is described with reference to FIG. 6. In certain embodiments and in response to communication from CDC 502, DMgr 504 provides a variety of functions to control data flow rate, data transfer size, data buffer size, data error monitoring or data error correction. For example, these functions or operations can be performed on-the-fly (while data is being transferred via the DMgr 504) or performed on buffered or stored data in DRAM or a buffer. In addition, one role of DMgr 504 is to provide interoperability among various memory subsystems or components and/or MCH 510.

In one embodiment, an exemplary host system operation begins with initialization. The CDC 502 receives a first command from the MCH 510 to initialize FDHDIMM 500 using a certain memory space. The memory space as would be controlled by MCH 510 can be configured or programmed during initialization or after initialization has completed. The MCH 510 can partition or parse the memory space in various ways that are optimized for a particular application that the host system needs to run or execute. In one embodiment, the CDC 502 maps the actual physical Flash 506 and DRAM 508 memory space using the information sent by MCH 510 via the first command. In one embodiment, the CDC 502 maps the memory address space of any one of the Flash 506 and DRAM 508 memory subsystems using memory address space information that is received from the host system, stored in a register within FDHDIMM 500, or stored in a memory location of a non-volatile memory subsystem, for example a portion of Flash 506 or a separate non-volatile memory subsystem. In one embodiment, the memory address space information corresponds to a portion of initialization information of the FDHDIMM 500.

In one embodiment, MCH 510 may send a command to restore a certain amount of data information from Flash 506 to DRAM 508. The CDC 502 provides control information to DMgr 504 to appropriately copy the necessary information from Flash 506 to the DRAM 508. This operation can provide support for various host system booting operations and/or a special host system power up operation.

In one embodiment, MCH 510 sends a command which may include various fields comprising control information regarding data transfer size, data format options, and/or startup time. CDC 502 receives and interprets the command and provides control signals to DMgr 504 to control the data traffic between the Flash 506, the DRAM 508, and the MCH 510. For example, DMgr 504 receives the data transfer size, formatting information, direction of data flow (via one or more multiplexers such as 611, 612, 621, 622 as detailed below), and the starting time of the actual data transfer from CDC 502. DMgr 504 may also receive additional control information from the CDC 502 to establish a data flow path and/or to correctly establish the data transfer fabric. In certain embodiments, DMgr 504 also functions as a bi-directional data transfer fabric. For example, DMgr 504 may have more than 2 sets of data ports facing the Flash 506 and the DRAM 508. Multiplexers 611 and 612 provide controllable data paths from any one of the DRAMs 508(1) and 508(2) (DRAM-A and DRAM-B) to any one of the MCH 510 and the Flash 506. Similarly multiplexers 621 and 622 provide controllable data paths from any one of the MCH and the Flash memory to any one of the DRAMs 508(1) and 508(2) (DRAM-A and DRAM-B). In one embodiment, DRAM 508(1) is a segment of DRAM 508, while in other

US 12,373,366 B2

15

embodiments, DRAM 508(1) is a separate DRAM memory subsystem. It will be understood that each memory segment can comprise one or more memory circuits, a memory devices, and/or memory integrated circuits. Of course other configurations for DRAM 508 are possible, and other data transfer fabrics using complex data paths and suitable types of multiplexing logic are contemplated.

In accordance with one embodiment, the two sets of multiplexors 611, 612 and 621, 622 allow independent data transfer to Flash 506 from DRAM-A 508(1) and DRAM-B 508(2). For example, in response to one or more control signals or a command from CDC 502, DMgr 504 can transfer data from DRAM-A 508(1) to MCH 510, via multiplexer 611, at the same time as from DRAM-B 508(2) to the Flash 506, via multiplexer 612; or data is transferred from DRAM-B 508(2) to MCH 510, via multiplexer 611, and simultaneously data is transferred from the Flash 506 to DRAM-A 508(1), via multiplexer 621. Further, in the same way that data can be transferred to or from the DRAM in both device-wide or segment-by-segment fashion, data can be transferred to or from the flash memory in device-wide or segment-by-segment fashion, and the flash memory can be addressed and accessed accordingly.

In accordance with one embodiment the illustrated arrangement of data transfer fabric of DMgr 504 also allows the CDC 502 to control data transfer from the Flash memory to the MCH by buffering the data from the Flash 506 using a buffer 602, and matching the data rate and/or data format of MCH 510. The buffer 602 is shown in FIG. 6 as a portion of a data format module 604; however, buffer 602 may also be a distributed buffer such that one buffer is used for each one of the set of multiplexer logic elements shown as multiplexers 611, 612, 621, and 622. Various buffer arrangements may be used, such as a programmable size buffer to meet the requirement of a given system design requirement, for example the disparity between read/write access time; or overall system performance, for example latency. In certain embodiments, the buffer 604 may introduce one or more clock cycle delays into a data communication path between MCH 510, DRAM 508, and Flash 506.

In certain embodiments, data format module 604 contains a data formatting subsystem (not shown) to enable DMgr 504 to format and perform data transfer in accordance with control information received from CDC 502. Data buffer 604 of data format module 602, discussed above, also supports a wide data bus 606 coupled to the Flash memory 506 operating at a first frequency, while receiving data from DRAM 508 using a relatively smaller width data bus 608 operating at a second frequency, the second frequency being larger than the first frequency in certain embodiments. The buffer 602 is designed to match the data flow rate between the DRAM 508 and the Flash 506.

A register 690 provides the ability to register commands received from MCH 510 via C/A 560 (FIG. 5A). The register 690 may communicate these commands to CDC 502 and/or to the DRAM 508 and/or Flash 506. The register 690 communicates these registered commands to CDC 502 for processing. The register 690 may also include multiple registers (not shown), such that it can provide the ability to register multiple commands, a sequence of commands, or provide a pipeline delay stage for buffering and providing a controlled execution of certain commands received from MCH 510.

In certain embodiments, the register 690 may register commands from MCH 510 and transmit the registered commands to DRAM 508 and/or Flash 506 memory subsystems. In certain embodiments, the CDC 502 monitors

16

commands received from MCH 510, via control and address bus C/A 560, and provides appropriate control information to DMgr 504, DRAM 508, or Flash 506 to execute these commands and perform data transfer operations between MCH 510 and FDHDIMM 500 via MCH data bus 610.

FIG. 7 illustrates a functional block diagram of the CDC 502. In certain embodiments, the major functional blocks of the CDC 502 are a DRAM control block DRAMCtrl 702, Flash control block FlashCtrl 704, MCH command interpreter CmdInt 706, DRAM-Flash interface scheduler Scheduler 708, and DMgr control block (DMgrCtrl) 710.

In accordance with one embodiment, DRAMCtrl 702 generates DRAM commands that are independent from the commands issued by the MCH 510. In accordance with one embodiment, when the MCH 510 initiates a read/write operation from/to the same DRAM 508 that is currently executing a command from the DRAMCtrl 702, then the CDC 502 may choose to instruct DRAMCtrl 702 to abort its operation in order to execute the operation initiated by the MCH. However, the CDC 502 may also pipeline the operation so that it causes DRAMCtrl 702 to either halt or complete its current operation prior to executing that of the MCH. The CDC 502 may also instruct DRAMCtrl 702 to resume its operation once the command from MCH 510 is completed.

In accordance with one embodiment, the FlashCtrl 704 generates appropriate Flash commands for the proper read/write operations. The CmdInt 706 intercepts commands received from MCH 510 and generates the appropriate control information and control signals and transmit them to the appropriate FDHDIMM functional block. For example, CmdInt 706 issues an interrupt signal to the DRAMCtrl 702 when the MCH issues a command that collides (conflicts) with the currently executing or pending commands that DRAMCtrl 702 has initiated independently from MCH 510, thus subordinating these commands to those from the MCH. The Scheduler 708 schedules the Flash-DRAM interface operation such that there is no resource conflict in the DMgr 504. In accordance with one embodiment, the Scheduler 708 assigns time slots for the DRAMCtrl 702 and FlashCtrl 704 operation based on the current status and the pending command received or to be received from the MCH. The DMgrCtrl 710 generates and sends appropriate control information and control signals for the proper operation and control of the data transfer fabric to enable or disable data paths between Flash 506, DRAM 508, and the MCH 510.

FIG. 8A is a block diagram showing a Flash-DRAM hybrid DIMM (FDHDIMM). As seen from FIG. 8A, this Flash-DRAM hybrid DIMM requires two separate and independent address buses to separately control the address spaces: one for the Flash memory Flash 506 and the other for the DRAM memory DRAM 508. The MCH treats the DRAM 508 and Flash 506 as separate memory subsystems, for example DRAM and SSD/HD memory subsystems. The memory in each address space is controlled directly by the MCH. However, the on-DIMM data path between Flash 506 and DRAM 508 allows for direct data transfer to occur between the Flash 506 and the DRAM 508 in response to control information from Ctrl 502. In this embodiment, this data transfer mechanism provides direct support for executing commands from the MCH without having the MCH directly controlling the data transfer, and thus improving data transfer performance from Flash 506 to the DRAM 508. However, the MCH needs to manage two address spaces and two different memory protocols simultaneously. Moreover, the MCH needs to map the DRAM memory space into the Flash memory space, and the data interface time suffers due

US 12,373,366 B2

17

to the difference in the data access time between the Flash memory and the DRAM memory.

In accordance with one embodiment, a memory space mapping of a Flash-DRAM hybrid DIMM is shown in FIG. 8B. A memory controller of a host system (not shown) controls both of the DRAM 508 address space and the Flash 506 address space using a single unified address space. The CDC 502 receives memory access commands from the MCH and generates control information for appropriate mapping and data transfer between Flash and DRAM memory subsystem to properly carry out the memory access commands. In one embodiment, the memory controller of the host system views the large Flash memory space as a DRAM memory space, and accesses this unified memory space with a standard DDR (double data rate) protocol used for accessing DRAM. The unified memory space in this case can exhibit overlapping memory address space between the Flash 506 and the DRAM 508. The overlapping memory address space may be used as a temporary storage or buffer for data transfer between the Flash 506 and the DRAM 508. For example, the DRAM memory space may hold a copy of data from the selected Flash memory space such that the MCH can access this data normally via DDR memory access commands. The CDC 502 controls the operation of the Flash 506 and DRAM 508 memory subsystems in response to commands received from a memory controller of a host system.

In one embodiment, the unified memory space corresponds to a contiguous address space comprising a first portion of the address space of the Flash 506 and a first portion of the address space of the DRAM 508. The first portion of the address space of the Flash 506 can be determined via a first programmable register holding a first value corresponding to the desired Flash memory size to be used. Similarly, the first portion of the address space of the DRAM 508 can be determined via a second programmable register holding a second value corresponding to the desired DRAM memory size to be used. In one embodiment, any one of the first portion of the address space of the Flash 506 and the first portion of the address space of the DRAM 508 is determined via a first value corresponding to a desired performance or memory size, the first value being received by the CDC 502 via a command sent by memory controller of the host system.

In accordance with one embodiment, a flow diagram directed to the transfer of data from Flash memory to DRAM memory and vice versa in an exemplary FDHIMM is shown in FIG. 9. In certain embodiments, data transfer from the Flash 506 to the DRAM 508 occurs in accordance with memory access commands which the CDC 502 receives from the memory controller of the host system. In certain embodiments, the CDC 502 controls the data transfer from the DRAM 508 to the Flash 506 so as to avoid conflict with any memory operation that is currently being executed. For example, when all the pages in a particular DRAM memory block are closed. The CDC 502 partitions the DRAM memory space into a number of blocks for the purpose of optimally supporting the desired application. The controller can configure memory space in the memory module based on at least one of one or more commands received from the MCH, instructions received from the MCH, a programmable value written into a register, a value corresponding to a first portion of the volatile memory subsystem, a value corresponding to a first portion of the non-volatile memory subsystem, and a timing value. Furthermore, the block size can be configurable by the memory controller of the host system, such that the number pages in a block can be

18

optimized to support a particular application or a task. Furthermore, the block size may be configured on-the-fly, e.g. CDC 502 can receive instruction regarding a desired block size from the memory controller via a memory command, or via a programmable value.

In certain embodiments, a memory controller can access the memory module using a standard access protocol, such as JEDEC's DDR DRAM, by sending a memory access command to the CDC 502 which in turn determines what type of a data transfer operation it is and the corresponding target address where the data information is stored, e.g. data information is stored in the DRAM 508 or Flash 506 memory subsystems. In response to a read operation, if the CDC 502 determines that data information, e.g. a page (or block), does not reside in the DRAM 508 but resides in Flash 506, then the CDC 502 initiates and controls all necessary data transfer operations from Flash 506 to DRAM 508 and subsequently to the memory controller. In one embodiment, once the CDC 502 completes the data transfer operation of the requested data information from the Flash 506 to the DRAM 508, the CDC 502 alerts the memory controller to retrieve the data information from the DRAM 508. In one embodiment, the memory controller initiates the copying of data information from Flash 506 to DRAM 508 by writing, into a register in the CDC 502, the target Flash address along with a valid block size. The CDC 502 in turn, executes appropriate operations and generates control information to copy the data information to the DRAM 508. Consequently, the memory controller can access or retrieve the data information using standard memory access commands or protocol.

An exemplary flow chart is shown in FIG. 9, a starting step or power up 902, is followed by an initialization step 904, the memory controller initiates, at step 906, a data move from the Flash 506 to the DRAM 508 by writing target address and size, to a control register in the CDC 502, which then copies, at 908, data information from the Flash 506 to the DRAM 508 and erases the block in the Flash. Erasing the data information from Flash may be accomplished independently from (or concurrently with) other steps that CDC 502 performs in this flow chart, i.e. other steps can be executed concurrently with the Erase the Flash block step. Once the data information or a block of data information is thus moved to the DRAM 508, the memory controller can operate on this data block using standard memory access protocol or commands at 910. The CDC 502 checks, at 912, if any of the DRAM 508 blocks, or copied blocks, are closed. If the memory controller closed any open blocks in DRAM 508, then the CDC 502 initiate a Flash write to write the closed block from the DRAM 508 to the Flash 506, at 914. In addition, the memory controller, at 916, reopens the closed block that is currently being written into the Flash 506, then the CDC 502 stops the Flash write operation and erases the Flash block which was being written to, as shown at 918. Otherwise, the CDC 502 continues and completes the writing operation to the Flash at 920.

The dashed lines in FIG. 9 indicate independent or parallel activities that can be performed by the CDC 502. At any time the CDC 502 receives a DRAM load command from a memory controller which writes a Flash target address and/or block size information into the RC register(s) at 922, as described above, then the CDC 502 executes a load DRAM w/RC step 906 and initiates another branch (or a thread) of activities that includes steps 908-922. In one embodiment, the CDC 502 controls the data transfer operations between DRAM 508 and Flash 506 such that the Flash 506 is completely hidden from the memory controller. The

US 12,373,366 B2

19

CDC 502 monitors all memory access commands sent by the memory controller using standard DRAM protocol and appropriately configures and manipulates both Flash 506 and DRAM 508 memory subsystems to perform the requested memory access operation and thus achieve the desired results. The memory controller does not interface directly with the Flash memory subsystem. Instead, the memory controller interfaces with the CDC 502 and/or DMgr 504 as shown in FIG. 5 and FIG. 6. Moreover, the memory controller may use one or more protocol, such as DDR, DDR2, DDR3, DDR4 protocols or the like.

In accordance with one embodiment, an example of mapping a DRAM address space to Flash memory address space is shown in FIG. 10. Two sets (1002, 1004) of address bits AD6 to AD17, forming a 24 bit extended memory page address, are allocated for the block address. For example, assuming a Block size of 256K Bytes, then a 24-bit block address space (using the two sets of AD6 to AD17 1002 and 1004) would enable access to 4 TB of Flash memory storage space. If a memory module has 1 GB of DRAM storage capacity, then it can hold approximately 4K Blocks of data in the DRAM memory, each Block comprise 256 K Bytes of data. The DRAM address space, corresponding to the 4K blocks, can be assigned to different virtual ranks and banks, where the number of virtual ranks and banks is configurable and can be manipulated to meet a specific design or performance needs. For example, if a 1 G Bytes memory module is configured to comprise two ranks with eight banks per rank, then each bank would hold two hundred fifty (250) blocks or the equivalent of 62 M Bytes or 62K pages, where each page correspond to a 1K Bytes. Other configurations using different page, block, banks, or ranks numbers may also be used. Furthermore, an exemplary mapping of 24-bit DDR DIMM block address to Flash memory address, using Block addressing as described above, is shown in FIG. 10. The 24-bit can be decomposed into fields, such as a logical unit number LUN address 1061 field, a Block address 1051 field, a Plane address 1041, a Page address 1031, and a group of least significant address bits A₀A₁ 1021. The Plane address 1041 is a sub address of the block address, and it may be used to support multiple page IO so as to improve Flash memory subsystem operation. In this example, it is understood that different number of bits may be allocated to each field of the 24-bit

The CDC 502 manages the block write-back operation by queuing the blocks that are ready to be written back to the Flash memory. As described above, if any page in a queued block for a write operation is reopened, then the CDC 502 will stop the queued block write operation, and remove the block from the queue. Once all the pages in a block are closed, then the CDC 502 restarts the write-back operation and queue the block for a write operation.

In accordance with one embodiment, an exemplary read operation from Flash 506 to DRAM 508 can be performed in approximately 400 μs, while a write operation from DRAM 508 to Flash 506 can be performed in approximately 22 ms resulting in a read to write ratio of 55 to 1. Therefore, if the average time a host system's memory controller spends accessing data information in a Block of DRAM is about 22 ms (that is the duration that a Block comprises one or more pages that are open), then the block write-back operation from DRAM to Flash would not impact performance and hence the disparity between read and write access may be completely hidden from the memory controller. If the block usage time is 11 ms instead of 22 ms, then the CDC 502 control the data transfer operation between DRAM 508 and Flash 506 such that there are no more than

20

9 closed blocks in the queue to be written-back to the Flash memory, hence approximately an average of 100 ms can be maintained for a standard DDR DRAM operation. Moreover, the number of closed Blocks in the queue to be written-back to the Flash memory subsystem varies with the average block usage time and the desired performance for a specific host system or for a specific application running using the host system resources.

Consequently, the maximum number of closed Blocks to be written-back to Flash can be approximated to be

$$\frac{(\text{#of blocks per bank})/(\text{ratio of 'Flash_block_write_time' to 'Flash_read_time'})}{(\text{Block usage time})/(\text{'Flash_block_write_time'})}$$

In order to maintain less than 100 ms time period for queued write-back Blocks, then using a Flash memory subsystem having 22 ms write access time per Block would result in a maximum number of four Blocks to be queued for write operation to Flash 506. Therefore, on average approximately 88 ms (=22 ms*4) for blocks means that each bank should not have more than four Blocks that need to be written back to the Flash 506.

The above equation also indicates that bigger DRAM memory space can support shorter block usage times. For example, 2 GB of DRAM memory allows the 8 closed blocks to be written-back to Flash. The table in FIG. 11 provides an estimation of the maximum allowed closed blocks in the queue to be written back to the Flash memory for different DRAM density using various average block use time.

Certain embodiments described herein include a memory system which can communicate with a host system such as a disk controller of a computer system. The memory system can include volatile and non-volatile memory, and a controller. The controller backs up the volatile memory using the non-volatile memory in the event of a trigger condition. Trigger conditions can include, for example, a power failure, power reduction, request by the host system, etc. In order to power the system in the event of a power failure or reduction, the memory system can include a secondary power source which does not comprise a battery and may include, for example, a capacitor or capacitor array.

In certain embodiments, the memory system can be configured such that the operation of the volatile memory is not adversely affected by the non-volatile memory or by the controller when the volatile memory is interacting with the host system. For example, one or more isolation devices may isolate the non-volatile memory and the controller from the volatile memory when the volatile memory is interacting with the host system and may allow communication between the volatile memory and the non-volatile memory when the data of the volatile memory is being restored or backed-up. This configuration generally protects the operation of the volatile memory when isolated while providing backup and restore capability in the event of a trigger condition, such as a power failure.

In certain embodiments described herein, the memory system includes a power module which provides power to the various components of the memory system from different sources based on a state of the memory system in relation to a trigger condition (e.g., a power failure). The power module may switch the source of the power to the various components in order to efficiently provide power in the event of the power failure. For example, when no power failure is detected, the power module may provide power to certain components, such as the volatile memory, from system power while charging a secondary power source (e.g., a capacitor array). In the event of a power failure or other

US 12,373,366 B2

21

trigger condition, the power module may power the volatile memory elements using the previously charged secondary power source.

In certain embodiments, the power module transitions relatively smoothly from powering the volatile memory with system power to powering it with the secondary power source. For example, the power system may power volatile memory with a third power source from the time the memory system detects that power failure is likely to occur until the time the memory system detects that the power failure has actually occurred.

In certain embodiments, the volatile memory system can be operated at a reduced frequency during backup and/or restore operations which can improve the efficiency of the system and save power. In some embodiments, during backup and/or restore operations, the volatile memory communicates with the non-volatile memory by writing and/or reading data words in bit-wise slices instead of by writing entire words at once. In certain embodiments, when each slice is being written to or read from the volatile memory the unused slice(s) of volatile memory is not active, which can reduce the power consumption of the system.

In yet other embodiments, the non-volatile memory can include at least 100 percent more storage capacity than the volatile memory. This configuration can allow the memory system to efficiently handle subsequent trigger conditions.

FIG. 12 is a block diagram of an example memory system 1010 compatible with certain embodiments described herein. The memory system 1010 can be coupled to a host computer system and can include a volatile memory subsystem 1030, a non-volatile memory subsystem 1040, and a controller 1062 operatively coupled to the non-volatile memory subsystem 1040. In certain embodiments, the memory system 1010 includes at least one circuit 1052 configured to selectively operatively decouple the controller 1062 from the volatile memory subsystem 1030.

In certain embodiments, the memory system 1010 comprises a memory module. The memory system 1010 may comprise a printed-circuit board (PCB) 1020. In certain embodiments, the memory system 1010 has a memory capacity of 512-MB, 1-GB, 2-GB, 4-GB, or 8-GB. Other volatile memory capacities are also compatible with certain embodiments described herein. In certain embodiments, the memory system 1010 has a non-volatile memory capacity of 512-MB, 1-GB, 2-GB, 4-GB, 8-GB, 16-GB, or 32-GB. Other non-volatile memory capacities are also compatible with certain embodiments described herein. In addition, memory systems 1010 having widths of 4 bytes, 8 bytes, 16 bytes, 32 bytes, or 32 bits, 64 bits, 128 bits, 256 bits, as well as other widths (in bytes or in bits), are compatible with embodiments described herein. In certain embodiments, the PCB 1020 has an industry-standard form factor. For example, the PCB 1020 can have a low profile (LP) form factor with a height of 30 millimeters and a width of 133.35 millimeters. In certain other embodiments, the PCB 1020 has a very high profile (VHP) form factor with a height of 50 millimeters or more. In certain other embodiments, the PCB 1020 has a very low profile (VLP) form factor with a height of 18.3 millimeters. Other form factors including, but not limited to, small-outline (SO-DIMM), unbuffered (UDIMM), registered (RDIMM), fully-buffered (FB-DIMM), miniDIMM, mini-RDIMM, VLP mini-DIMM, micro-DIMM, and SRAM DIMM are also compatible with certain embodiments described herein. For example, in other embodiments, certain non-DIMM form factors are possible

22

such as, for example, single in-line memory module (SIMM), multi-media card (MMC), and small computer system interface (SCSI).

In certain preferred embodiments, the memory system 1010 is in electrical communication with the host system. In other embodiments, the memory system 1010 may communicate with a host system using some other type of communication, such as, for example, optical communication. Examples of host systems include, but are not limited to, blade servers, I U servers, personal computers (PCs), and other applications in which space is constrained or limited. The memory system 1010 can be in communication with a disk controller of a computer system, for example. The PCB 1020 can comprise an interface 1022 that is configured to be in electrical communication with the host system (not shown). For example, the interface 1022 can comprise a plurality of edge connections which fit into a corresponding slot connector of the host system. The interface 1022 of certain embodiments provides a conduit for power voltage as well as data, address, and control signals between the memory system 1010 and the host system. For example, the interface 1022 can comprise a standard 240-pin DDR2 edge connector.

The volatile memory subsystem 1030 comprises a plurality of volatile memory elements 1032 and the non-volatile memory subsystem 1040 comprises a plurality of non-volatile memory elements 1042. Certain embodiments described herein advantageously provide nonvolatile storage via the non-volatile memory subsystem 1040 in addition to high-performance (e.g., high speed) storage via the volatile memory subsystem 1030. In certain embodiments, the first plurality of volatile memory elements 1032 comprises two or more dynamic random-access memory (DRAM) elements. Types of DRAM elements 1032 compatible with certain embodiments described herein include, but are not limited to, DDR, DDR2, DDR3, and synchronous DRAM (SDRAM). For example, in the block diagram of FIG. 12, the first memory bank 1030 comprises eight 64Mx8 DDR2 SDRAM elements 1032. The volatile memory elements 1032 may comprise other types of memory elements such as static random-access memory (SRAM). In addition, volatile memory elements 1032 having bit widths of 4, 8, 16, 32, as well as other bit widths, are compatible with certain embodiments described herein. Volatile memory elements 1032 compatible with certain embodiments described herein have packaging which include, but are not limited to, thin small-outline package (TSOP), ball-grid-array (BGA), fine-pitch BGA (FBOA), micro-BOA (1.1,BGA), mini-BGA (mBGA), and chip-scale packaging (CSP).

In certain embodiments, the second plurality of non-volatile memory elements 1042 comprises one or more flash memory elements. Types of flash memory elements 1042 compatible with certain embodiments described herein include, but are not limited to, NOR flash, NAND flash, ONE-NAND flash, and multi-level cell (MLC). For example, in the block diagram of FIG. 12, the second memory bank 1040 comprises 512 MB of flash memory organized as four 128 Mbx8 NAND flash memory elements 1042. In addition, nonvolatile memory elements 1042 having bit widths of 4, 8, 16, 32, as well as other bit widths, are compatible with certain embodiments described herein. Non-volatile memory elements 1042 compatible with certain embodiments described herein have packaging which include, but are not limited to, thin small-outline package (TSOP), ball-grid-array (BOA), fine-pitch BOA (FBGA), micro-BOA (POA), mini-BGA (mBGA), and chip-scale packaging (CSP).

US 12,373,366 B2

23

FIG. 13 is a block diagram of an example memory module 10 with ECC (error-correcting code) having a volatile memory subsystem 1030 with nine volatile memory elements 1032 and a non-volatile memory subsystem 1040 with five non-volatile memory elements 1042 in accordance with certain embodiments described herein. The additional memory element 1032 of the first memory bank 1030 and the additional memory element 1042 of the second memory bank 1040 provide the ECC capability. In certain other embodiments, the volatile memory subsystem 1030 comprises other numbers of volatile memory elements 1032 (e.g., 2, 3, 4, 5, 6, 7, more than 9). In certain embodiments, the non-volatile memory subsystem 1040 comprises other numbers of nonvolatile memory elements 1042 (e.g., 2, 3, more than 5).

Referring to FIG. 12, in certain embodiments, the logic element 1070 comprises a field-programmable gate array (FPGA). In certain embodiments, the logic element 1070 comprises an FPGA available from Lattice Semiconductor Corporation which includes an internal flash. In certain other embodiments, the logic element 1070 comprises an FPOA available from another vendor. The internal flash can improve the speed of the memory system 1010 and save physical space. Other types of logic elements 1070 compatible with certain embodiments described herein include, but are not limited to, a programmable-logic device (PLD), an application-specific integrated circuit (ASIC), a custom-designed semiconductor device, a complex programmable logic device (CPLD). In certain embodiments, the logic element 1070 is a custom device. In certain embodiments, the logic element 1070 comprises various discrete electrical elements, while in certain other embodiments, the logic element 1070 comprises one or more integrated circuits. FIG. 14 is a block diagram of an example memory module 1010 having a microcontroller unit 1060 and logic element 1070 integrated into a single controller 1062 in accordance with certain embodiments described herein. In certain embodiments, the controller 1062 includes one or more other components. For example, in one embodiment, an FPGA without an internal flash is used and the controller 1062 includes a separate flash memory component which stores configuration information to program the FPGA.

In certain embodiments, the at least one circuit 1052 comprises one or more switches coupled to the volatile memory subsystem 1030, to the controller 1062, and to the host computer (e.g., via the interface 1022, as schematically illustrated by FIGS. 12-14). The one or more switches are responsive to signals (e.g., from the controller 1062) to selectively operatively decouple the controller 1062 from the volatile memory subsystem 1030 and to selectively operatively couple the controller 1062 to the volatile memory subsystem 1030. In addition, in certain embodiments, the at least one circuit 1052 selectively operatively couples and decouples the volatile memory subsystem 1030 and the host system.

In certain embodiments, the volatile memory subsystem 1030 can comprise a registered DIMM subsystem comprising one or more registers 1160 and a plurality of DRAM elements 1180, as schematically illustrated by FIG. 15A. In certain such embodiments, the at least one circuit 1052 can comprise one or more switches 1172 coupled to the controller 1062 (e.g., logic element 1070) and to the volatile memory subsystem 1030 which can be actuated to couple and decouple the controller 1062 to and from the volatile memory subsystem 1030, respectively. The memory system 1010 further comprises one or more switches 1170 coupled to the one or more registers 1160 and to the plurality of

24

DRAM elements 1180 as schematically illustrated by FIG. 15A. The one or more switches 1170 can be selectively switched, thereby selectively operatively coupling the volatile memory subsystem 1030 to the host system 1150. In certain other embodiments, as schematically illustrated by FIG. 15B, the one or more switches 1174 are also coupled to the one or more registers 1160 and to a power source 1162 for the one or more registers 1160. The one or more switches 1174 can be selectively switched to turn power on or off to the one or more registers 1160, thereby selectively operatively coupling the volatile memory subsystem 1030 to the host system 1150. As schematically illustrated by FIG. 15C, in certain embodiments the at least one circuit 1052 comprises a dynamic on-die termination (ODT) 1176 circuit of the logic element 1070. For example, the logic element 1070 can comprise a dynamic ODT circuit 1176 which selectively operatively couples and decouples the logic element 1070 to and from the volatile memory subsystem 1030, respectively. In addition, and similar to the example embodiment of FIG. 15A described above, the one or more switches 1170 can be selectively switched, thereby selectively operatively coupling the volatile memory subsystem 1030 to the host system 1150.

Certain embodiments described herein utilize the non-volatile memory subsystem 1040 as a flash "mirror" to provide backup of the volatile memory subsystem 1030 in the event of certain system conditions. For example, the non-volatile memory subsystem 1040 may backup the volatile memory subsystem 1030 in the event of a trigger condition, such as, for example, a power failure or power reduction or a request from the host system. In one embodiment, the nonvolatile memory subsystem 1040 holds intermediate data results in a noisy system environment when the host computer system is engaged in a long computation. In certain embodiments, a backup may be performed on a regular basis. For example, in one embodiment, the backup may occur every millisecond in response to a trigger condition. In certain embodiments, the trigger condition occurs when the memory system 1010 detects that the system voltage is below a certain threshold voltage. For example, in one embodiment, the threshold voltage is 10 percent below a specified operating voltage. In certain embodiments, a trigger condition occurs when the voltage goes above a certain threshold value, such as, for example, 10 percent above a specified operating voltage. In some embodiments, a trigger condition occurs when the voltage goes below a threshold or above another threshold. In various embodiments, a backup and/or restore operation may occur in reboot and/or non-reboot trigger conditions.

As schematically illustrated by FIGS. 12 and 13, in certain embodiments, the controller 1062 may comprise a microcontroller unit (MCU) 1060 and a logic element 1070. In certain embodiments, the MCU 1060 provides memory management for the non-volatile memory subsystem 1040 and controls data transfer between the volatile memory subsystem 30 and the nonvolatile memory subsystem 1040. The MCU 1060 of certain embodiments comprises a 16-bit microcontroller, although other types of microcontrollers are also compatible with certain embodiments described herein. As schematically illustrated by FIGS. 12 and 13, the logic element 1070 of certain embodiments is in electrical communication with the non-volatile memory subsystem 1040 and the MCU 1060. The logic element 1070 can provide signal level translation between the volatile memory elements 1032 (e.g., 1.8V SSTL-2 for DDR2 SDRAM elements) and the non-volatile memory elements 1042 (e.g., 3V TTL for NAND flash memory elements). In certain embodi-

US 12,373,366 B2

25

ments, the logic element **1070** is also programmed to perform address/address translation between the volatile memory subsystem **1030** and the non-volatile memory subsystem **1040**. In certain preferred embodiments, 1-NAND type flash are used for the non-volatile memory elements **1042** because of their superior read speed and compact structure.

The memory system **1010** of certain embodiments is configured to be operated in at least two states. The at least two states can comprise a first state in which the controller **1062** and the non-volatile memory subsystem **1040** are operatively decoupled (e.g., isolated) from the volatile memory subsystem **1030** by the at least one circuit **1052** and a second state in which the volatile memory subsystem **1030** is operatively coupled to the controller **1062** to allow data to be communicated between the volatile memory subsystem **1030** and the nonvolatile memory subsystem **1040** via the controller **1062**. The memory system **1010** may transition from the first state to the second state in response to a trigger condition, such as when the memory system **1010** detects that there is a power interruption (e.g., power failure or reduction) or a system hang-up.

The memory system **1010** may further comprise a voltage monitor **1050**. The voltage monitor circuit **1050** monitors the voltage supplied by the host system via the interface **1022**. Upon detecting a low voltage condition (e.g., due to a power interruption to the host system), the voltage monitor circuit **1050** may transmit a signal to the controller **1062** indicative of the detected condition. The controller **1062** of certain embodiments responds to the signal from the voltage monitor circuit **1050** by transmitting a signal to the at least one circuit **1052** to operatively couple the controller to the volatile memory system **1030**, such that the memory system **1010** enters the second state. For example, the voltage monitor **1050** may send a signal to the MCU **1060** which responds by accessing the data on the volatile memory system **1030** and by executing a write cycle on the nonvolatile memory subsystem **1040**. During this write cycle, data is read from the volatile memory subsystem **1030** and is transferred to the non-volatile memory subsystem **1040** via the MCU **1060**. In certain embodiments, the voltage monitor circuit **1050** is part of the controller **1062** (e.g., part of the MCU **1060**) and the voltage monitor circuit **1050** transmits a signal to the other portions of the controller **1062** upon detecting a power threshold condition.

The isolation or operational decoupling of the volatile memory subsystem **1030** from the non-volatile memory subsystem in the first state can preserve the integrity of the operation of the memory system **1010** during periods of operation in which signals (e.g., data) are transmitted between the host system and the volatile memory subsystem **1030**. For example, in one embodiment during such periods of operation, the controller **1062** and the nonvolatile memory subsystem **1040** do not add a significant capacitive load to the volatile memory system **1030** when the memory system **1010** is in the first state. In certain such embodiments, the capacitive load of the controller **1062** and the non-volatile memory subsystem **1040** do not significantly affect the signals propagating between the volatile memory subsystem **1030** and the host system. This can be particularly advantageous in relatively high-speed memory systems where loading effects can be significant. In one preferred embodiment, the at least one circuit **1052** comprises an FSA1208 Low-Power, Eight-Port, Hi-Speed Isolation Switch from Fairchild Semiconductor. In other embodiments, the at least one circuit **1052** comprises other types of isolation devices.

26

Power may be supplied to the volatile memory subsystem **1030** from a first power supply (e.g., a system power supply) when the memory system **1010** is in the first state and from a second power supply **1080** when the memory system **1010** is in the second state. In certain embodiments, the memory system **1010** is in the first state when no trigger condition (e.g., a power failure) is present and the memory system **1010** enters the second state in response to a trigger condition. In certain embodiments, the memory system **1010** has a third state in which the controller **1062** is operatively decoupled from the volatile memory subsystem **1030** and power is supplied to the volatile memory subsystem **1030** from a third power supply (not shown). For example, in one embodiment the third power supply may provide power to the volatile memory subsystem **1030** when the memory system **1010** detects that a trigger condition is likely to occur but has not yet occurred.

In certain embodiments, the second power supply **1080** does not comprise a battery. Because a battery is not used, the second power supply **1080** of certain embodiments may be relatively easy to maintain, does not generally need to be replaced, and is relatively environmentally friendly. In certain embodiments, as schematically illustrated by FIGS. **12-14**, the second power supply **1080** comprises a step-up transformer **1082**, a step-down transformer **1084**, and a capacitor bank **1086** comprising one or more capacitors (e.g., double-layer capacitors). In one example embodiment, capacitors may take about three to four minutes to charge and about two minutes to discharge. In other embodiments, the one or more capacitors may take a longer time or a shorter time to charge and/or discharge. For example, in certain embodiments, the second power supply **1080** is configured to power the volatile memory subsystem **1030** for less than thirty minutes. In certain embodiments, the second power supply **1080** may comprise a battery. For example, in certain embodiments, the second power supply **1080** comprises a battery and one or more capacitors and is configured to power the volatile memory subsystem **1030** for no more than thirty minutes.

In certain embodiments, the capacitor bank **1086** of the second power supply **1080** is charged by the first power supply while the memory system **1010** is in the first state. As a result, the second power supply **1080** is fully charged when the memory system **1010** enters the second state. The memory system **1010** and the second power supply **1080** may be located on the same printed circuit board **1020**. In other embodiments, the second power supply **1080** may not be on the same printed circuit board **1020** and may be tethered to the printed circuit board **1020**, for example.

When operating in the first state, in certain embodiments, the step-up transformer **1082** keeps the capacitor bank **1086** charged at a peak value. In certain embodiments, the step-down transformer **1084** acts as a voltage regulator to ensure that regulated voltages are supplied to the memory elements (e.g., 1.8V to the volatile DRAM elements **1032** and 3.0V to the non-volatile flash memory elements **1042**) when operating in the second state (e.g., during power down). In certain embodiments, as schematically illustrated by FIGS. **12-14**, the memory module **1010** further comprises a switch **1090** (e.g., FET switch) that switches power provided to the controller **1062**, the volatile memory subsystem **1030**, and the non-volatile memory subsystem **1040**, between the power from the second power supply **1080** and the power from the first power supply (e.g., system power) received via the interface **1022**. For example, the switch **1090** may switch from the first power supply to the second power supply **1080** when the voltage monitor **1050** detects a low voltage con-

US 12,373,366 B2

27

dition. The switch **1090** of certain embodiments advantageously ensures that the volatile memory elements **1032** and non-volatile memory elements **1042** are powered long enough for the data to be transferred from the volatile memory elements **1032** and stored in the non-volatile memory elements **1042**. In certain embodiments, after the data transfer is complete, the switch **1090** then switches back to the first power supply and the controller **1062** transmits a signal to the at least one circuit **1052** to operatively decouple the controller **1062** from the volatile memory subsystem **1030**, such that the memory system **1010** reenters the first state.

When the memory system **1010** re-enters the first state, data may be transferred back from the non-volatile memory subsystem **1040** to the volatile memory subsystem **1030** via the controller **1062**. The host system can then resume accessing the volatile memory subsystem **1030** of the memory module **1010**. In certain embodiments, after the memory system **1010** enters or re-enters the first state (e.g., after power is restored), the host system accesses the volatile memory subsystem **1030** rather than the non-volatile memory subsystem **1040** because the volatile memory elements **1032** have superior read/write characteristics. In certain embodiments, the transfer of data from the volatile memory bank **1030** to the nonvolatile memory bank **1040**, or from the non-volatile memory bank **1040** to the volatile memory bank **1030**, takes less than one minute per GB.

In certain embodiments, the memory system **1010** protects the operation of the volatile memory when communicating with the host-system and provides backup and restore capability in the event of a trigger condition such as a power failure. In certain embodiments, the memory system **1010** copies the entire contents of the volatile memory subsystem **1030** into the nonvolatile memory subsystem **1040** on each backup operation. Moreover, in certain embodiments, the entire contents of the non-volatile memory subsystem **1040** are copied back into the volatile memory subsystem **1030** on each restore operation. In certain embodiments, the entire contents of the non-volatile memory subsystem **1040** are accessed for each backup and/or restore operation, such that the non-volatile memory subsystem **1040** (e.g., flash memory subsystem) is used generally uniformly across its memory space and wear-leveling is not performed by the memory system **1010**. In certain embodiments, avoiding wear-leveling can decrease cost and complexity of the memory system **1010** and can improve the performance of the memory system **1010**. In certain other embodiments, the entire contents of the volatile memory subsystem **1030** are not copied into the non-volatile memory subsystem **1040** on each backup operation, but only a partial copy is performed. In certain embodiments, other management capabilities such as bad-block management and error management for the flash memory elements of the non-volatile memory subsystem **1040** are performed in the controller **1062**.

The memory system **1010** generally operates as a write-back cache in certain embodiments. For example, in one embodiment, the host system (e.g., a disk controller) writes data to the volatile memory subsystem **1030** which then writes the data to non-volatile storage which is not part of the memory system **1010**, such as, for example, a hard disk. The disk controller may wait for an acknowledgment signal from the memory system **1010** indicating that the data has been written to the hard disk or is otherwise secure. The memory system **1010** of certain embodiments can decrease delays in the system operation by indicating that the data has been written to the hard disk before it has actually done so. In certain embodiments, the memory system **1010** will still

28

be able to recover the data efficiently in the event of a power outage because of the backup and restore capabilities described herein. In certain other embodiments, the memory system **1010** may be operated as a write-through cache or as some other type of cache.

FIG. **16** schematically illustrates an example power module **1100** of the memory system **1010** in accordance with certain embodiments described herein. The power module **1100** provides power to the various components of the memory system **1010** using different elements based on a state of the memory system **1010** in relation to a trigger condition. In certain embodiments, the power module **1100** comprises one or more of the components described above with respect to FIG. **12**. For example, in certain embodiments, the power module **1100** includes the second power supply **1080** and the switch **1090**.

The power module **1100** provides a plurality of voltages to the memory system **1010** comprising non-volatile and volatile memory subsystems **1030**, **1040**. The plurality of voltages comprises at least a first voltage **1102** and a second voltage **1104**. The power module **1100** comprises an input **1106** providing a third voltage **1108** to the power module **1100** and a voltage conversion element **1120** configured to provide the second voltage **1104** to the memory system **1010**. The power module **1100** further comprises a first power element **1130** configured to selectively provide a fourth voltage **1110** to the conversion element **1120**. In certain embodiments, the first power element **1130** comprises a pulse-width modulation power controller. For example, in one example embodiment, the first power element **1130** is configured to receive a 1.8V input system voltage as the third voltage **1108** and to output a modulated 5V output as the fourth voltage **1110**.

The power module **1100** further comprises a second power element **1140** can be configured to selectively provide a fifth voltage **1112** to the conversion element **1120**. The power module **1100** can be configured to selectively provide the first voltage **1102** to the memory system **1010** either from the conversion element **1120** or from the input **1106**.

The power module **1100** can be configured to be operated in at least three states in certain embodiments. In a first state, the first voltage **1102** is provided to the memory system **1010** from the input **1106** and the fourth voltage **1110** is provided to the conversion element **1120** from the first power element **1130**. In a second state, the fourth voltage **1110** is provided to the conversion element **1120** from the first power element **1130** and the first voltage **1102** is provided to the memory system **1010** from the conversion element **1120**. In the third state, the fifth voltage **1112** is provided to the conversion element **1120** from the second power element **1140** and the first voltage **1104** is provided to the memory system **1010** from the conversion element **1120**.

In certain embodiments, the power module **1100** transitions from the first state to the second state upon detecting that a trigger condition is likely to occur and transitions from the second state to the third state upon detecting that the trigger condition has occurred. For example, the power module **1100** may transition to the second state when it detects that a power failure is about to occur and transitions to the third state when it detects that the power failure has occurred. In certain embodiments, providing the first voltage **1102** in the second state from the first power element **1130** rather than from the input **1106** allows a smoother transition from the first state to the third state. For example, in certain embodiments, providing the first voltage **1102** from the first power element **1130** has capacitive and other smoothing effects. In addition, switching the point of power transition

US 12,373,366 B2

29

to be between the conversion element **1120** and the first and second power elements **1130**, **1140** (e.g., the sources of the pre-regulated fourth voltage **1110** in the second state and the pre-regulated fifth voltage **1112** in the third state) can smooth out potential voltage spikes.

In certain embodiments, the second power element **1140** does not comprise a battery and may comprise one or more capacitors. For example, as schematically illustrated in FIG. **16**, the second power element **1140** comprises a capacitor array **1142**, a buck-boost converter **1144** which adjusts the voltage for charging the capacitor array and a voltage/current limiter **1146** which limits the charge current to the capacitor array **1142** and stops charging the capacitor array **1142** when it has reached a certain charge voltage. In one example embodiment, the capacitor array **1142** comprises two 50 farad capacitors capable of holding a total charge of 4.6V. For example, in one example embodiment, the buck-boost converter **1144** receives a 1.8V system voltage (first voltage **1108**) and boosts the voltage to 4.3V which is outputted to the voltage current limiter **1146**. The voltage/current limiter **1146** limits the current going to the capacitor array **1142** to 1 A and stops charging the array **1142** when it is charged to 4.3V. Although described with respect to certain example embodiments, one of ordinary skill will recognize from the disclosure herein that the second power element **1140** may include alternative embodiments. For example, different components and/or different value components may be used. For example, in other embodiments, a pure boost converter may be used instead of a buck-boost converter. In another embodiment, only one capacitor may be used instead of a capacitor array **1142**.

The conversion element **1120** can comprise one or more buck converters and/or one or more buck-boost converters. The conversion element **1120** may comprise a plurality of sub-blocks **1122**, **1124**, **1126** as schematically illustrated by FIG. **16**, which can provide more voltages in addition to the second voltage **1104** to the memory system **1010**. The sub-blocks may comprise various converter circuits such as buck-converters, boost converters, and buck-boost converter circuits for providing various voltage values to the memory system **1010**. For example, in one embodiment, sub-block **1122** comprises a buck converter, sub-block **1124** comprises a dual buck converter, and sub-block **1126** comprises a buck-boost converter as schematically illustrated by FIG. **16**. Various other components for the sub-blocks **1122**, **1124**, **1126** of the conversion element **1120** are also compatible with certain embodiments described herein. In certain embodiments, the conversion element **1120** receives as input either the fourth voltage **1110** from the first power element **1130** or the fifth voltage **1112** from the second power element **1140**, depending on the state of the power module **1100**, and reduces the input to an appropriate amount for powering various components of the memory system. For example, the buck-converter of sub-block **1122** can provide 1.8V at 2 A for about 60 seconds to the volatile memory elements **1032** (e.g., DRAM), the non-volatile memory elements **1042** (e.g., flash), and the controller **1062** (e.g., an FPGA) in one embodiment. The sub-block **1124** can provide the second voltage **1104** as well as another reduced voltage **1105** to the memory system **1010**. In one example embodiment, the second voltage **1104** is 2.5V and is used to power the at least one circuit **1052** (e.g., isolation device) and the other reduced voltage **1105** is 1.2V and is used to power the controller **1062** (e.g., FPGA). The subblock **1126** can provide yet another voltage **1107** to the memory system **1010**.

30

For example, the voltage **1107** may be 3.3V and may be used to power both the controller **1062** and the at least one circuit **1052**.

Although described with respect to certain example embodiments, one of ordinary skill will recognize from the disclosure herein that the conversion element **1120** may include alternative embodiments. For example, there may be more or less sub-blocks which may comprise other types of converters (e.g., pure boost converters) or which may produce different voltage values. In one embodiment, the volatile memory elements **1032** and nonvolatile memory elements **1042** are powered using independent voltages and are not both powered using the first voltage **1102**.

FIG. **17** is a flowchart of an example method **1200** of providing a first voltage **1102** and a second voltage **1104** to a memory system **1010** including volatile and nonvolatile memory subsystems **1030**, **1040**. While the method **1200** is described herein by reference to the memory system **1010** schematically illustrated by FIGS. **12-15**, other memory systems are also compatible with embodiments of the method **1200**. During a first condition, the method **1200** comprises providing the first voltage **1102** to the memory system **1010** from an input power supply **1106** and providing the second voltage **1104** to the memory system **1010** from a first power subsystem in operational block **1210**. For example, in one embodiment, the first power subsystem comprises the first power element **1130** and the voltage conversion element **1120** described above with respect to FIG. **16**. In other embodiments, other first power subsystems are used.

The method **1200** further comprises detecting a second condition in operational block **1220**. In certain embodiments, detecting the second condition comprises detecting that a trigger condition is likely to occur. During the second condition, the method **1200** comprises providing the first voltage **1102** and the second voltage **1104** to the memory system **1010** from the first power subsystem in an operational block **1230**. For example, referring to FIG. **16**, a switch **1148** can be toggled to provide the first voltage **1102** from the conversion element **1120** rather than from the input power supply.

The method **1200** further comprises charging a second power subsystem in operational block **1240**. In certain embodiments, the second power subsystem comprises the second power element **1140** or another power supply that does not comprise a battery. For example, in one embodiment, the second power subsystem comprises the second power element **1140** and the voltage conversion element **1120** described above with respect to FIG. **16**. In other embodiments, some other second power subsystem is used.

The method **1200** further comprises detecting a third condition in an operational block **1250** and during the third condition, providing the first voltage **1102** and the second voltage **1104** to the memory system **1010** from the second power subsystem **1140** in an operational block **1260**. In certain embodiments, detecting the third condition comprises detecting that the trigger condition has occurred. The trigger condition may comprise various conditions described herein. In various embodiments, for example, the trigger condition comprises a power reduction, power failure, or system hang-up. The operational blocks of the method **1200** may be performed in different orders in various embodiments. For example, in certain embodiments, the second power subsystem **1140** is charged before detecting the second condition.

In certain embodiments, the memory system **1010** comprises a volatile memory subsystem **1030** and a non-volatile

US 12,373,366 B2

31

memory subsystem **1040** comprising at least 100 percent more storage capacity than does the volatile memory subsystem. The memory system **1010** also comprises a controller **1062** operatively coupled to the volatile memory subsystem **1030** and operatively coupled to the non-volatile memory subsystem **1040**. The controller **1062** can be configured to allow data to be communicated between the volatile memory subsystem **1030** and the host system when the memory system **1010** is operating in a first state and to allow data to be communicated between the volatile memory subsystem **1030** and the non-volatile memory subsystem **1040** when the memory system **1010** is operating in a second state.

Although the memory system **1010** having extra storage capacity of the non-volatile memory subsystem **1040** has been described with respect to certain embodiments, alternative configurations exist. For example, in certain embodiments, there may be more than 100 percent more storage capacity in the non-volatile memory subsystem **1040** than in the volatile memory subsystem **1030**. In various embodiments, there may be at least 200, 300, or 400 percent more storage capacity in the non-volatile memory subsystem **1040** than in the volatile memory subsystem **1030**. In other embodiments, the non-volatile memory subsystem **1040** includes at least some other integer multiples of the storage capacity of the volatile memory subsystem **1030**. In some embodiments, the non-volatile memory subsystem **1040** includes a non-integer multiple of the storage capacity of the volatile memory subsystem **1030**. In one embodiment, the non-volatile memory subsystem **1040** includes less than 100 percent more storage capacity than does the volatile memory subsystem **1030**.

The extra storage capacity of the non-volatile memory subsystem **1040** can be used to improve the backup capability of the memory system **1010**. In certain embodiments in which data can only be written to portions of the non-volatile memory subsystem **1040** which do not contain data (e.g., portions which have been erased), the extra storage capacity of the nonvolatile memory subsystem **1040** allows the volatile memory subsystem **1030** to be backed up in the event of a subsequent power failure or other trigger event. For example, the extra storage capacity of the non-volatile memory subsystem **1040** may allow the memory system **1010** to backup the volatile memory subsystem **1030** efficiently in the event of multiple trigger conditions (e.g., power failures). In the event of a first power failure, for example, the data in the volatile memory system **1030** is copied to a first, previously erased portion of the nonvolatile memory subsystem **1040** via the controller **1062**. Since the non-volatile memory subsystem **1040** has more storage capacity than does the volatile memory subsystem **1030**, there is a second portion of the non-volatile memory subsystem **1040** which does not have data from the volatile memory subsystem **1030** copied to it and which remains free of data (e.g., erased). Once system power is restored, the controller **1062** of the memory system **1010** restores the data to the volatile memory subsystem **1030** by copying the backed-up data from the non-volatile memory subsystem **1040** back to the volatile memory subsystem **1030**. After the data is restored, the memory system **1010** erases the non-volatile memory subsystem **1040**. While the first portion of the non-volatile memory subsystem **1040** is being erased, it may be temporarily inaccessible.

If a subsequent power failure occurs before the first portion of the non-volatile memory subsystem **1040** is completely erased, the volatile memory subsystem **1030** can be backed-up or stored again in the second portion of the

32

non-volatile memory subsystem **1040** as described herein. In certain embodiments, the extra storage capacity of the non-volatile memory subsystem **1040** may allow the memory system **1010** to operate more efficiently. For example, because of the extra storage capacity of the non-volatile memory subsystem **1040**, the memory system **1010** can handle a higher frequency of trigger events that is not limited by the erase time of the non-volatile memory subsystem **1040**.

FIG. **18** is a flowchart of an example method **1300** of controlling a memory system **1010** operatively coupled to a host system and which includes a volatile memory subsystem **1030** and a non-volatile memory subsystem **1040**. In certain embodiments, the non-volatile memory subsystem **1040** comprises at least 100 percent more storage capacity than does the volatile memory subsystem **1030** as described herein. While the method **1300** is described herein by reference to the memory system **1010** schematically illustrated by FIGS. **12-14**, the method **1300** can be practiced using other memory systems in accordance with certain embodiments described herein. In an operational block **1310**, the method **1300** comprises communicating data between the volatile memory subsystem **1030** and the host system when the memory system **1010** is in a first mode of operation. The method **1300** further comprises storing a first copy of data from the volatile memory subsystem **1030** to the non-volatile memory subsystem **1040** at a first time when the memory system **1010** is in a second mode of operation in an operational block **1320**.

In an operational block **1330**, the method **1300** comprises restoring the first copy of data from the non-volatile memory subsystem **1040** to the volatile memory subsystem **1030**. The method **1300** further comprises erasing the first copy of data from the non-volatile memory subsystem **1040** in an operational block **1340**. The method further comprises storing a second copy of data from the volatile memory subsystem **1030** to the non-volatile memory subsystem **1040** at a second time when the memory system **1010** is in the second mode of operation in an operational block **1350**. Storing the second copy begins before the first copy is completely erased from the non-volatile memory subsystem **1040**.

In some embodiments, the memory system **1010** enters the second mode of operation in response to a trigger condition, such as a power failure. In certain embodiments, the first copy of data and the second copy of data are stored in separate portions of the nonvolatile memory subsystem **1040**. The method **1300** can also include restoring the second copy of data from the non-volatile memory subsystem **1040** to the volatile memory subsystem **1030** in an operational block **1360**. The operational blocks of method **1300** referred to herein may be performed in different orders in various embodiments. For example, in some embodiments, the second copy of data is restored to the volatile memory subsystem **1030** at operational block **1360** before the first copy of data is completely erased in the operational block **1340**.

FIG. **19** schematically illustrates an example clock distribution topology **1400** of a memory system **1010** in accordance with certain embodiments described herein. The clock distribution topology **1400** generally illustrates the creation and routing of the clock signals provided to the various components of the memory system **1010**. A clock source **1402** such as, for example, a 25 MHz oscillator, generates a clock signal. The clock source **1402** may feed a clock generator **1404** which provides a clock signal **1406** to the controller **1062**, which may be an FPGA. In one embodi-

US 12,373,366 B2

33

ment, the clock generator **1404** generates a 125 MHz clock signal **1406**. The controller **1062** receives the clock signal **1406** and uses it to clock the controller **1062** master state control logic. For example, the master state control logic may control the general operation of an FPGA controller **1062**.

The clock signal **1406** can also be input into a clock divider **1410** which produces a frequency-divided version of the clock signal **1406**. In an example embodiment, the clock divider **1410** is a divide by two clock divider and produces a 62.5 MHz clock signal in response to the 125 MHz clock signal **1406**. A non-volatile memory phase-locked loop (PLL) block **1412** can be included (e.g., in the controller **1062**) which distributes a series of clock signals to the non-volatile memory subsystem **1040** and to associated control logic. For example, a series of clock signals **1414** can be sent from the controller **1062** to the non-volatile memory subsystem **1040**. Another clock signal **1416** can be used by the controller logic which is dedicated to controlling the non-volatile memory subsystem **1040**. For example, the clock signal **1416** may clock the portion of the controller **1062** which is dedicated to generating address and/or control lines for the non-volatile memory subsystem **1040**. A feedback clock signal **1418** is fed back into the non-volatile memory PLL block **1412**. In one embodiment, the PLL block **1412** compares the feedback clock **1418** to the reference clock **1411** and varies the phase and frequency of its output until the reference **1411** and feedback **1418** clocks are phase and frequency matched.

A version of the clock signal **1406** such as the backup clock signal **1408** may be sent from the controller to the volatile memory subsystem **1030**. The clock signal **1408** may be, for example, a differential version of the clock signal **1406**. As described herein, the backup clock signal **1408** may be used to clock the volatile memory subsystem **1030** when the memory system **1010** is backing up the data from the volatile memory subsystem **1030** into the non-volatile memory subsystem **1040**. In certain embodiments, the backup clock signal **1408** may also be used to clock the volatile memory subsystem **1030** when the memory system **1010** is copying the backed-up data back into the volatile memory subsystem **1030** from the nonvolatile memory subsystem **1040** (also referred to as restoring the volatile memory subsystem **1030**). The volatile memory subsystem **1030** may normally be run at a higher frequency (e.g., DRAM running at 400 MHz) than the nonvolatile memory subsystem **1040** (e.g., flash memory running at 62.5 MHz) when communicating with the host system (e.g., when no trigger condition is present). However, in certain embodiments the volatile memory subsystem **1030** may be operated at a reduced frequency (e.g., at twice the frequency of the non-volatile memory subsystem **1040**) without introducing significant delay into the system during backup operation and/or restore operations. Running the volatile memory subsystem **1030** at the reduced frequency during a backup and/or restore operation may advantageously reduce overall power consumption of the memory system **1010**.

In one embodiment, the backup clock **1408** and the volatile memory system clock signal **1420** are received by a multiplexer **1422**, as schematically illustrated by FIG. **19**. The multiplexer **1422** can output either the volatile memory system clock signal **1420** or the backup clock signal **1408** depending on the backup state of the memory system **1010**. For example, when the memory system **1010** is not performing a backup or restore operation and is communicating with the host system (e.g., normal operation), the volatile memory system clock signal **1420** may be provided by the

34

multiplexer **422** to the volatile memory PLL block **1424**. When the memory system **1010** is performing a backup (or restore) operation, the backup clock signal **1408** may be provided.

The volatile memory PLL block **1424** receives the volatile memory reference clock signal **1423** from the multiplexer **1422** and can generate a series of clock signals which are distributed to the volatile memory subsystem **1030** and associated control logic. For example, in one embodiment, the PLL block **1424** generates a series of clock signals **1426** which clock the volatile memory elements **1032**. A clock signal **1428** may be used to clock control logic associated with the volatile memory elements, such as one or more registers (e.g., the one or more registers of a registered DIMM). Another clock signal **1430** may be sent to the controller **1062**. A feedback clock signal **1432** is fed back into the volatile memory PLL block **1424**. In one embodiment, the PLL block **1424** compares the feedback clock signal **1432** to the reference clock signal **1423** and varies the phase and frequency of its output until the reference clock signal **1423** and the feedback clock signal **1432** clocks are phase and frequency matched.

The clock signal **1430** may be used by the controller **1062** to generate and distribute clock signals which will be used by controller logic which is configured to control the volatile memory subsystem **1030**. For example, control logic in the controller **1062** may be used to control the volatile memory subsystem **1030** during a backup or restore operation. The clock signal **1430** may be used as a reference clock signal for the PLL block **1434** which can generate one or more clocks **1438** used by logic in the controller **1062**. For example, the PLL block **1434** may generate one or more clock signals **1438** used to drive logic circuitry associated with controlling the volatile memory subsystem **1030**. In certain embodiments, the PLL block **1434** includes a feedback clock signal **1436** and operates in a similar manner to other PLL blocks described herein.

The clock signal **1430** may be used as a reference clock signal for the PLL block **1440** which may generate one or more clock signals used by a sub-block **1442** to generate one or more other clock signals **1444**. In one embodiment, for example, the volatile memory subsystem **1030** comprises DDR2 SDRAM elements and the sub-block **1442** generates one or more DDR2 compatible clock signals **1444**. A feedback clock signal **1446** is fed back into the PLL block **1440**. In certain embodiments, the PLL block **1440** operates in a similar manner to other PLL blocks described herein.

While described with respect to the example embodiment of FIG. **19**, various alternative clock distribution topologies are possible. For example, one or more of the clock signals have a different frequency in various other embodiments. In some embodiments, one or more of the clocks shown as differential signals are single ended signals. In one embodiment, the volatile memory subsystem **1030** operates on the volatile memory clock signal **1420** and there is no backup clock signal **1408**. In some embodiments, the volatile memory subsystem **1030** is operated at a reduced frequency during a backup operation and not during a restore operation. In other embodiments, the volatile memory subsystem **1030** is operated at a reduced frequency during a restore operation and not during a backup operation.

FIG. **20** is a flowchart of an example method **1500** of controlling a memory system **1010** operatively coupled to a host system. Although described with respect to the memory system **1010** described herein, the method **1500** is compatible with other memory systems. The memory system **1010** may include a clock distribution topology **1400** similar to

US 12,373,366 B2

35

the one described above with respect to FIG. 19 or another clock distribution topology. The memory system 1010 can include a volatile memory subsystem 30 and a non-volatile memory subsystem 1040.

In an operational block 1510, the method 1500 comprises operating the volatile memory subsystem 1030 at a first frequency when the memory system 1010 is in a first mode of operation in which data is communicated between the volatile memory subsystem 1030 and the host system. In an operational block 1520, the method 1500 comprises operating the non-volatile memory subsystem 1040 at a second frequency when the memory system 1010 is in a second mode of operation in which data is communicated between the volatile memory subsystem 1030 and the non-volatile memory subsystem 1040. The method 1500 further comprises operating the volatile memory subsystem 1030 at a third frequency in an operational block 1530 when the memory system 1010 is in the second mode of operation. In certain embodiments, the memory system 1010 is not powered by a battery when it is in the second mode of operation. The memory system 1010 may switch from the first mode of operation to the second mode of operation in response to a trigger condition. The trigger condition may be any trigger condition described herein such as, for example, a power failure condition. In certain embodiments, the second mode of operation includes both backup and restore operations as described herein. In other embodiments, the second mode of operation includes backup operations but not restore operations. In yet other embodiments, the second mode of operation includes restore operations but not backup operations.

The third frequency can be less than the first frequency. For example, the third frequency can be approximately equal to the second frequency. In certain embodiments, the reduced frequency operation is an optional mode. In yet other embodiments, the first, second and/or third frequencies are configurable by a user or by the memory system 1010.

FIG. 21 schematically illustrates an example topology of a connection to transfer data slices from two DRAM segments 1630, 1640 of a volatile memory subsystem 1030 of a memory system 1010 to a controller 1062 of the memory system 1010. While the example of FIG. 21 shows a topology including two DRAM segments 1630, 1640 for the purposes of illustration, each address location of the volatile memory subsystem 1030 comprises more than the two segments in certain embodiments. The data lines 1632, 1642 from the first DRAM segment 1630 and the second DRAM segment 1640 of the volatile memory subsystem 1030 are coupled to switches 1650, 1652 which are coupled to the controller 1062 (e.g., logic element 1070) of the memory system 1010. The chip select lines 1634, 1644 and the self-refresh lines 1636, 1646 (e.g., CKe signals) of the first and second DRAM segments 1630, 1640, respectively, are coupled to the controller 1062. In certain embodiments, the controller 1062 comprises a buffer (not shown) which is configured to store data from the volatile memory subsystem 1030. In certain embodiments, the buffer is a first-in, first out buffer (FIFO). In certain embodiments, data slices from each DRAM segment 1630, 1640 comprise a portion of the volatile memory subsystem data bus. In one embodiment, for example, the volatile memory subsystem 1030 comprises a 72-bit data bus (e.g., each data word at each addressable location is 72 bits wide and includes, for example, 64 bits of accessible SDRAM and 8 bits of ECC), the first data slice from the first DRAM segment 1630 may comprise 40 bits of the data word, and the second data slice from the second DRAM segment 1640 may comprise the remaining 32 bits

36

of the data word. Certain other embodiments comprise data buses and/or data slices of different sizes.

In certain embodiments, the switches 1650, 1652 can each be selectively switched to selectively operatively couple the data lines 1632, 1642, respectively from the first and second DRAM segments 1630, 1640 to the controller 1062. The chip select lines 1634, 1644 enable the first and second DRAM segments 1630, 1640, respectively, of the volatile memory subsystem 1030, and the self-refresh lines 1636, 1646 toggle the first and second DRAM segments 1630, 1640, respectively, from self-refresh mode to active mode. In certain embodiments, the first and second DRAM segments 1630, 1640 maintain stored information but are not accessible when they are in self-refresh mode, and maintain stored information and are accessible when they are in active mode.

In certain embodiments, when the memory system 1010 is backing up the volatile memory system 1030, data slices from only one of the two DRAM segments 1630, 1640 at a time are sent to the controller 1062. For example, when the first slice is being written to the controller 1062 during a back-up, the controller 1062 sends a signal via the CKe line 1636 to the first DRAM segment 1630 to put the first DRAM segment 1630 in active mode. In certain embodiments, the data slice from the first DRAM segment 1630 for multiple words (e.g., a block of words) is written to the controller 1062 before writing the second data slice from the second DRAM segment 1640 to the controller 1062. While the first data slice is being written to the controller 1062, the controller 1062 also sends a signal via the CKe line 1646 to put the second DRAM segment 1640 in self-refresh mode. Once the first data slice for one word or for a block of words is written to the controller 1062, the controller 1062 puts the first DRAM segment 1630 into self-refresh mode by sending a signal via the CKe line 1636 to the first DRAM segment 1640. The controller 1062 also puts the second DRAM segment 1640 into active mode by sending a signal via the CKe line 1646 to the DRAM segment 1640. The second slice for a word or for a block of words is written to the controller 1062. In certain embodiments, when the first and second data slices are written to the buffer in the controller 1062, the controller 1062 combines the first and second data slices 1630, 1640 into complete words or blocks of words and then writes each complete word or block of words to the non-volatile memory subsystem 1040. In certain embodiments, this process is called "slicing" the volatile memory subsystem 1030.

In certain embodiments, the data may be sliced in a restore operation as well as, or instead of, during a backup operation. For example, in one embodiment, the nonvolatile memory elements 1042 write each backed-up data word to the controller 1062 which writes a first slice of the data word to the volatile memory subsystem 1030 and then a second slice of the data word to the volatile memory subsystem 1030. In certain embodiments, slicing the volatile memory subsystem 1030 during a restore operation may be performed in a manner generally inverse to slicing the volatile memory subsystem 1030 during a backup operation.

FIG. 22 is a flowchart of an example method 1600 of controlling a memory system 1010 operatively coupled to a host system and which includes a volatile memory subsystem 1030 and a non-volatile memory subsystem 1040. Although described with respect to the memory system 1010 described herein with respect to FIGS. 12-14 and 21, the method 1600 is compatible with other memory systems. The method 1600 comprises communicating data words between the volatile memory subsystem 1030 and the host system

US 12,373,366 B2

37

when the memory system **1010** is in a first mode of operation in an operational block **1610**. For example, the memory system **1010** may be in the first mode of operation when no trigger condition has occurred and the memory system is not performing a backup and/or restore operation or is not being powered by a secondary power supply.

In an operational block **1620**, the method further comprises transferring data words from the volatile memory subsystem **1030** to the non-volatile memory subsystem **1040** when the memory system **1010** is in a second mode of operation. In certain embodiments, each data word comprises the data stored in a particular address of the memory system **1010**. The memory system **1010** may enter the second mode of operation, for example, when a trigger condition (e.g., a power failure) occurs. In certain embodiments, transferring each data word comprises storing a first portion (also referred to as a slice) of the data word in a buffer in an operational block **1622**, storing a second portion of the data word in the buffer in an operational block **1624**, and writing the entire data word from the buffer to the non-volatile memory subsystem **1040** in an operational block **1626**.

In one example embodiment, the data word may be a 72 bit data word (e.g., 64 bits of accessible SDRAM and 8 bits of ECC), the first portion (or "slice") may comprise 40 bits of the data word, and the second portion (or "slice") may comprise the remaining 32 bits of the data word. In certain embodiments, the buffer is included in the controller **1062**. For example, in one embodiment, the buffer is a first-in, first-out buffer implemented in the controller **1062** which comprises an FPGA. The method **1600** may generally be referred to as "slicing" the volatile memory during a backup operation. In the example embodiment, the process of "slicing" the volatile memory during a backup includes bringing the 32-bit slice out of self-refresh, reading a 32-bit block from the slice into the buffer, and putting the 32-bit slice back into self-refresh. The 40-bit slice is then brought out of self-refresh and a 40-bit block from the slice is read into a buffer. Each block may comprise a portion of multiple words. For example, each 32-bit block may comprise 32-bit portions of multiple 72-bit words. In other embodiments, each block comprises a portion of a single word. The 40-bit slice is then put back into self-refresh in the example embodiment. The 32-bit and 40-bit slices are then combined into a 72-bit block by the controller **1062** and ECC detection/correction is performed on each 72-bit word as it is read from the buffer and written into the non-volatile memory subsystem (e.g., flash).

In some embodiments, the entire data word may comprise more than two portions. For example, the entire data word may comprise three portions instead of two and transferring each data word further comprises storing a third portion of each data word in the buffer. In certain other embodiments, the data word may comprise more than three portions.

In certain embodiments, the data may be sliced in a restore operation as well as, or instead of, during a backup operation. For example, in one embodiment, the nonvolatile memory elements **1040** write each backed-up data word to the controller **1062** which writes a first portion of the data word to the volatile memory subsystem **1030** and then a second portion of the data word to the volatile memory subsystem **1030**. In certain embodiments, slicing the volatile memory subsystem **1030** during a restore operation may be performed in a manner generally inverse to slicing the volatile memory subsystem **1030** during a backup operation.

The method **1600** can advantageously provide significant power savings and can lead to other advantages. For

38

example, in one embodiment where the volatile memory subsystem **1030** comprises DRAM elements, only the slice of the DRAM which is currently being accessed (e.g., written to the buffer) during a backup is configured in full-operational mode. The slice or slices that are not being accessed may be put in self-refresh mode. Because DRAM in self-refresh mode uses significantly less power than DRAM in full-operational mode, the method **1600** can allow significant power savings. In certain embodiments, each slice of the DRAM includes a separate self-refresh enable (e.g., CKE) signal which allows each slice to be accessed independently.

In addition, the connection between the DRAM elements and the controller **1062** may be as large as the largest slice instead of as large as the data bus. In the example embodiment, the connection between the controller **1062** and the DRAM may be 40 bits instead of 72 bits. As a result, pins on the controller **1062** may be used for other purposes or a smaller controller may be used due to the relatively low number of pin-outs used to connect to the volatile memory subsystem **1030**. In certain other embodiments, the full width of the data bus is connected between the volatile memory subsystem **1030** and the controller **1062** but only a portion of it is used during slicing operations. For example, in some embodiments, memory slicing is an optional mode.

While embodiments and applications have been shown and described, it would be apparent to those skilled in the art having the benefit of this disclosure that many more modifications than mentioned above are possible without departing from the inventive concepts disclosed herein. The invention, therefore, is not to be restricted except in the spirit of the appended claims.

What is claimed is:

1. A dual in-line memory module (DIMM) configured to fit into a memory slot connector of a system board of a computer system, the memory slot connector including a set of power conduits, a set of data conduits, and a set of address and control conduits, wherein data signals are transferred between the system board and the DIMM via the set of data conduits, and address and control signals are delivered from the system board to the DIMM via the set of address and control conduits, the DIMM comprising:

a printed circuit board (PCB) having edge connections configured to fit into the memory slot connector of the system board, the PCB including:

a first set of edge connections configured to be electrically connected to the set of power conduits and operable to deliver to the DIMM power from power rails of the system board via the set of power conduits, wherein the power delivered to the DIMM from the power rails of the system board is the only power received by the DIMM;

a second set of edge connections configured to be electrically connected to the set of data conduits, and a third set of edge connections configured to be electrically connected to the set of address and control conduits,

a set of voltage supply lines electrically connected to the first set of edge connections and configured to receive the power delivered to the DIMM from the power rails of the system board by way of the memory slot connector of the system board and the first set of edge connections of the DIMM,

a set of data lines electrically connected to the second set of edge connections, and

a set of address and control lines electrically connected to the third set of edge connections;

US 12,373,366 B2

39

a controller including a voltage monitor circuit and non-volatile memory, the controller coupled to the PCB and to an input voltage supply line of the set of voltage supply lines, the voltage monitor circuit configured to:

- (i) monitor an input voltage supplied via the input voltage supply line;
- (ii) generate a trigger signal upon detecting a trigger condition;
- (iii) transmit the trigger signal to at least one other portion of the controller, wherein the controller is configured to perform, in response to the trigger signal, a write operation to write data into the nonvolatile memory; and wherein the trigger condition occurs when the input voltage exceeds a first threshold voltage;

a set of components coupled to the PCB, the set of components including a plurality of double data rate (DDR) synchronous dynamic random access memory (SDRAM) devices coupled to the set of address and control lines and to the set of data lines, the DDR SDRAM devices being operable to receive or output data signals via the set of data lines based on address and control signals received from the system board via the set of address and control lines; and

first, second, third, and fourth converter circuits coupled to the PCB and to the input voltage supply line, and configured to receive power from the input voltage supply line and to deliver power via first, second, third, and fourth regulated voltage lines,

wherein the set of components includes:

- a first component configured to receive power via at least the first regulated voltage line;
- a second component configured to receive power via at least the second regulated voltage line;
- a third component configured to receive power via at least the third regulated voltage line; and
- a fourth component configured to receive power via at least the fourth regulated voltage line;

wherein each component of the set of components:

- (i) is electrically connected to one or more of the first, second, third, and fourth regulated voltage lines; and
- (ii) receives power only via the one or more of the first, second, third, and fourth regulated voltage lines,

wherein the plurality of DDR SDRAM devices includes:

- a first group of at least five DDR SDRAM devices that are each connected to a first chip select line configured to electrically conduct a first chip select signal;
- a second group of at least four DDR SDRAM devices that are each connected to a second chip select line configured to electrically conduct a second chip select signal;

wherein when power from the power rails of the system board is provided to the DIMM and each group of the first and second groups of DDR SDRAM devices is not in self-refresh mode:

the first group of DDR SDRAM devices is configured to be enabled, in response to the first chip select signal, to receive or output at least 40 1-bit DDR data signals in parallel via at least 40 data lines of the set of data lines independently of whether the second group of DDR SDRAM devices is enabled, in response to the second chip select signal, to receive or output an integer number of 1-bit DDR data signals in parallel via an integer number of data lines of the set of data lines;

the at least 40 1-bit DDR data signals include at least 32 1-bit wide DDR data signals and at least eight 1-bit DDR error-correcting code (ECC) data signals; and

40

wherein the sum of the at least 40 data lines and the integer number of data lines equals a total number of data conduits of the set of data conduits of the memory slot connector, and the total number of data conduits is at least 72 data conduits.

2. The DIMM of claim 1, wherein the controller further includes:

- a data storage element configured to store a plurality of data bits and a logic element coupled to the nonvolatile memory and the data storage element and configured to electrically communicate with the nonvolatile memory and the data storage element, and wherein the write operation conducted by the controller in response to the trigger signal includes:
 - a first write operation to write a first set of data bits into a first portion of the data storage element;
 - a second write operation to write a second set of data bits into a second portion of the data storage element; and subsequently
 - a third write operation to write a third set of data bits from the data storage element into the nonvolatile memory, wherein the third set of data bits includes the first and second sets of data bits written during the first and second write operations.

3. The DIMM of claim 2, wherein, prior to performing the third write operation, the logic element is configured to generate the third set of data bits based on a combination of the first and second sets of data bits written during the first and second write operations.

4. The DIMM of claim 2, wherein, prior to performing the third write operation, the logic element is configured to generate the third set of data bits based on groups of data bits stored in the data storage element including the first and second sets of data bits written during the first and second write operations.

5. The DIMM of claim 2, wherein, prior to performing the third write operation, the logic element is configured to perform, in response to the detected trigger signal, a plurality of read operations to read the first and second sets of data bits from the first and second portions of the data storage element, respectively, and to generate the third set of data bits based on the read first and second sets of data bits.

6. The DIMM of claim 2, wherein, prior to performing the third write operation, the logic element is configured to perform, in response to the detected trigger signal, a read operation to read the first and second sets of data bits from the first and second portions of the data storage element, respectively, and to generate the third set of data bits based on the read first and second sets of data bits.

7. The DIMM of claim 2, wherein, prior to performing the first and second write operations, the logic element is configured to perform, in response to the detected trigger signal, one or more read operations.

8. The DIMM of claim 7, wherein the logic element is configured to perform a first read operation of the one or more read operations by reading the first and second sets of data bits from the first and second groups of DDR SDRAM devices, respectively.

9. The DIMM of claim 2, wherein the third write operation comprises a transfer of the third set of data bits from the data storage element to the nonvolatile memory via a first connection, and wherein the first and second write operations comprise a transfer of the first and second sets of data bits into the data storage element via a second connection.

US 12,373,366 B2

41

10. The DIMM of claim 2, wherein the data storage element is a buffer configured to store a plurality of data bits or a First-In First-Out (FIFO) buffer configured to store the plurality of data bits.

11. The DIMM of claim 1, wherein when power is provided to the DIMM from the power rails of the system board the voltage monitor circuit compares the input voltage to a plurality of threshold voltages including the first threshold voltage and a second threshold voltage.

12. The DIMM of claim 1, wherein the second component is further connected to the first regulated voltage line and is powered by the first and second regulated voltages when power from the power rails of the system board is provided to the DIMM.

13. The DIMM of claim 12, wherein, when power from the power rails of the system board is provided to the DIMM, an additional trigger condition occurs when the input voltage is above the first threshold and below the second threshold; and

wherein, in response to the additional trigger condition, the voltage monitor circuit is further configured to generate and transmit a second trigger signal.

14. The DIMM of claim 1, wherein, when power from the power rails of the system board is provided to the DIMM, the nonvolatile memory is not powered using the first regulated voltage.

15. The DIMM of claim 1, wherein each of the first, second, and third converter circuits is a buck converter, and the fourth converter circuit is a voltage reducing circuit.

16. A dual in-line memory module (DIMM) configured to fit into a memory slot connector of a system board of a computer system, the memory slot connector including a set of power conduits, a set of data conduits, and a set of address and control conduits, wherein data signals are transferred between the system board and the DIMM via the set of data conduits, and address and control signals are delivered from the system board to the DIMM via the set of address and control conduits, the DIMM comprising:

a printed circuit board (PCB) having edge connections configured to fit into the memory slot connector of the system board, the PCB including:

a first set of edge connections configured to be electrically connected to the set of power conduits and operable to deliver the DIMM power from the power rails of the system board via the set of power conduits, wherein the power delivered to the DIMM from the power rails of the system board is the only power received by the DIMM,

a second set of edge connections configured to be electrically connected to the set of data conduits,

a third set of edge connections configured to be electrically connected to the set of address and control conduits,

a set of voltage supply lines electrically connected to the first set of edge connections and configured to receive the power delivered to the DIMM from the power rails of the system board by way of the memory slot connector of the system board and the first set of edge connections of the DIMM,

a set of data lines electrically connected to the second set of edge connections, and

a set of address and control lines electrically connected to the third set of edge connections;

a controller including a voltage monitor circuit and a data storage element configured to store a plurality of data bits, the controller coupled to the PCB and to an input

42

voltage supply line of the set of voltage supply lines, the voltage monitor circuit configured to:

(i) monitor an input voltage supplied via the input voltage supply line;

(ii) generate a trigger signal upon detecting a trigger condition; and

(iii) transmit the trigger signal to at least one other portion of the controller,

wherein, in response to the detected trigger condition, the controller is configured to periodically perform, at regular intervals of time, a plurality of write operations, including a first write operation to write a first set of data bits into a first portion of the data storage element;

a set of components coupled to the PCB, the set of components including a plurality of double data rate (DDR) synchronous dynamic random access memory (SDRAM) devices coupled to the set of address and control lines and to the set of data lines, the DDR SDRAM devices being operable to receive or output data signals via the set of data lines based on address and control signals received from the system board via the set of address and control lines; and

first, second, third, and fourth converters coupled to the PCB and to the input voltage supply line and configured to receive power from the input voltage supply line and to deliver power via first, second, third, and fourth regulated voltage lines,

wherein the set of components includes:

a first component configured to receive power via the first regulated voltage line;

a second component configured to receive power via the second regulated voltage line;

a third component configured to receive power via the third regulated voltage line; and

a fourth component configured to receive power via the fourth regulated voltage line;

wherein each component of the set of components:

(i) is electrically connected to one or more of the first, second, third, and fourth regulated voltage lines;

(ii) receives power only via the one or more of the first, second, third, and fourth regulated voltage lines, and

wherein the plurality of DDR SDRAM devices includes:

a first group of at least five DDR SDRAM devices that are each connected to a first chip select line configured to electrically conduct a first chip select signal;

a second group of at least four DDR SDRAM devices that are each connected to a second chip select line configured to electrically conduct a second chip select signal;

wherein when power from the power rails of the system board is provided to the DIMM and each group of the first and second groups of DDR SDRAM devices is not in self-refresh mode:

the first group of DDR SDRAM packages is configured to be enabled, in response to the first chip select signal, to receive or output at least 40 1-bit DDR data signals in parallel via at least 40 data lines of the set of data lines independently of whether the second group of DDR SDRAM packages is enabled, in response to the second chip select signal, to receive or output an integer number of 1-bit DDR data signals in parallel via an integer number of data lines of the set of data lines; and the at least 40 1-bit DDR data signals include 32 1-bit DDR data signals and eight 1-bit DDR error-correcting code (ECC) data signals; and

US 12,373,366 B2

43

wherein the sum of the at least 40 data lines and the integer number of data lines matches a total number of data conduits of the set of data conduits of the memory slot connector, and the total number of data conduits is at least 72 data conduits.

17. The DIMM of claim 16, wherein the controller further includes a nonvolatile memory and a logic element, the logic element coupled to the nonvolatile memory and the data storage element and configured to electrically communicate with the nonvolatile memory and the data storage element.

18. The DIMM of claim 17, wherein the logic element is configured to perform a second write operation to write a second set of data bits into a second portion of the data storage element.

19. The DIMM of claim 18, wherein the plurality of write operations further includes a third write operation subsequent to the first and second write operations to write a third set of data bits from the data storage element into the nonvolatile memory, wherein the third set of data bits includes the first and second set of data bits written during the first and second write operations, and wherein, in response to the detected trigger condition, the logic element is configured to periodically perform every millisecond all of the first and second write operations followed by the third write operation.

20. The DIMM of claim 19, wherein, prior to performing the third write operation, the logic element is configured to perform, in response to the detected trigger condition, periodically at regular intervals of time one or more read operations to read the first and second sets of data bits from the first and second portions of the data storage element, respectively, and to generate the third set of data bits based on the read first and second sets of data bits.

21. The DIMM of claim 19, wherein the logic element is configured to perform the third write operation by writing the third set of data bits from the data storage element into the nonvolatile memory, and wherein the third set of data bits includes the first and second sets of data bits written during the first and second write operations.

22. The DIMM of claim 19, wherein, prior to performing the third write operation, the logic element is configured to generate the third set of data bits based on the groups of data bits stored in the data storage element.

23. The DIMM of claim 19, wherein the third write operation comprises a transfer of the first and second set of data bits from the data storage element to the nonvolatile memory via a first connection, and wherein the first and second write operations comprise a transfer of the first and second set of data bits into the data storage element via a second connection.

24. The DIMM of claim 18, wherein, prior to performing the first and second write operations, the logic element is configured to perform, in response to the detected trigger condition, periodically at regular intervals of time one or more read operations to read the first and second sets of data bits from the first and second groups of DDR SDRAM devices, respectively.

25. The DIMM of claim 16, wherein, when power is provided to the DIMM from the power rails of the system board, the nonvolatile memory is not powered using the first regulated voltage.

26. The DIMM of claim 16, wherein when power is provided to the DIMM from the power rails of the system board the voltage monitor circuit compares the input voltage to a plurality of threshold voltages including a first threshold voltage and a second threshold voltage.

44

27. The DIMM of claim 26, wherein when power is provided to the DIMM from the power rails of the system board the trigger condition occurs when the input voltage goes above a specified operating voltage.

28. The DIMM of claim 26, wherein when power is provided to the DIMM from the power rails of the system board, an additional trigger condition occurs when the input voltage is above the first threshold and below the second threshold; and

wherein, in response to the additional trigger condition, the voltage monitor circuit is further configured to generate and transmit a second trigger signal.

29. A dual in-line memory module (DIMM) configured to fit into a memory slot connector of a system board of a computer system, the memory slot connector including a set of power conduits, a set of data conduits, and a set of address and control conduits, wherein data signals are transferred between the system board and the DIMM via the set of data conduits, and address and control signals are delivered from the system board to the DIMM via the set of address and control conduits, the DIMM comprising:

a printed circuit board (PCB) having edge connections configured to fit into the memory slot connector, the PCB including:

a first set of edge connections configured to be electrically connected to the set of power conduits and operable to deliver to the DIMM power from the power rails of the system board via the set of power conduits, wherein the power delivered to the DIMM from the power rails of the system board is the only power received by the DIMM,

a second set of edge connections configured to be electrically connected to the set of data conduits,

a third set of edge connections configured to be electrically connected to the set of address and control conduits,

a set of voltage supply lines electrically connected to the first set of edge connections and configured to receive the power delivered to the DIMM from the power rails of the system board by way of the memory slot connector of the system board and the first set of edge connections of the DIMM,

a set of data lines electrically connected to the second set of edge connections, and

a set of address and control lines electrically connected to the third set of edge connections, wherein the set of data lines are electrically coupled to respective edge connections of the second set of edge connection, wherein each data line of the set of data lines is configured to electrically conduct a 1-bit DDR data signal;

a controller including a voltage monitor circuit and a data storage element configured to store a plurality of data bits, the controller coupled to the PCB and an input voltage supply line of the set of voltage supply lines, the voltage monitor circuit configured to:

(i) monitor an input voltage supplied via the input voltage supply line;

(ii) generate a trigger signal upon detecting a trigger condition; and

(iii) transmit the trigger signal to at least one other portion of the controller,

wherein the controller is configured to perform, in response to the trigger signal,

one or more write operations to write data into the data storage element;

US 12,373,366 B2

45

a set of components coupled to the PCB, the set of components including a plurality of double data rate (DDR) synchronous dynamic random access memory (SDRAM) devices coupled to the set of address and control lines and to the set of data lines, the DDR SDRAM devices being operable to receive or output data signals via the set of data lines based on address and control signals received from the system board via the set of address and control lines; and
 first, second, third, and fourth converter circuits coupled to the PCB and to the input voltage supply line and configured to receive power from the input voltage supply line and to deliver power via first, second, third, and fourth regulated voltage lines,
 wherein the set of components includes:
 a first component configured to receive power via the first regulated voltage line;
 a second component configured to receive power via the second regulated voltage line;
 a third component configured to receive power via the third regulated voltage line; and
 a fourth component configured to receive power via the fourth regulated voltage line;
 wherein each component of the set of components:
 (i) is electrically connected to one or more of the first, second, third, and fourth regulated voltage lines; and
 (ii) receives power only via the one or more of the first, second, third, and fourth regulated voltage lines;
 wherein the plurality of DDR SDRAM devices includes:
 a first group of at least five DDR SDRAM devices that are each connected to a first chip select line configured to electrically conduct a first chip select signal; and
 a second group of at least four DDR SDRAM devices that are each connected to a second chip select line configured to electrically conduct a second chip select signal, and
 wherein when power from the power rails of the system board is provided to the DIMM and each group of the first and second groups of DDR SDRAM devices is not in self-refresh mode:
 the first group of DDR SDRAM devices is configured to be enabled, in response to the first chip select signal, to receive or output at least 40 1-bit DDR data signals in parallel via at least 40 data lines of the set of data lines independently of whether the second group of DDR SDRAM devices is enabled, in response to the second chip select signal, to receive or output an integer number of 1-bit DDR data signals in parallel via an integer number of data lines of the set of data lines; and
 the at least 40 1-bit DDR data signals include 32 1-bit DDR data signals and eight 1-bit DDR error-correcting code (ECC) data signals; and

46

wherein the sum of the at least 40 data lines and the integer number of data lines matches a total number of data conduits of the set of data conduits of the memory slot connector, and the total number of data conduits is at least 72 data conduits.
 30. The DIMM of claim 29, wherein, when power is provided to the DIMM from the power rails of the system board, the nonvolatile memory is not powered using the first regulated voltage.
 31. The DIMM of claim 29, wherein when power is provided to the DIMM from the power rails of the system board, the nonvolatile memory is not powered using the first regulated voltage.
 32. The DIMM of claim 29, when power is provided to the DIMM from the power rails of the system board the voltage monitor circuit compares the input voltage to a plurality of threshold voltages including a first threshold voltage and a second threshold voltage.
 33. The DIMM of claim 32, wherein when power is provided to the DIMM from the power rails of the system board the trigger condition occurs when the input voltage goes above a specified operating voltage.
 34. The DIMM of claim 32, wherein, when power is provided to the DIMM from the power rails of the system board, an additional trigger condition occurs when the input voltage is above the first threshold and below the second threshold; and
 wherein, in response to the additional trigger condition, the voltage monitor circuit is further configured to generate and transmit a second trigger signal.
 35. The DIMM of claim 29, wherein the controller further includes a logic element coupled to the nonvolatile memory and the data storage element and configured to electrically communicate with the nonvolatile memory and the data storage element, and the one or more write operations include a first write operation to write a first set of data bits into a first portion of the data storage element and a second write operation to wire a second set of data bits into a second portion of the data storage element.
 36. The DIMM of claim 35, wherein the logic element is configured to perform a third write operation subsequent to the first and second write operations by writing a third set of data bits from the data storage element into the nonvolatile memory, and wherein the third set of data bits includes the first and second set of data bits written during the first and second write operations.
 37. The DIMM of claim 35, wherein the logic element is configured to perform, prior to the first and second write operations, first and second read operations by reading the first and second sets of data bits from the first and second groups of DDR SDRAM devices, respectively.
 38. The DIMM of claim 29, wherein each of the first, second, and third converter circuits is a buck converter.

* * * * *

EXHIBIT 4

JEDEC MANUAL

JEDEC Manual of Organization and Procedure

JM21T

(Revision of JM21S, November 2017)

SEPTEMBER 2020

JEDEC SOLID STATE TECHNOLOGY ASSOCIATION



NOTICE

JEDEC standards and publications contain material that has been prepared, reviewed, and approved through the JEDEC Board of Directors level and subsequently reviewed and approved by the JEDEC legal counsel.

JEDEC standards and publications are designed to serve the public interest through eliminating misunderstandings between manufacturers and purchasers, facilitating interchangeability and improvement of products, and assisting the purchaser in selecting and obtaining with minimum delay the proper product for use by those other than JEDEC members, whether the standard is to be used either domestically or internationally.

JEDEC standards and publications are adopted without regard to whether or not their adoption may involve patents or articles, materials, or processes. By such action JEDEC does not assume any liability to any patent owner, nor does it assume any obligation whatever to parties adopting the JEDEC standards or publications.

The information included in JEDEC standards and publications represents a sound approach to product specification and application, principally from the solid state device manufacturer viewpoint. Within the JEDEC organization there are procedures whereby a JEDEC standard or publication may be further processed and ultimately become an ANSI standard.

No claims to be in conformance with this standard may be made unless all requirements stated in the standard are met.

Inquiries, comments, and suggestions relative to the content of this JEDEC standard or publication should be addressed to JEDEC at the address below, or refer to www.jedec.org under Standards and Documents for alternative contact information.

Published by
©JEDEC Solid State Technology Association 2020
3103 North 10th Street
Suite 240 South
Arlington, VA 22201-2107

This document may be downloaded free of charge; however JEDEC retains the copyright on this material. By downloading this file the individual agrees not to charge for or resell the resulting material.

PRICE: Contact JEDEC

Printed in the U.S.A.
All rights reserved

PLEASE!

DON'T VIOLATE
THE
LAW!

This document is copyrighted by JEDEC and may not be
reproduced without permission.

For information, contact:

JEDEC Solid State Technology Association
3103 North 10th Street
Suite 240 South
Arlington, VA 22201-2107

or refer to www.jedec.org under Standards-Documents/Copyright Information.

JEDEC MANUAL OF ORGANIZATION AND PROCEDURE

Contents

		Page
1	JEDEC	1
1.1	Mission	1
1.2	Scope	1
1.3	Membership	1
1.4	Board of Directors	2
1.4.1	Eligibility	2
1.4.2	Election of Directors	2
1.4.3	Board representation changes	3
1.4.3.1	Merger	3
1.4.3.2	Acquisition	4
1.4.3.3	Split	4
1.4.4	Member Company and Director resignation	4
1.4.5	Officers of the Board of Directors	5
1.4.6	Board of Directors meetings	5
1.4.6.1	Board of Directors Executive session	5
1.4.6.2	Board of Directors voting	5
1.4.6.3	Emergency action by the Board of Directors	5
1.4.7	Board of Directors guest policy	5
1.5	Governance/Interpretation	6
1.6	Scope of participation	6
1.6.1	Interest categories	6
1.6.1.1	Producer	6
1.6.1.2	User	6
1.6.1.3	General Interest	6
1.7	Limits and restrictions	7
1.8	Modifications to the Manual of Organization and Procedure	7
 2	 Committees	 8
2.1	Committee responsibilities	8
2.2	Executive committee	9
2.3	Election of Officers	9
2.4	Government entities	9
 3	 Meetings	 10
3.1	General requirements	10
3.2	Notices	10
3.3	Agendas	10
3.4	Meeting attendance	11
3.5	Quorum	11
3.6	Recording devices	11
3.7	Minutes	12
3.7.1	Policy statement on JEDEC sign-in/attendance rosters	12
3.8	Regular meetings	13
3.9	Special (Limited Scope)	13
3.10	Joint meetings	13
3.10.1	Between JEDEC Committees	13
3.10.2	Between JEDEC and any non-JEDEC organization	13
3.10.3	Memorandum of Understanding (MoU)	13
3.11	Task groups	14
3.11.1	Joint task groups between JEDEC committees	14
3.11.2	Joint task groups between a JEDEC committee and other organizations	14

JEDEC MANUAL OF ORGANIZATION AND PROCEDURE

Contents (cont'd)

		Page
4	Document development procedure	15
4.1	Copyrights	16
5	Voting	16
5.1	Ballot voting	17
5.2	Ballot approval by acclamation	17
6	Ballots	18
6.1	Editorial changes	18
6.2	Ballot lobbying	18
6.3	Copyrights (see 4.1)	18
6.4	Regular committee ballot	18
6.5	Limited committee ballot	19
6.6	Board of Directors ballot	19
6.6.1	Committee ballots sent to the Board of Directors	19
6.7	Survey ballots	20
7	JEDEC documents	20
7.1	Interpretation	20
7.2	Document Maintenance	20
7.3	Document Types	21
7.3.1	Standards (JESD)	21
7.3.2	Publications (JEP)	21
7.3.3	Specifications (JES)	21
7.3.4	Registered Outlines (TO, MO, CO, PO, etc.)	21
7.3.5	Registration Data Formats (RDF)	21
7.3.6	Manuals (JM)	21
7.3.7	Joint Documents (J-STD, JP, and JS)	21
7.3.8	ANSI/JEDEC Standards	21
8	Legal Guidelines	22
8.1	JEDEC documents and publications	22
8.2	Patent policy	22
8.2.1	Terms and definitions	22
8.2.2	General provisions	24
8.2.2.1	Committee members	24
8.2.2.2	Committees	24
8.2.3	Disclosure of potentially essential patents	25
8.2.3.1	Disclosure of unwillingness to license work of a committee	25
8.2.4	RAND patent licensing commitment	26
8.2.5	Licensing assurance/disclosure form	26
8.2.6	Reference to patented products in JEDEC standards and publications	27
8.2.7	Special legal disclaimer	28
8.2.8	Miscellaneous	28
8.2.9	Violations	29
8.2.10	Interpretation and governing law	29

JEDEC MANUAL OF ORGANIZATION AND PROCEDURE

Contents (cont'd)

		Page
9	Appeals	29
9.1	Appeals related to ANSI/JEDEC standards and publications	29
9.1.1	Right of appeal	29
9.1.2	Criteria for appeals mechanism	30
9.1.3	Appeals procedures	30
9.1.3.1	Complaints	30
9.1.3.2	Response	30
9.1.3.3	Hearing	30
9.1.3.4	Conduct of the Hearing	31
9.1.3.5	Decision	31
9.1.4	Further appeal	31
9.2	Appeals not related to ANSI/JEDEC standards and publications	31
9.2.1	Right of appeal	31
9.2.2	Criteria for appeals mechanism	32
9.2.3	Appeals procedures	32
9.2.3.1	Complaints	32
9.2.3.2	Response	32
9.2.3.3	Hearing	32
9.2.3.4	Conduct of the Hearing	33
9.2.3.5	Decision	33
10	External publicity	33
11	JEDEC ethics rules for electronic communications	34
	Annex A (normative) Legal Guidelines Summary	35
A.1	JEDEC patent policy summary	35
A.2	Copyright material	35
A.3	License assurance/disclosure form	36
A.4	Notice of refusal to offer license and RAND terms	39
	Annex B (informative) Patent policy application guidelines	41
	Annex C (informative) ANSI Standardization Procedures	42

JEDEC Manual No. 21T

JEDEC MANUAL OF ORGANIZATION AND PROCEDURE

(From JEDEC Board Ballot JCB-02-31A, JCB-04-75, JCB-04-89, JCB-05-60, JCB-05-103, JCB-05-124, JCB-05-99C, JCB-06-02, JCB-06-36, JCB-08-11, JCB-09-65, JCB-09-73, JCB-09-82, JCB-09-86, JCB-10-15, JCB-10-21, JCB-10-88, JCB-12-38, JCB-13-46, JCB-13-47, JCB-15-14, JCB-17-24, JCB-19-04, JCB-19-27, JCB-20-14, and JCB-20-15, formulated under the cognizance of the JEDEC Board of Directors.)

1 JEDEC

JEDEC is an independent incorporated Association governed by a Board of Directors (Board). The Association facilitates standardization within the solid-state, microelectronics, and associated industries, and other related activities through special purpose committees.

The Association and its members are committed to foster open competition in the development of products and services.

1.1 Mission

The mission of JEDEC is to serve the solid state industry by creating, publishing, and promoting global acceptance of standards, and by providing a forum for technical exchange on leading industry topics.

1.2 Scope

The scope of JEDEC includes but is not necessarily limited to areas relating to (1) solid state devices, (2) integrated circuits, (3) electronic modules and associated electronic components, and (4) various manufacturing functions/processes.

1.3 Membership

Any company, organization, consumer, or individual conducting business that itself or through a related entity manufactures electronic equipment or electronics-related products, or provides electronics or electronics-related services, shall¹ be eligible for membership in JEDEC. Membership on a given committee shall be open to any JEDEC member company in good standing dependent solely on the dues structure defined by the Board². Government representatives may also participate as described in 2.4.

Each member company, organization, or individual shall be entitled to appoint one member and multiple alternates for each committee it joins. The member and alternates should be capable of making a technical contribution to the standards-setting process of the committee. To preserve the “one-company one-vote” concept within JEDEC, wholly-owned or majority-owned subsidiaries shall not be granted separate voting privileges within any single committee or subcommittee. The JEDEC Office shall maintain a list of all member companies and representatives for each committee.

¹ Special word usage. The word "shall" indicates a requirement; see JM7, Annex J.

² JEDEC membership is not a prerequisite to committee participation. Non-member participation fees will be charged. Before attending a JEDEC meeting, guests and non-members must agree to comply with all JEDEC rules and procedures, including the rules and procedures set forth in this Manual. Further details are available from the JEDEC Office. See 3.4.

1.3 Membership (cont'd)

All changes in committee membership shall be the responsibility of the member company, who shall so advise the JEDEC Office and the committee chairperson in writing or by electronic means.

NOTE Government regulations restrict the participation of non-U.S. Nationals at certain meetings. Every effort shall be made to notify member companies, well in advance, when a meeting will cover restricted topics.

1.4 Board of Directors

The governing body of JEDEC is the JEDEC Board of Directors (Board) which is comprised of individual Directors (or their alternates) representing JEDEC member companies.

1.4.1 Eligibility

Any maximum dues paying Member Company of JEDEC, having an identified representative in two or more committees for at least the past two consecutive years, may nominate one candidate for Board membership. The nomination consists of the nominee's resume accompanied by a letter from a company executive to the Board chairperson (in care of the JEDEC Office). The support letter is to request that the nominee be considered for Board membership and committing that if elected the member company will support the nominee's active attendance and participation in JEDEC Board meetings and affairs. If the application meets the above requirements, the JEDEC Office will hold the application on file with other (qualified) nominees. The JEDEC Office will maintain an official list of nominations to the Board. A list of the current Board Committees with a brief description and the membership of each will be sent to the nominee with an invitation for the nominee to participate where he/she has interest.

Seats on the Board are occupied by Directors representing individual JEDEC Member Companies. Once elected to a Seat on the Board, the Member Company maintains its position until the seat is voluntarily relinquished or forfeited due to other events as outlined in 1.4.4. Mergers, divestitures, acquisitions and other corporate changes that do not substantially alter the nature of a member company will not be grounds to forfeit a seat on the Board so long as the new entity continues to meet all Board Eligibility Requirements, see 1.4.3.

The total number of Directors shall not exceed thirty (30).

1.4.2 Election of Directors

The Membership Committee (chaired by appointment of the Chairman of the Board with membership open to any Board Member) shall review with the Board circumstances of a vacated seat and recommend qualified nominees from the official nominations list. In order to choose an appropriate nominee to fulfill objectives of the Board, the Membership Committee will review activities that are being conducted in active Board Committees with their respective Chairs.

1.4.2 Election of Directors (cont'd)

The following process will occur:

- Board Committee Chairs will be informed of the vacant seat by the Membership Committee.
- It will be the responsibility of the Membership Committee to qualify the level of support required to achieve ongoing Board Committee objectives.
- The Membership Committee will compile inputs and prepare a review to the Board at the next Board meeting including a recommendation on how to fill any vacated seats.

The following procedures and guidelines will be employed by the Membership Committee during the preparation of its Director recommendations:

- Gathering of Board Committee Chair inputs on current activities and resource requirements. Chairs will have an opportunity to review with the Membership Committee special requirements, shortcomings, etc. that should be made evident during the election process.
- Identification of the nominee's JEDEC experience and previous contributions. Tenure and amount of activity within a Committee or Task group will be identified.
- Identification of Board demographics in regards to representation of committees.
- Identification of the leadership qualities the candidate will bring to the Board.
- Identification of the unique value of the individual and the member company. (Value based on Committee Objectives identified to properly support JEDEC.)
- Identification of the individual's education and professional background (Engineering, Management, Finance, Marketing, Business development, etc.) and the balance the candidate will bring to the Board.

The Board will act on the recommendations of the Membership Committee. New Directors are elected to the Board by a 75% majority of the Board members voting at the meeting. There is no requirement that the maximum of 30 Seats be filled. The Membership Committee will send a response letter to the elected candidate/company and inform them of their responsibilities as outlined in this section. The elected Director's qualifications will be presented in regards to support as an active member on one or more of the Board Committees. The Board may stipulate a contingency period during which the newly-elected Director is required to serve on one or more Board Committees.

1.4.3 Board Representation Changes

The following outlines events that may occur that affect Board Representation.

1.4.3.1 Merger

Two companies on the Board merge or consolidate:

- It is up to the new surviving company (or the new parent company) to decide which Board representative remains on the Board to comply with the "one vote per company" rule.

1.4.3 Board Representation Changes (cont'd)

1.4.3.2 Acquisition

One Board company acquires another Board company:

- It is up to the acquiring company to decide which Board representative remains on the Board to comply with the “one vote per company” rule.

A Board company acquires a non-Board company:

- The Board company can elect to nominate a different Board representative if it wishes, but final approval is still required by the Board.

1.4.3.3 Merger or Acquisition by Non-Board Company

A non-Board company merges with or acquires a Board company:

- If the non-Board company is eligible for Board membership (ref. 1.4.1), it can assume the Board seat of the acquired company subject to Board approval.
- If the non-Board company is not eligible for Board membership, it cannot assume the Board seat of the acquired company and the Board seat is considered open.

1.4.3.4 Split

A Board company splits into two or more companies:

- If the Board representative associated with the split (or spun-off subsidiary) is still employed by the member company then his/her status is unchanged.
- If the Board representative associated with the split is now a member of the new (or non-member) company then representative in question loses his/her seat.
 - Exceptions: Board decides that the Board representative in question is vital to the Board and the new member company complies with all JEDEC requirements to be on the Board.
- The member company that still has a seat on the Board can nominate a new representative to the Board who is then subject to Board approval.

1.4.4 Member company and Director resignation

The Director will relinquish a position on the Board in the instance of any of the following events:

- 1) The Member Company formally resigns the position. This includes a statement from the Member Company that confirms the Director and Member Company resignation.
- 2) The Member Company which the Director represents is no longer eligible as defined in 1.4.1.
- 3) The Member Company is removed by a majority vote of the Board for gross violation of JEDEC policy.
- 4) The Director is absent from and not represented by a Member Company's approved alternate at two consecutive Board meetings. It is the expectation that Directors attend all Board meetings, at minimum Directors must attend at least one face-to-face Board meeting, in person, each calendar year.

1.4.4 Member company and Director resignation (cont'd)

In the event that a Director no longer represents the Member Company for which the individual was elected (including but not limited to retirement or change of employment) or should the Director be removed due to gross violation of JEDEC policy outside of the control of the Member Company, then the Board seat shall not be considered vacated and the Member Company shall have the right to maintain its position on the Board by nominating a replacement Director. The replacement Director must meet all nomination requirements as set forth in 1.4.1 and be elected by the full Board by a 75% majority of the Board members voting at the meeting. Should the replacement Director fail to be elected, the Member Company may submit a second candidate and repeat the process.

1.4.5 Officers of the Board of Directors

The officers of the Board shall follow the same guidelines as the officers of a committee with the following exceptions: There shall be at least one Vice Chairperson on the Board. The number of Vice Chairpersons on the Board shall be determined by regular vote at any regularly scheduled Board meeting.

1.4.6 Board of Directors meetings

All meetings of the Board shall follow the same procedures and guidelines as outlined for regular committee meetings, with exceptions per 1.4.6.1, 1.4.6.2, and 1.4.6.3.

1.4.6.1 Board of Directors executive session

The needs of the Association occasionally require the Board to meet in executive session. Such executive sessions are to be defined as closed meetings of Board member company representatives and guests invited for specific purposes only.

1.4.6.2 Board of Directors voting

All actions within the Board require an affirmative vote of at least 75%; members abstaining shall not be considered for determination of the 75% criterion. When defined and agreed to by a duly recorded vote of the Board, specific action may require polling of all members whether present or not and/or may require a higher percentage affirmative vote.

1.4.6.3 Emergency action by the Board of Directors

If action normally requiring Board consideration must be taken, and a meeting cannot be reasonably convened, then the following modifications apply to the approval process: 1) A reasonable attempt must be made to notify 100% of Board members, and 2) The decision shall become binding only upon receipt by the JEDEC Office of affirmative votes from 75% of the members of the Board.

1.4.7 Board of Directors guest policy

All guests must be approved by the chairman and must have a purpose related to the topics of discussions at the Board meeting. The guest must be introduced to the Board at the meeting and the purpose of the guest's visit should be stated in the agenda for that meeting.

At the discretion of the Board, a guest may be allowed to stay beyond the allotted time slot if the Board feels there is a need for that guest's continued attendance. Otherwise, all guests may only attend the portion of the Board meeting that is directly related to their purpose in the meeting. At no time may a guest speak to the Board beyond the intended purpose of his/her invitation without the explicit approval of the Board.

At anytime, any member of the Board may request that the chairman, at his discretion, ask a guest to leave the meeting.

1.5 Governance/Interpretation

The Board shall have the sole authority to interpret, in whole or in part, these operating procedures. In any case where conflict arises from ambiguity, or lack of specificity, or misunderstanding within or related to this document, the officers of the Board of the Association may make temporary rulings to be confirmed or modified at or before the next regularly scheduled Board meeting by vote of the Directors. In any case where there is not agreement between 100% of the officers of the Board concerning an interpretation, then the matter must be referred to the entire Board for resolution.

1.6 Scope of participation

The standards development process shall not be dominated by any single interest category, individual, or organization. Dominance means a position or exercise of dominate authority, leadership, or influence by reason of superior leverage, strength, or representation to the exclusion of fair and equitable consideration of other viewpoints.

The standards development process should have a balance of interests. Participants from diverse interest categories shall be sought with the objective of achieving balance.

NOTE In the case of an ANSI American National Standard, JEDEC shall comply with the criteria for balance as contained in ANSI Essential Requirements.

1.6.1 Interest categories

For a formulating group to work effectively, all interest categories should be represented, and no interest category shall be intentionally excluded. Because standards development depends on consensus rather than the will of a majority, participation in the consensus-building process is of far greater significance than the final act of voting in a formulating group. Unless it is claimed by a directly and materially affected interest that a single interest category dominated the standards development process, no test for dominance is required.

Interest categories are defined relative to each standards activity: producer; user; and general interest. See ANSI Essential Requirements for further information.

1.6.1 Interest categories (cont'd)

1.6.1.1 Producer

Producers are formulating group members that manufacture electronic equipment or electronics-related products covered by the standard.

1.6.1.2 User

Users are members that utilize electronic equipment or electronics-related products covered by the standard to provide services to an end user.

1.6.1.3 General interest

General Interest members are neither producers nor users. This category includes, but is not limited to, regulatory agencies (state and federal), researchers, other organizations and associations, end users and consumers.

1.7 Limits and restrictions

The Members (particularly where they are direct competitors in certain lines of business, and including their employees and representatives), shall, when engaged in the activities and affairs of the Association, act in a manner that does not violate any state, federal or international laws or regulations.

Member companies must abide by all applicable export control regulations of their respective countries while attending a JEDEC meeting, irrespective of where the meeting is actually held. It is recommended that the member seek guidance from his or her company if involved in "state of the art" engineering work, as significant changes are made to the export control regulations from time to time.

The Association prohibits any discussion on specific costs, prices, quantity or yield of production levels, methods or channels of distribution, markets, customers or any other topic, or any conduct that may be construed as a violation of antitrust laws. Each respective member company shall advise its respective participant representatives and employees on the importance of compliance with applicable state, federal or international antitrust laws or regulations and require the scope of any discussions or exchanges amongst representatives of Members and/or the Association, be limited to the mission and scope of the Association stated above.

1.8 Modifications to the Manual of Organization and Procedure

Modifications to this manual may be made in whole or in part, upon approval by the Board by the normal ballot process. Any changes to this Manual shall be distributed to the JEDEC membership at the earliest possible time upon adoption.

2 Committees

The Board shall have the authority to establish committees to perform any task(s) for the association. JEDEC committees (main and sub) shall be authorized only by the Board, and the Board shall approve the scope of each³. A main committee shall be designated by the prefix JC- plus a two-digit designator, such as committee JC-11. Main committees are separated into two types: (1) service committees, whose scope is limited to a specific subject that may impact several product types, and (2) product committees, whose scope is limited to a specific type product. For service committees, the first digit of the two-digit designator shall be the number 1, and for product committees, the first digit shall be the number 2 or higher.

When the main committee finds it necessary to divide its work into parts, subcommittees may be formed under the scope of the main committee and such subcommittees shall be designated by adding a period to the main committee number followed by a number or a number-and-letter designator (e.g., JC-42.3 or JC-42.3B).

Subcommittees shall have all the same authority and responsibilities as committees herein defined, except that a subcommittee scope shall be a subset of a committee scope, to be authorized, reviewed, and approved by the Board.

NOTE Participation shall be open to all persons who are directly and materially affected by the activity in question.

2.1 Committee responsibilities

Each committee shall have the authority to consider all matters within its approved scope, including but not limited to the following actions:

- a) Develop proposed standards, guidelines, registrations, and related technical publications for committee ballot.
- b) Authorize committee ballots on all matters within the committee's defined scope.
- c) Initiate proposals for Board ballots of previously approved committee ballots.
- d) Evaluate the results of all committee ballots, as well as comments on related Board ballots.
- e) Request approval of the Board to convert standards or publications into ANSI or IEC standards.
- f) Maintain liaison with other national and international organizations.
- g) Initiate press releases to announce a significant advancement in the work of solid state technology covered within the committee scope such as: registrations, standardization, or publication. All press releases are to be issued by the JEDEC Office after review and concurrence by the JEDEC Legal Counsel.

³ The committee scope shall go through the standard balloting process and must be completed within a reasonable timeframe and prior to any publication of a JEDEC document authored by the committee.

2.2 Executive committee

A JEDEC committee or subcommittee may form an executive committee to manage its administrative work. The executive committee shall meet the following criteria:

- a) The executive committee should be chaired by the committee or subcommittee chairperson and shall include as a minimum all chairpersons and vice chairpersons of that committee and its subcommittees.
- b) Meetings of the executive committee shall be open to any committee member who gives reasonable advance notice to the executive committee chairperson or secretary.
- c) Actions of the executive committee shall be limited to handling administrative matters dealing with such items as meeting schedule, meeting agenda, future meeting dates and locations, personnel requirements, proposed work priorities, and the like. Any decision made by the executive committee is subject to review and modification by the full committee at the request of any committee member or alternate. Decisions on the final disposition of all technical matters shall be made only by the overseeing committee. Recommendations concerning future meeting dates and locations should be presented to the overseeing committee.

2.3 Election of officers

All JEDEC committee officers shall be members or alternates of their respective committees. All JEDEC committees shall have a chairperson and may also have a vice chairperson as officer(s) of the committee. All officers shall be directly elected either by secret vote of the committee members present during a regular meeting or by ballot as determined by vote of the committee. Committee officers shall serve a term of two (2) years. The JEDEC staff shall count the secret ballot and announce the result at the meeting.

At least two meetings prior to the end of a chairperson's term, he/she shall appoint a nominating committee to propose candidates to fill the next term of office(s) for the committee. The committee's candidates must be presented to the full committee at the next meeting and nominations from the floor shall also be accepted. All nominees must agree to accept the additional responsibilities before they can be considered.

If any of the officers cannot fulfill his/her term then elections shall be held at the earliest opportunity of the committee members. Similarly, if a new committee is formed, then the chairperson shall be selected by vote of the members attending the first meeting as a first order of business of the committee. These special case elections shall become effective immediately.

2.4 Government entities

JEDEC desires and encourages the active participation in its standards-developing activities of all parties having a direct and material interest in the standards, including U.S. federal, state, and local government entities when they have such a direct and material interest via the JC-13 Committee, Government Liaison.

Contribution to the standards development process may take the form of liaison with appropriate formulating groups, participation in the ANSI public review process, and active participation in formulating groups.

3 Meetings

Communication between JEDEC member companies is essential for the work within the Association. All committees shall hold regular meetings for the purpose of conducting JEDEC business.

The use of teleconference or other electronic conferencing equipment is permitted for any type of JEDEC meeting.

3.1 General requirements

Every JEDEC meeting, regardless of type, shall be held in accordance with the directives contained within this manual and in accordance with all applicable laws.

All meeting notices, agendas, ballots, and minutes shall be distributed only by the JEDEC Office. All correspondence from the JEDEC Office shall only be by electronic means unless otherwise required by law. If a company requests distribution by other means, the JEDEC Office will charge a fee for such additional service as directed by the JEDEC Board. The JEDEC website, www.jedec.org, shall be the primary method of electronic communication.

3.2 Notices

Notices of regular meetings shall be distributed solely by the JEDEC Office no less than 28 calendar days prior to the start of such meeting, unless a “special” meeting is authorized by committee vote and the date and agenda for such meeting is fixed at a regularly scheduled meeting. Notices of such special meetings shall be distributed as soon as possible after said regular meeting and in no case less than 10 calendar days prior to the start of such special meeting.

3.3 Agendas

All meetings shall have prepared agendas distributed as soon as possible, but in no case less than 10 days prior to a meeting and should be included with the meeting notice if possible. All special meeting notices shall include the agenda at time of distribution of such notice.

3.4 Meeting attendance

All JEDEC committee and subcommittee meetings are open to members, their designated alternates, and guests invited by the committee or subcommittee chairperson.

A JEDEC staff member or a designated alternate must participate in every committee meeting, either in person or by electronic means. The duties of the staff person shall include but not be limited to taking minutes, advising the chairperson regarding the content of this manual and JEDEC legal policies and guidelines, and distributing the minutes. JEDEC staff attendance is not required at every subcommittee meeting.

A potential new member company may be a guest for one meeting with prior approval of the chairperson. A timely submission of such a request stating the purpose of attendance shall be given so that appropriate accommodations can be made. If there are any issues to be raised by a guest, those issues must be specified in the purpose statement.

Nonmembers may participate in committee meetings if they pay a nonmember participation fee, the fee shall be predetermined, fixed, and reasonable. A procedure for requesting a fee waiver or fee reduction shall be available. Nonmembers do not have the right to vote. Nonmembers should contact the JEDEC Office prior to the meeting date for further details.

Before attending a JEDEC meeting, guests and non-members must agree to comply with all JEDEC rules and procedures, including the rules and procedures set forth in this Manual.

An individual representative may represent only one company at any committee, subcommittee or task group meeting, and may vote on behalf of only one company on any ballot placed on the electronic voting machine. It is the responsibility of the committee chairman to enforce this rule during a committee or subcommittee meeting. For the electronic voting machine, it is the responsibility of the ballot sponsor to enforce this rule when the ballot tally is reported.

3.5 Quorum

The JEDEC Office shall maintain a list of active members for the sole purpose of determining the quorum requirements for a committee meeting. An active member is defined as a member company that has been represented at either or both of the previous two consecutive regular meetings of the committee. Special meetings shall not be considered for the purpose of determining “active” member status. The quorum requirements shall be 50% by count of a committee’s active member list, but in no case shall the quorum requirement be greater than 12 member companies regardless of committee size. All committee member companies present at a meeting are counted toward quorum requirement.

3.6 Recording devices

Voice or video recording devices of any kind are forbidden at any type of JEDEC meeting.

3.7 Minutes

The principal activities and attendees at all JEDEC meetings shall be duly recorded in meeting minutes. Meeting minutes shall be taken by the assigned JEDEC staff for that committee or by a substitute designated by the JEDEC staff. In the absence of a JEDEC staff person at a committee meeting, a temporary secretary may be assigned by the committee chairperson.

Minutes shall be termed “unconfirmed” until reviewed and confirmed by the committee chairperson and JEDEC Legal Counsel. In the interest of member company needs, minutes should be distributed to members and alternates within 14 calendar days after a meeting. If the chairperson is unable to review and confirm the minutes within that time, the JEDEC Office should distribute the unconfirmed minutes clearly labeled as unconfirmed minutes. Unconfirmed minutes shall be approved by JEDEC Legal Counsel prior to distribution. Confirmed meeting minutes from regular meetings shall be distributed no later than 28 days prior to the next regular meeting. Confirmed minutes from “special” meetings should be distributed as soon as possible after such special meeting and in no case less than 10 days prior to the next regular meeting.

At the end of each set of minutes the following statements shall appear:

“Reviewed/corrected and authorized for release” followed by the Chairperson’s and Secretary’s name and date.

Chairperson, Name / date

Secretary, Name / date

Confirmed Legal Approval Date

“This meeting was conducted in accordance with JEDEC Legal Guides and the JEDEC Manual of Organization and Procedure.”

For those situations where a report from an independent group is included in the minutes, for information purposes, then the following disclaimer may be added to the minutes:

“Disclaimer: All liaison, government, technical, and task group reports contained herein are believed to be accurate but the accuracy or completeness thereof is not guaranteed.”

Minutes of a task group meeting shall include a list of the attendees, the meeting agenda, a status review of the agenda items, and any final proposals including discussion related to the proposals. The time, the place, and any required follow-up should be included. Task group minutes are reported and approved as part of the minutes for the next meeting of the sponsoring committee.

3.7.1 Policy statement on JEDEC sign-in/attendance rosters

The policy statements reproduced on the reverse side of the JEDEC sign-in/attendance rosters shall become an integral part of the minutes of all committee, subcommittee, and task group meetings; and a statement shall be included in the minutes indicating compliance with the foregoing policy.

3.8 Regular meetings

Every JEDEC committee shall set up an annual calendar for regular meetings at which any item covered under the scope of the committee may be discussed. Under the JEDEC guidelines, committees must hold at least one (1) regular meeting each year to retain status as an active committee.

3.9 Special (Limited Scope)

All JEDEC committees shall have the authority to approve meetings of limited scope. Schedule for the meeting and specific agenda items must be approved by formal vote during a regular committee meeting. During a limited scope meeting, no other committee business shall be conducted outside of the specific items on the previously agreed to agenda. Items shall not be added to the agenda of a special meeting.

3.10 Joint meetings

3.10.1 Between JEDEC committees

Joint meetings between two or more JEDEC committees may be authorized by agreement of the committees by vote. Selection of the joint meeting chairperson shall occur by vote of the committees' attendees. The joint chairperson shall ensure that all meeting notices, agendas, and minutes are distributed to the relevant JEDEC committees.

3.10.2 Between JEDEC and any non-JEDEC organization

Joint meetings between a JEDEC committee and any non-JEDEC organization are subject to the same requirements as 3.10.1. In addition, JEDEC Legal Counsel shall establish a Memorandum of Understanding (MoU), see 3.10.3, between the organizations and review all release of information from the meetings to any outside person or organization.

JEDEC committees may meet jointly with other organizations or agencies for the purpose of permanent or ongoing liaison. In these cases, those nonmembers attending the committee meeting are expected to pay their proportionate part of the total meeting costs. Apportionment of the meeting costs shall be agreed upon between the committee chair and the representatives of those groups participating in the meeting(s), subject to the approval of the JEDEC Board.

3.10.3 Memorandum of Understanding (MoU) with any non-JEDEC organization

Prior to the formation of a joint meeting/task group with a non-JEDEC committee, a Memorandum of Understanding (MoU) must be in effect between JEDEC and the other organization covering the following items:

- a) Balloting procedure (acceptance criteria, resolution of comments, and disapprovals),
- b) Patents procedure (including joint ownership),
- c) Distribution rights (copyright, use of association logos, available on JEDEC website without charge),
- d) Procedures for future revisions of the document(s),
- e) Selection of chairperson(s) and appointment of task group membership, and
- f) Meetings and related legal requirements (notices, agendas, motions, minutes, costs, etc.).

3.11 Task groups

A task group may be authorized by any committee chairperson for a specific non-continuing objective. The task group chairperson, appointed by the committee chairperson, shall be a member or alternate of the overseeing committee. Membership within a task group is open to any member or alternate of the committee within which the task group is formed. In addition, members of the task group may also be appointed by the task group chairperson from nonmember JEDEC companies with approval of the overseeing committee chairperson; appointment shall be based solely on ability to contribute to the work; however, all members of a task group shall conduct their work in accordance with this manual.

A task group is not permitted to hold any formal votes or establish any final results relative to standards, publications, or ballots. All task group activities shall be presented to the overseeing committee for review, modification, and approval in the form of reports generated from every task meeting. By duly recorded committee vote, committees may authorize a task group to issue ballots or other specific tasks as required, the specific exception being that committees cannot authorize a task group to conduct a formal vote. A task group must be dissolved upon completion of its objectives.

Scheduling of task group meetings is subject to approval of the sponsoring committee's chairperson, in order to avoid conflict with regularly scheduled meetings.

3.11.1 Joint task groups between JEDEC committees

Joint task groups between two or more JEDEC committees may be authorized by agreement of the committee chairpersons. Selection of the task group chairperson shall be part of this agreement and that chairperson must be a member or alternate of one or more committees involved in the task group. The task group chairperson shall keep all sponsoring committees on the distribution list for all meeting notices, agendas, and reports. Joint task groups between JEDEC committees shall comply with all other requirements for task groups.

3.11.2 Joint task groups between a JEDEC committee and other organizations

Joint task groups between a JEDEC committee and other non-JEDEC organizations are subject to the same requirements as task groups per 3.11. Unless otherwise authorized by the Board and approved by JEDEC Legal Counsel, such joint task groups shall establish a Memorandum of Understanding (MoU), see 3.10.3, in the same manner as joint committees covered in 3.10.2.

4 Document development procedure

The development of all JEDEC documents shall follow the minimum sequential procedure outlined herein.

- 1) Presentation – initial showing or distribution of material proposed for publication. Presentations may occur at any regular meeting or special meeting if the limited scope and agenda include such information. Presentations may also occur entirely electronically without being at a meeting. An additional showing of material proposed for publication may be required.
- 2) Motion for ballot – may occur at any time after a presentation. The motion must have a second before being considered for vote. The company agreeing to second the motion for ballot should be willing to accept responsibility for sponsoring the material should the original sponsor be unable or unwilling to do so. The motion for ballot should be discussed after receiving a second and then a formal vote shall be taken.
- 3) All committee ballots shall be subjected to the committee’s editorial process prior to being posted on the JEDEC voting machine.
- 4) Committee ballot – if approved, material shall be sent out to ballot through the JEDEC voting machine. Details of balloting procedure are covered in section 6.
- 5) Motion to Board ballot – if a ballot passes and all comments are addressed then a motion may be made to send the material on to the Board for further consideration and approval. Alternatively, the motion to forward the ballot to the Board is unnecessary if the ballot passes by acclamation, see 5.2.
- 6) Board ballot – when approved, the committee material shall be sent out to ballot through the JEDEC voting machine. Board ballots follow the same procedures as committee ballots, see 6.6.1(c). When reviewing material passed by committee, the Board primarily considers whether due process has been followed.
- 7) Publication – once proposed material has successfully completed all the above steps and obtained approval through the appropriate balloting processes, it is approved for publication.
- 8) All JEDEC documents shall be published by the JEDEC Office.

NOTE 1 Steps may be repeated as many times as required to satisfy the wishes of the member companies participating in the process.

NOTE 2 JEDEC Registrations are special-case publications that do not have to be reviewed by the JEDEC Board of Directors, and the registration procedure ends upon the completion of step 4 above.

4.1 Copyrights

JEDEC owns the copyright to all documents created through JEDEC, and use of that copyrighted material is at the sole discretion of the JEDEC Board.

For the protection of member companies, all JEDEC publications and standards shall be copyrighted. As a precaution, draft proposals circulated for comment and/or vote shall display the symbol ©, followed by the year and “JEDEC”. Stapled or bound documents require this designation only on the first page or title page; short documents that may be distributed unstapled shall be so labeled on every page. When final documents are ready for publication, the JEDEC Office shall register the copyright and maintain all appropriate records.

JEDEC member companies may reproduce any JEDEC document for internal use, without restriction. Use of JEDEC material outside of member companies is at the discretion of the Board. Generally, all JEDEC documents are available on the JEDEC website: www.jedec.org.

If, in developing a standard, a publication, or a specification, a committee proposes to incorporate verbatim material from a publication copyrighted by another organization, the committee shall obtain written permission from the owner of the copyright for JEDEC to copy the material. The letter granting permission shall be sent to the JEDEC Office, and the standard or publication shall include a reference to such permission as a footnote.

5 Voting

JEDEC shall formally operate under the rule of “One company, one vote” wherein all formal, binding votes will be restricted to one vote per company. Nonbinding fact-finding “straw” or survey votes are not considered formal votes. Unless specifically authorized by the Board, all approval vote counts within JEDEC shall be a minimum of 2/3 of votes cast. All Board approval vote counts shall be a minimum of 3/4 of votes cast. Abstention votes shall not be considered in the calculation.

Votes taken during committee meetings shall be duly recorded in the meeting minutes, and, except for unanimous votes, the vote count shall be included in those same meeting minutes.

In order to maintain fairness within every committee, the committee chairperson shall refrain from voting during a meeting for all but secret votes. If a chairperson has a specific company position on a topic being discussed, then control of the meeting should be turned over to the vice chairperson or other less biased member. Chairpersons may vote on any ballot to represent his/her company position.

5.1 Ballot voting

The period for voting on committee ballots shall be at least 21 calendar days from the date of issuance, and shall not exceed 90 days. In any case, the ballot closes at midnight, Eastern Time, on the final day of the voting period. All committee ballots shall be processed through the JEDEC electronic voting machine, and shall include the following vote options:

- a) approve, comments are optional with approval
- b) disapprove, comment is required and shall include a full explanation of the reason(s) for the disapproval
- c) abstain, comments are optional with abstain

NOTE All committee ballots shall contain the following patent statement:

“If anyone receiving this ballot is aware of any patents (granted or pending) involving in this ballot, check here and notify the committee, citing the applicable patent numbers.”

No one can have access to the ballot results before the voting period ends. As part of the deliberative, consensus-building process, commenters may send questions to the ballot sponsor independent of their votes via email, telephone, etc.

5.2 Ballot approval by acclamation

A committee ballot is considered approved by acclamation in either of the following two cases:

- a) The ballot is approved unanimously with no comment at the close of the voting period.
- b) The ballot is approved by count at the end of the voting period, and all comments whether from approve, disapprove, or abstain votes are resolved to the satisfaction of the sponsor and the commenting company. If any company is not satisfied with the proposed resolution of a comment, then resolution of the ballot must await committee action.

Any ballot approved by acclamation shall automatically advance to the next level in the publication process (i.e., committee ballots are automatically submitted to the Board, or Board ballots automatically become approved for publication). At the discretion of the committee chairperson, ballots that otherwise meet the requirement for approval by acclamation may be referred back to committee for further action if special conditions exist (e.g., uncharacteristically low vote count for the committee).

6 Ballots

All committee ballots shall be distributed ONLY by the JEDEC Office with Board-approved electronic distribution methods to all committee members and alternates for vote and to all related committees for information and comment. A submission for ballot shall be made through the JEDEC Office ONLY by Board approved editable and non-editable “soft copy” format of the complete ballot content.

Reference to any individual(s) and/or member company name(s) is prohibited from appearing on any committee ballot unless it is part of a return address.

In order to facilitate document identification and retrieval, keywords shall be submitted with all ballots. A keyword is defined as a word, term, or acronym that is descriptive of the essential contents of the ballot or a major portion thereof.

6.1 Editorial changes

During a committee meeting, provided there is no disagreement within the committee that the changes are editorial, these changes may be incorporated into committee-approved balloted material without reballoting.

6.2 Ballot lobbying

Member companies should not lobby or coerce other member companies to approve or disapprove ballots. Open discussion of balloted material is permitted and encouraged, but such discussions should focus on the technical aspects of balloted material.

The use of verbal, physical or business threats, intimidation, harassment or coercion is unprofessional behavior and unacceptable conduct at or in connection with any activity sponsored by JEDEC, including committee and task group meetings.

6.3 Copyrights

See 4.1

6.4 Regular committee ballot

A “regular” ballot is one in which the entire standard, publication, registration, guideline, or specification is open for comment.

Prompt consideration shall be given to the written comments and objections of all participants during the committee balloting process. An effort to resolve all expressed objections accompanied by comments related to the proposal under consideration shall be made, and the disposition of each such comment or objection and the reasons therefore shall be documented in writing (including electronic minutes and ballot material).

6.5 Limited committee ballot

If a limited number of revisions to an existing document are being proposed, a “limited” ballot may be issued. Limited ballots may be restricted to considering only the proposed areas for modification, addition, or deletion. In the cases of limited ballots both the original material and the proposed revisions thereto shall be clearly shown in the ballot material. Examples of appropriate occasions for using limited ballots are (1) a proposal to add a new definition to a long list of definitions or (2) a new package outline to JEDEC Publication No. 95.

On occasion, after a committee ballot has been approved, the committee may wish to propose minor technical changes to the ballot before forwarding the approved material on to the Board for final consideration. In such cases, a motion may be made for a limited ballot in which only these technical changes may be addressed. If the limited ballot is passed, these changes shall be incorporated. If it fails, the original material stands as approved.

NOTE This process is not intended for use in the development of an ANS.

6.6 Board of Directors ballot

In general, all ballots originating at the Board level shall follow the same guidelines outlined in 6.4. Board ballots are restricted in scope to non-technical Association business and specifically may not include new technical for publication creation as such material must first be submitted by the appropriate technical committee.

6.6.1 Committee ballots sent to the Board of Directors

All passed committee ballots relating to proposed JEDEC standards, guidelines, or specifications must be reviewed and approved by the JEDEC Board before becoming official JEDEC publications. Upon approval by acclamation or by committee vote⁴, the originating committee shall submit such passed committee ballot⁵ to the JEDEC Office for distribution as a Board ballot.

All ballots submitted to the JEDEC Board must include the committee ballot material, and:

- a) a brief description of the recent committee background regarding the material.
- b) the complete voting report including the names and votes of the member companies voting and all comments on the material.
- c) a description of the resolution of all of the comments on the material. Negative votes must be addressed, however if they cannot be resolved, the ballot can be submitted to the Board with the approval of the committee.

The Board may make only editorial modifications to any material proposed by committee for further action. Any technical modifications of committee material must be referred back to the originating committee for possible action.

⁴ Unless the committee ballot passes by acclamation, a formal committee vote must be taken and recorded in the minutes to send the committee ballot to the Board for further action.

⁵ Ballot material submitted for a Board ballot must be identical to the ballot material approved by the committee. The committee may authorize an editorial committee to review the document for editorial corrections only prior to issuing the Board ballot. In this situation, the Board ballot must be a marked up version of the committee-approved ballot material, such that all edits are clearly noted.

6.7 Survey ballots

Surveys may be issued to determine whether a proposal warrants further action. A survey may ask for information or for a support/nonsupport position. No committee vote is required for the distribution of a survey ballot; however, every survey must be approved by the committee chairperson and issued by the JEDEC Office following legal approval. The JEDEC Office shall monitor every survey to ensure adherence to JEDEC legal guidelines. Surveys that contain company-sensitive material that must remain confidential shall be stripped of any reference to a company by the JEDEC Office prior to use of the material. In all cases, survey results shall be presented to the full committee at the first meeting following the voting period expiration date.

Survey ballots are for gathering information or further clarifying member positions and survey votes are strictly nonbinding. All surveys should contain the following background statement:

“Companies are under no obligation to respond to this survey.”

The committee shall determine the method of data collection to ensure company confidentiality if required. In the case of anonymous surveys, the names of respondents shall be deleted by the JEDEC Office. The survey shall specify the method of data-collection.

7 JEDEC documents

JEDEC documents are available from the JEDEC website: www.jedec.org. JEDEC documents may not be sold without written authorization from JEDEC.

7.1 Interpretation

Official interpretations involving technical matters pertaining to JEDEC and/or ANSI/JEDEC standards and publications shall be made by the formulating committee of the document in question.

7.2 Document Maintenance

JEDEC documents shall be reviewed by the formulating committee within five years of the last publication date, after which the committee shall take one of the following actions with respect to the standards material:

- a) Reaffirm it as a JEDEC document
- b) Recommend its revision and reissuance as a JEDEC document
- c) Recommend its rescission
- d) Recommend its consideration as an ANSI or IEC document proposal.

7.3 Document Types

7.3.1 Standards (JESD)

Documents that are intended to represent a sound approach to product specification and application.

7.3.2 Publications (JEP)

Documents containing general engineering information on products, procedures, or services, that are not necessarily appropriate for standardization.

7.3.3 Specifications (JES)⁶

Documents prepared specifically to facilitate procurement, that clearly and accurately describes the essential technical requirements for purchased material or services.

7.3.4 Registered Outlines/Standards (CO, DO, GS, MO, MS, etc.)

Mechanical package outline drawings that have been registered and assigned a JEDEC registration number.

7.3.5 Registration Data Formats (RDF)

Engineering information in the form of a list of specific electrical and mechanical parameters necessary for the characterization of a particular class of solid state devices.

7.3.6 Manuals (JM)

Documents that describe internal JEDEC policies and procedures.

7.3.7 Joint Documents (J-STD, JP, and JS)

Documents created in cooperation between JEDEC and another organization.

7.3.8 ANSI/JEDEC Standards

At the request of the sponsoring committee and after approval by Board ballot, JEDEC standards and publications may be submitted to ANSI through the JEDEC Office for processing as an ANSI National Standard. The JEDEC Office will be responsible for providing appropriate documentation to ANSI. All policies and procedures set forth by ANSI per the “ANSI Essential Requirements: Due Process Requirements for American National Standards” shall be followed. See Annex D for more details.

⁶ This document type has not been in use by JEDEC since early 1990’s but is maintained here for historical purposes only.

8 Legal guidelines

All meetings of the JEDEC Board and its associated committees, subcommittees, task groups, and other units shall comply with JEDEC Legal Guidelines incorporated herein by reference. JEDEC Legal Counsel may advise the Board and committees from time to time concerning interpretation of Legal Guidelines.

8.1 JEDEC documents and publications

Prior to publication, the following documents including, but not limited to, all JEDEC standards, publications, manuals, minutes, surveys, questionnaires, press releases, registration data formats, and workshop proceedings, formulated by the Board, committees, and subcommittees, must be approved by JEDEC Legal Counsel.

8.2 Patent Policy

8.2.1 Terms and definitions

Affiliate: An entity owned or controlled by a Committee Member, or which controls or is under common control with a Committee Member for so long as such control exists.

NOTE For purposes of this definition, control means the ownership of greater than 50% of the voting securities of that Committee Member.

Approval by the Committee: An approval by vote of the committee to send a proposed Standard to the Board of Directors for review and approval. Some committees approve proposed Standards in parts, for example, by a "pass-hold" process. In those instances, "Approval by the Committee" means preliminary or final approval of the relevant part of a proposed Standard by a committee vote.

Committee Member: An entity that participates in a JEDEC committee and its Affiliates.

Contribution: Any submitted material, proposal or other submission offered by a Committee Member in the process of developing a Standard for the purpose of incorporating such submitted material, proposal, or other submission into a Standard, provided that such submitted material, proposal, or other submission either, a) exists in a tangible form of expression (including in electronic media); or b) is a verbal statement that is memorialized in written documentation (such as a proposed Standard or submitted material) and is either confirmed or not objected to by such Committee Member within thirty (30) calendar days after its receipt of or access to such memorialization, including without limitation, via posting on the JEDEC website.

8.2.1 Terms and definitions (cont'd)

Disclosure of Potentially Essential Patents (to Disclose Potentially Essential Patents): Disclosure to JEDEC, in writing, in the following form:

- a) for issued patents, disclose the patent owner, assignee (if available at time of disclosure) name, address and intellectual property rights (“IPR”) contact person; name or title of the patent; the patent number; the country or countries of issuance, and the Standard(s) on which the submitter of the disclosure believes an Essential Patent Claim may read,
- b) for published applications, disclose the patent application owner, assignee (if available at time of disclosure) name, address and IPR contact person; name or title of the patent; the patent application number; the country or countries of filing, and the Standard(s) on which the submitter of the disclosure believes an Essential Patent Claim may read, and
- c) for unpublished pending applications, disclose the name or title of the patent application, the patent application owner or assignee (if available at time of disclosure) name, address and the IPR contact person; the patent application number; the general subject matter of such unpublished application(s) and the Standard(s) on which the submitter of the disclosure believes an Essential Patent claim may read. Nothing herein precludes broader disclosure of unpublished patent applications on a voluntary basis or pursuant to a non-disclosure agreement.

Disclosure of a patent is deemed to include all patents claiming priority of a single filing. The listing of foreign equivalents is optional and at the discretion of the patent holder.

NOTE If any of a Committee Member’s disclosed patent or patent application contains Essential Patent Claims which, if licensed, would require a payment of royalties or other material consideration to an unaffiliated third party, the Committee Member must also note this fact in its disclosure statement.

Essential Patent: A Patent containing one or more Essential Patent Claims.

Essential Patent Claims: Those Patent claims the use of which would necessarily be infringed by the use, sale, offer for sale or other disposition of a portion of a product in order to be compliant with the required portions of a final approved JEDEC Standard.

NOTE Essential Patent Claims do not include Patent claims covering aspects that are not required to comply with a JEDEC Standard, or are required only for compliance with sections that are marked “example,” “non-normative,” or otherwise indicated as not being required for compliance, or related to underlying enabling technologies or manufacturing techniques not specified in the standard.

Participate (Participated, Participation): (a) joining a JEDEC committee or task group involved in the development of or amendments to a Standard; or (b) attending more than one meeting of such committee or task group (in person or by telephone) within a one-year period. “Participate” includes being present at any point in a meeting of a committee or task group even if the individual or entity has not formally joined the committee or task group.

Patents: All classes or types of patents other than design patents (including, without limitation, originals, divisions, continuations, continuations-in-part, extensions or reissues), and applications for these classes or types of patents throughout the world.

Potentially Essential Patent: A Patent that is reasonably believed by a subject person to contain one or more Essential Patent Claims.

8.2.1 Terms and definitions (cont'd)

Potential Licensees: All JEDEC Committee Members and non-members.

RAND: Reasonable and non-discriminatory licensing terms and conditions.

Representative: An individual who represents a Committee Member. The term includes entities that have not formally joined the committee or task group, if they participate.

Standard: Standards include publications and package outlines, provided that they are adopted (or intended to be adopted) by the JEDEC Board of Directors and have the effect of Standards. Standards work includes activities in JEDEC committees and task groups.

NOTE Standards include both proposed and adopted JEDEC technical specifications.

8.2.2 General Provisions

8.2.2.1 Committee Members

All Committee Members, as a condition of committee membership or committee Participation, agree to abide by JEDEC rules and procedures, including this JEDEC patent policy ("Patent Policy"). This commitment also applies to entities that Participate in JEDEC Committees or task groups, whether or not they have formally joined as JEDEC members.

All Committee Members, as a condition of committee membership and participation, agree to Disclose Potentially Essential Patents, as set forth more fully in 8.2.3, and to offer to license their Essential Patent Claims to all Potential Licensees on RAND terms and conditions, if and as consistent with 8.2.3 and 8.2.4.

The disclosure and licensing obligations of Committee Members and Participants are limited to Standards developed in the particular JEDEC committees in which and while they Participate.

As JEDEC members, Committee Members and Representatives are encouraged, but not required, both to disclose their Potentially Essential Patents and to license their Essential Patent Claims on RAND terms and conditions for Standards developed in all other JEDEC committees in which they are not Committee Members. Such disclosure shall be in the format described in 8.2.3 and be made to the JEDEC Legal Department.

Committee Members and Representatives are encouraged, but not required, to disclose the Potentially Essential Patents of other entities, including other Committee Members and third-parties.

Committee Members and Representatives are encouraged, but not required, to bring Patent Policy issues or concerns with respect to the Disclosure of Potentially Essential Patents and the licensing of Essential Patent Claims to the attention of the JEDEC Board of Directors for resolution, see 9.2.

8.2.2 General Provisions (cont'd)

8.2.2.2 Committees

JEDEC committees and task groups shall not consider any matters or proposals that are outside the approved scope of such committees or task groups. Furthermore, JEDEC committees and task groups are encouraged to use best efforts to avoid proposals for Standards that are within the approved scope of such committees and task groups and that would be likely to cause Committee Members to terminate their membership in the committee or task group for the purpose of discontinuing their disclosure or RAND licensing obligations.

JEDEC committees and task groups shall keep track of all notices, disclosures and licensing assurances with respect to Potentially Essential Patents and to review, on a periodic basis, the list of prior disclosures of Potentially Essential Patent Claims.

JEDEC committees and task groups shall consider reasonable workarounds and technical alternatives (but are not required to implement such workarounds or alternatives) within the earlier of: a) balloting, or b) one-hundred twenty (120) calendar days of receiving Notice of Refusal to offer Licenses on RAND Terms.

JEDEC committees and task groups are encouraged, but not required, to resolve patent policy issues before balloting. Unresolved issues relating to Potentially Essential Patents may be submitted to the JEDEC Board of Directors for resolution, see 9.2. The sole remedy for any issues or concerns about deviations from approved scope shall be an appeal to the JEDEC Board of Directors made prior to Approval by the Committee, see 9.2.

8.2.3 Disclosure of Potentially Essential Patents

At each committee meeting, the chairperson should call to the attention of all those present the requirements contained in the JEDEC Legal Guides and the obligation of all Representatives to inform the committee of any personal knowledge they have of any Potentially Essential Patents that are owned or controlled by that Committee Member and to call for the Disclosure of Potentially Essential Patents by Representatives. Annex A provides information to be displayed at the beginning of the committee meeting to satisfy the requirement. Additionally, all Representatives should be asked to read the statement attached to each JEDEC sign-in/attendance roster; see Annex A for patent policy application guidelines.

All Committee Members must Disclose Potentially Essential Patents, known to their Representative(s) to be Potentially Essential Patents that are owned or controlled by that Committee Member to the personal knowledge of the Representatives. However, neither Committee Members nor Representatives shall have any obligation to conduct a search for Potentially Essential Patents.

Disclosure of Potentially Essential Patents by a Representative or the Committee Member being represented shall be made as early as reasonably possible. The Disclosure of Potentially Essential Patents shall be in accordance with the definition of Disclosure of Potentially Essential Patents, see 8.2.1. Initial disclosure by a Committee Member or Representative may be made in a meeting of the committee or task group. The Representative is responsible for ensuring that such disclosure is properly recorded in the meeting minutes.

8.2.3 Disclosure of Potentially Essential Patents (cont'd)

The Representative or Committee Member shall document all known Potentially Essential Patents in either: a) License Assurance/Disclosure Forms, see Annex A.3, or b) Notice of Refusal to offer Licenses on RAND Terms forms, see Annex A.4. Such disclosures or notices of known Potentially Essential Patents shall be delivered to the JEDEC Legal Department within thirty (30) calendar days of Approval by the Committee in order to be effective.

8.2.3.1 Disclosure of unwillingness to license work of a committee

If a Committee Member knows that it would not be willing to offer to license Essential Patent Claims arising from the Contributions of another Committee Member on RAND terms to all Potential Licensees consistent with 8.2.4, the Committee Member shall notify the committee chair in writing as early as reasonably possible but no later than thirty (30) calendar days after Approval by the Committee in the form of a Notice of Refusal to offer Licenses on RAND Terms, see Annex A.4.

The committee shall consider working around the material or other technical alternatives with a goal of resolving the concern if at all possible. The Committee Member may (at its own discretion) withdraw its notification of Refusal to License at any time. However, if the Committee Member wishes to maintain its position that it would refuse to License the work of the committee, the Committee Member must withdraw from the committee no later than one-hundred twenty (120) calendar days after providing notice of its unwillingness to license.

8.2.4 RAND Patent Licensing Commitment

Subject to the terms and conditions of section 8.2.4, each Committee Member, as a condition of Participation, agrees to offer to license on RAND terms, to all Potential Licensees, such Committee Member's Essential Patent Claims for the use, sale, offer for sale or other disposition of a portion of a product in order to be compliant with the required portions of a final approved JEDEC Standard issued during the period of membership in that committee. The licensing commitment does not apply to Essential Patents of a Committee Member where notice of a Refusal to License has been given by the Committee Member in accordance with 8.2.3.1.

This commitment applies to any Standard that was pending in the committee or task group while the Committee Member Participated in that committee or task group. If a Committee Member terminates its committee membership, the commitment does not apply with respect to any new proposal or Standard introduced in the committee or task group after the Committee Member provided notice to the JEDEC Legal Department terminating its committee membership. If and as necessary, questions whether or not a particular proposal is a new Standard or merely a continuation of a prior Standard will be addressed in the first instance by the JEDEC Legal Department in accordance with 8.2.10.

Certain exceptions apply to the licensing requirement. A Committee Member will not be required to license or continue to license its Essential Patent Claims to a Potential Licensee for a particular Standard if: (i) such Potential Licensee does not agree to grant a license to all other Potential Licensees under such Potential Licensee's Essential Patent Claims of that Standard on RAND terms and conditions for the approved Standard and/or (ii) such Potential Licensee has commenced or has threatened to commence patent litigation targeting such Committee Member's products that are meant to comply with that Standard.

8.2.5 Licensing Assurance/Disclosure Form

A Committee Member or Representative agrees to offer the following licensing options for an Essential Patent Claim:

For any Essential Patent Claims held or controlled by the entity, pending or anticipated to be filed, which are or may be required to implement a Standard that may result from the JEDEC Standard Activity, the entity hereby makes one of the following commitments:

- (i) A license will be offered, without compensation, under reasonable terms and conditions that are free of any unfair discrimination to applicants desiring to utilize the license for the purpose of implementing the JEDEC Standard, subject to the terms and conditions in 8.2.4; *or*
- (ii) A license will be offered, to applicants desiring to utilize the license for the purpose of implementing the JEDEC Standard under reasonable terms and conditions that are free of any unfair discrimination, subject to the terms and conditions in 8.2.4.

Such assurance shall indicate that the patent holder (or third party authorized to make assurances on its behalf) will include in any documents transferring ownership of patents subject to the assurance, provisions sufficient to ensure (i) that the commitments in the assurance are binding on the transferee⁷, and (ii) that the transferee will similarly include appropriate provisions in the event of future transfers with the goal of binding each successor-in-interest.

The assurance shall also indicate that it is intended to be binding on successors-in-interest regardless of whether such provisions are included in the relevant transfer documents.

If a Committee Member, at its discretion, elects not to submit a License Assurance/Disclosure Form (see Annex A.3) at or before the time the ballot closes and does not otherwise provide notice of an unwillingness to license in accordance with 8.2.3.1, the Committee Member and its Affiliates will be deemed to have agreed to offer to grant licenses on RAND terms and conditions for all of its Essential Patent Claims of the balloted Standard, if and as consistent with 8.2.4.

8.2.6 Reference to patented products in JEDEC standards and publications

The following notice shall be included inside the front cover of all JEDEC Standards in which Essential Patent Claims are or may be involved:

“The user’s attention is called to the possibility that compliance with this document may implicate the use of an invention covered by patent rights.”

“By publication of this document, no position is taken with respect to the infringement or validity of any patents. However, JEDEC has received statements of a willingness to offer to license certain rights on reasonable and nondiscriminatory terms and conditions to applicants desiring to obtain such a license. Details may be obtained from JEDEC.”

⁷ Committee Members and Representatives may comply with (i) and (ii) by notifying the transferee in writing that the RAND licensing obligations under Section 8.2.4 are binding on the Committee Member or Representative, the transferee and subsequent transferees or a general statement in the transfer or assignment agreement that the Patent Claim(s) being transferred or assigned are subject to any encumbrances, e.g., a RAND licensing obligation, that may exist as of the effective date of such agreement.

8.2.7 Special legal disclaimer

The Board has the discretion to approve the issuance of a standard for which a patent owner or applicant has not provided written assurance that the Essential Patent Claims will be licensed, subject to special legal disclaimers.

When determining whether to approve the issuance of such a standard, the Board shall consider whether the committee used diligent efforts, if appropriate under the circumstances, to develop a standard that does not require the use of the Essential Patent Claim(s).

A special legal disclaimer might include the following language:

“Notice: JEDEC has received information that certain patents or patent applications may be essential to this standard. However, as of the publication date of this standard, no statements regarding an assurance or refusal to license such patents or patent applications have been provided. JEDEC does not make any determination as to the validity or relevancy of such patents or patent applications. Anyone making use of the standard assumes all liability resulting from such use. JEDEC and its members disclaim any representation or warranty, express or implied, relating to the standard and its use.”

If the Board determines that:

- (1) the claim by the patent owner or applicant that an identified patent or patent application may be essential to implement a Standard is a reasonable claim, and
- (2) there is a credible indication that the patent owner or applicant is unwilling or unable to grant licenses, with or without compensation, on RAND terms,

then the Board *shall not* approve the issuance of the standard except as provided in 8.2.7. In its discretion, the Board may seek the advice of an independent outside patent attorney or consultant in making the determinations in subsections (1) and/or (2) of 8.2.7.

8.2.8 Miscellaneous

The Patent Policy applies equally to situations involving the Essential Patent Claims that are discovered after adoption of the Standard.

JEDEC makes no representation as to the reasonableness of any terms or conditions of the license agreements offered by such patent rights holders, and all negotiations regarding such terms and conditions must take place between the individual parties outside the context of JEDEC.

To the extent that Potentially Essential Patent Claims are disclosed with respect to a JEDEC Standard, information regarding the parties who have asserted patent rights in the Standard may be obtained from the JEDEC Corporate Office.

It is possible that other patent rights that have not been disclosed to JEDEC may be implicated by implementation of or compliance with a JEDEC Standard. JEDEC is not responsible for identifying patent claims for which a license may be required in order to implement a JEDEC Standard, and takes no position and conducts no inquiries with respect to the legal validity or scope of those patents or patent claims that are brought to its attention.

8.2.9 Violations

If a Committee Member or Representative materially violates this Patent Policy and the violation is not cured within thirty (30) days following notification of the violation by the JEDEC Corporate Office, the JEDEC Board of Directors may in its discretion suspend the Committee Member or Representative's right to participate in JEDEC. If the Committee Member is allowed to continue participation and engages in another material violation of this Policy, the JEDEC Board of Directors may in its discretion revoke the Committee Member or Representative's right to participate in JEDEC.

8.2.10 Interpretation and Governing Law

The obligations and rights of parties to this agreement in the conduct of committee meetings and the approval of standards as set forth in this Patent Policy will be interpreted in the first instance by the JEDEC Legal Department. When specifically requested by a Committee Member or Representative, interpretations by the Legal Department will be in writing.

Written interpretations by the Legal Department can be appealed by a Committee Member or Representative to the JEDEC Board of Directors. The Board of Directors will make a final decision on appeal, as determined in their reasonable and good faith judgment. The final decision will be in writing.

In no event will JEDEC or the JEDEC Board make a determination that a Patent is or is not a Potentially Essential Patent or that a Patent Claim is or is not an Essential Patent Claim.

The Patent Policy will be interpreted and governed under the laws of the State of New York.

9 Appeals

All appeals related to standard or publication development are initially directed to the JEDEC Office. The provision for appeals is important for the protection of all parties involved within the JEDEC standards development process. An appeals mechanism safeguards both those who are directly and or materially affected by the process and the standards developing organization.

9.1 Appeals related to ANSI/JEDEC standards and publications

The following appeals procedures shall be used for actions taken within the JEDEC standards process for the development of ANSI/JEDEC standards and publications.

9.1.1 Right of appeal

Persons who have directly and or materially affected interests and who have been or will be adversely affected by a standard or publication within ANSI jurisdiction have the right to appeal substantive or procedural actions or inactions of JEDEC formulating groups.

The burden of proof to show an adverse effect shall be on the appellant. Appeals shall be directed to the JEDEC Office in accordance with the appeals procedures as defined in this section.

9.1.2 Criteria for appeals mechanism

The following general criteria are encompassed in these appeals procedures:

- a) Appeals shall be addressed promptly and a decision made expeditiously;
- b) The right of the involved parties to present their cases shall not be denied;
- c) These procedures shall provide for participation by all parties concerned without imposing an undue burden on them;
- d) Consideration of appeals shall be fair and unbiased and shall fully address the concerns expressed; and
- e) Records of appeals shall be kept and made available upon request. The JEDEC Office may levy an appropriate charge to cover the cost of reproduction, handling, and distribution of materials. The fee shall be predetermined, fixed, and reasonable. A procedure for requesting a fee waiver or fee reduction shall be available.

9.1.3 Appeals procedures

The following subsections outline the JEDEC Appeals Procedures. In the case of an appeal to ANSI of a JEDEC action or inaction on an ANSI/JEDEC standard, the appellant may waive the right of the JEDEC appeals process, see 9.1.4.

9.1.3.1 Complaints

The appellant shall file a written complaint with the JEDEC Office within 30 calendar days after the date of the action being appealed or at any time with respect to inaction. The complaint shall state the nature of the objection(s) including any adverse effects, the section(s) of these procedures or the standard(s) that are at issue, action(s) or instances of inaction that are at issue, and the specific remedial action(s) that would satisfy the appellant's concerns. Previous efforts to resolve the objection(s) and the outcome of each shall be noted.

9.1.3.2 Response

Within 30 calendar days after receipt of the complaint, the respondent (formulating group chairman or JEDEC representative) shall respond in writing to the appellant, specifically addressing each allegation of fact in the complaint to the extent of the respondent's knowledge. Complaints that are technical in nature shall be referred to the appropriate JEDEC formulating committee having jurisdiction over the material for review and possible action. Complaints that are procedural in nature shall be referred to the JEDEC Board for review and possible action.

9.1.3.3 Hearing

If the appellant and the respondent are unable to resolve the written complaint informally in a manner consistent with these procedures within 15 calendar days, the JEDEC Office shall schedule a hearing with an appeals panel* on a date agreeable to all participants but within 30 calendar days, giving at least 10 calendar days' notice. Hearing on the complaint will take place at the next regular meeting of the JEDEC formulating committee or the JEDEC Board depending on the nature of the complaint.

NOTE In the case of an ANSI appeal, see Annex D.10 for clarification of an appeals panel.

9.1.3.4 Conduct of the hearing

The appellant has the burden of demonstrating adverse effects, improper actions or inactions, and the efficacy of the requested remedial action.

9.1.3.5 Decision

The appeal disposition shall be documented by the formulating committee for technical matters and by the JEDEC Board for procedural matters. A record of the appeals shall be kept by JEDEC and made available to the involved parties.

NOTE In the case of an ANSI appeal, the appeals panel shall render its decision in writing within 30 days of the hearing, based upon a preponderance of the evidence, stating its findings of fact and conclusions, with reasons therefore and citing the evidence. The Committee Secretary shall notify the appellant and the Committee of the decision of the appeals panel, which shall be binding and final on all concerned.

If, for whatever reason(s), the appellant chooses to forego the process detailed above, the JEDEC Office shall maintain all documents pertaining to JEDEC attempts to resolve the matter at hand. These documents may be used in any further appeals proceedings.

9.1.4 Further appeal

If the matter under appeal relates to a JEDEC Standard that has been or is expected to be recognized as an American National Standard, subsequent further appeal may be made directly to ANSI. If the appellant gives notice to the JEDEC Office that further appeal to ANSI is intended, a full record of the complaint, response and decision shall be submitted by the JEDEC Office to ANSI.

NOTE ANSI will not normally hear an appeal of an action or inaction by JEDEC until the appeals process has been completed. Such appeals shall be directed to ANSI in accordance with the procedures of the appropriate ANSI entity (e.g., Board of Standards Review, Executive Standards Council, etc.).

9.2 Appeals not related to ANSI/JEDEC standards and publications

The following appeals procedures shall be used for actions taken within the JEDEC standards process that are not related to the development of ANSI/JEDEC standards and publications.

9.2.1 Right of appeal

Committee Members or Participants who have directly and or materially affected interests and who have been or will be adversely affected by the activities of another Committee Member or Representative, a JEDEC committee, or a JEDEC standard or publication have the right to appeal technical or procedural actions or inactions.

The burden of proof to show an adverse effect shall be on the appellant. Appeals shall be directed to the JEDEC Office in accordance with the appeals procedures as defined in this section.

9.2.2 Criteria for appeals mechanism

The following general criteria are encompassed in these appeals procedures:

- a) Appeals shall be addressed promptly and a decision made expeditiously;
- b) The right of the involved parties to present their cases shall not be denied;
- c) These procedures shall provide for participation by all parties concerned without imposing an undue burden on them;
- d) Consideration of appeals shall be fair and unbiased and shall fully address the concerns expressed; and
- e) Records of appeals shall be kept and made available upon request. The JEDEC Office may levy an appropriate charge to cover the cost of reproduction, handling, and distribution of materials.

9.2.3 Appeals procedures

The following subsections outline the JEDEC Appeals Procedures for appeals that are not related to ANSI/JEDEC standards and publications.

9.2.3.1 Complaints

The appellant shall file a written complaint with the JEDEC Office within 30 calendar days after the date of the action being appealed or at any time with respect to inaction. The complaint shall state the nature of the objection(s) including any adverse effects, the section(s) of these procedures or the Standard(s) that are at issue, action(s) or instances of inaction that are at issue, and the specific remedial action(s) that would satisfy the appellant's concerns. Previous efforts to resolve the objection(s) and the outcome of each shall be noted.

9.2.3.2 Response

Within 30 calendar days after receipt of the complaint, the respondent(s) shall respond in writing to the appellant, specifically addressing each allegation of fact in the complaint to the extent of the respondent's knowledge.

Complaints that are technical in nature shall be referred to the appropriate JEDEC committee having jurisdiction over the material for review and possible action. Complaints that are non-technical or procedural in nature shall be referred to the JEDEC Board for review and possible action. Complaints that relate to the patent policy shall be referred to the JEDEC Board for review and possible action.

9.2.3.3 Hearing

If the appellant and the respondent are unable to resolve the written complaint informally in a manner consistent with these procedures within 30 calendar days, the JEDEC Office shall schedule a hearing with an impartial appeals panel consisting of three individuals appointed by the committee or the JEDEC Board, as appropriate, on a date agreeable to all parties, giving at least 10 calendar days' notice. Hearing on the complaint will take place in person or by telephone.

9.2.3 Appeals procedures (cont'd)

9.2.3.4 Conduct of the hearing

The appellant has the burden of demonstrating adverse effects, improper actions or inactions, and the efficacy of the requested remedial action.

9.2.3.5 Decision

The appeal disposition shall be documented by the JEDEC appeals panel. A record of the appeals shall be kept by JEDEC and made available to the involved parties.

Appeal dispositions on complaints that are technical in nature may be further appealed to the appropriate JEDEC committee. Appeal dispositions on complaints that are non-technical or procedural in nature may be further appealed to the JEDEC Board of Directors.

10 External publicity

All articles, presentations, news releases, or other publicity must be approved by JEDEC Legal Counsel prior to release.

Articles, speeches, and other presentations for publication or delivery by members of the Board, its committees, or JEDEC staff that cover specific JEDEC work or JEDEC policy or purports to speak on behalf of JEDEC must be approved by the Board, if time permits; otherwise by the Board chairperson who shall be responsible for ensuring that the content is accurate and reflects current JEDEC policy. Any publicity that speaks for the Board shall require similar approval.

The committee chairperson may talk to the press and approve technical and tradeshow presentations about ongoing committee work when it is appropriate to promote a particular standard or the work of the committee. This public disclosure should be limited to the goal of the standard under development and any details of the standard that have been approved by committee ballot(s).

Companies may not advertise or give press releases stating that their products comply with a JEDEC standard, publication, or registration prior to Board approval of the documents.

11 JEDEC ethics rules for electronic communications

No JEDEC member or alternate shall post to an e-mail reflector, bulletin board, or Web site owned or sponsored by JEDEC, any of the following:

- a) any advertising or marketing information,
- b) any cost or pricing information,
- c) any proprietary information,
- d) any non-JEDEC copyright information or material that is in excess of that allowed under current USA copyright laws,
- e) any technical information that requires a USA government export license before it is transferred outside of the United States,
- f) information covered by USA secrecy laws,
- g) solicitations, or requests for contributions to “for profit” or “non-profit” (charitable) organizations,
- h) chain letters,
- i) job search or hiring information,
- j) computer games or other software products,
- k) files with known software viruses,
- l) offensive material covering:
 - 1) libel (i.e., false or malicious statements),
 - 2) threats, retaliation, or intimidation,
 - 3) use of language that is obscene or in poor taste, including jokes,
 - 4) sexual harassment or inappropriate sexual content, such as pictures, images, messages, or cartoons,
 - 5) discriminatory communication relating to race, color, national origin, sex, marital status, sexual orientation, age, veteran’s status, disability, religious or political beliefs.

All electronic communications and posting must comply with this JEDEC Manual and the JEDEC Legal Guides.

Annex A (normative) Legal guidelines summary

A.1 JEDEC patent policy summary

- All JEDEC Member Companies and their Representatives agree to abide by the JEDEC patent policy.
- The commitment applies to all entities, including non-committee members, that have a Representative present at any point in a meeting of a JEDEC Committee or Task Group.
- The commitment continues as long as the Committee Member or Representative is a member of or a participant in a particular JEDEC Committee.
- For the purposes of the patent policy, "Committee" refers to the full committee, except for those committees that have subcommittees. In that context, "committee" refers to any group that has a separate scope.
- All companies and their representatives are required to disclose all known Potentially Essential Patent Claims owned or controlled by the company.
- No company has a duty to conduct a search for Potentially Essential Patent Claims.
- All companies, as a condition of committee membership or participation, agree to license their Essential Patent Claims on RAND terms and conditions.
- The licensing assurance must indicate that the patent holder (or third party authorized to make assurances on its behalf) will include in any documents transferring ownership of patents subject to the assurance, provisions sufficient to ensure (i) that the commitments in the assurance are binding on the transferee, and (ii) that the transferee will similarly include appropriate provisions in the event of future transfers with the goal of binding each successor-in-interest.
- If a Member Company knows it would be unwilling to license its Essential Patent Claims on RAND terms, the member must notify the committee chair and withdraw from the committee within 120 days after giving the notice. This option is not available to a Member Company with respect to its own Contribution.
- Companies and their representatives are required to use a standard form for the disclosures and assurances.
- The disclosure and licensing obligations of companies and their representatives are limited to Standards developed in the particular JEDEC committees in which they are members or in which they participate.
- The obligations apply to all entities, including non-committee members, that have a representative present at any point in a meeting of a JEDEC Committee or Task Group.
- Member Companies and their Representatives are strongly encouraged to review their current committee and subcommittee assignments, and ensure they are prepared to accept the commitments defined by the patent policy.

See, 8.2 for additional information.

A.2 Copyright material

If, in developing a standard, a publication, or a specification, a committee proposes to incorporate verbatim material from a publication copyrighted by another organization, it shall obtain written permission from the owner of the copyright for JEDEC to copy the material. The letter granting permission shall be sent to the JEDEC Office, and the standard or publication shall include a reference to such permission as a footnote.

JEDEC Manual No. 21T
Page 36

A.3 License Assurance/Disclosure Form (See 8.2.3 and 8.2.5 of the JEDEC Patent Policy)

For the electronic version of the License Assurance/Disclosure Form go to: [http://www.jedec.org/sites/default/files/License Assurance-Disclosure Form 20150710.pdf](http://www.jedec.org/sites/default/files/License_Assurance-Disclosure_Form_20150710.pdf)

Part 1

Date: _____

Entity Name: _____

Address: _____

IPR Contact: _____

Phone: _____

E-Mail: _____

Part 2

Identify the relevant JEDEC Standard:

Part 3

For Issued Patents:

Patent No: _____

Patent Name or Title: _____

Country of Issuance: _____

For Published Patent Applications:

Published Patent Application No. : _____

Patent Application Name or Title: _____

Country of Filing: _____

For Unpublished Patent Applications:

Subject Matter of Patent Application:

A.3 License Assurance/Disclosure Form (cont'd)

You must complete Part 4 if the entity holds a Patent or has applied for a Patent on an invention the use of which is or may be required to comply with a Standard that may result from the JEDEC Standard Activity.

Part 4

For any Essential Patent Claims held or controlled by the entity, pending or anticipated to be filed, the entity states:

- ___ (i) A license will be offered, without compensation, under reasonable terms and conditions that are demonstrably free of any unfair discrimination to applicants desiring to utilize the license for the purpose of implementing the JEDEC Standard; *or*
- ___ (ii) A license will be offered, with compensation, to applicants desiring to utilize the license for the purpose of implementing the JEDEC Standard under reasonable terms and conditions that are demonstrably free of any unfair discrimination.

Such assurance shall indicate that the patent holder (or third party authorized to make assurances on its behalf) will include in any documents transferring ownership of patents subject to the assurance, provisions sufficient to ensure (i) that the commitments in the assurance are binding on the transferee⁸, and (ii) that the transferee will similarly include appropriate provisions in the event of future transfers with the goal of binding each successor-in-interest.

The assurance shall also indicate that it is intended to be binding on successors-in-interest regardless of whether such provisions are included in the relevant transfer documents.

⁸ Committee Members and Representatives may comply with (i) and (ii) by notifying the transferee in writing that the RAND licensing obligations under Section 8.2.4 of the JEDEC Patent Policy are binding on the Committee Member or Representative, the transferee and subsequent transferees or a general statement in the transfer or assignment agreement that the Patent Claim(s) being transferred or assigned are subject to any encumbrances, e.g., a RAND licensing obligation, that may exist as of the effective date of such agreement.

JEDEC Manual No. 21T
Page 38

A.3 License Assurance/Disclosure Form (cont'd)

Part 5

For any disclosed Patent or Patent Application that contains Essential Patent Claims which, if licensed, would require a payment of royalties or other material consideration to an unaffiliated third party, provide the following information:

Name of Third Party: _____

Address: _____

IPR Contact: _____

Phone: _____

E-Mail: _____

AGREED, ON BEHALF OF THE ENTITY:

(Signature)

(Name printed)

(Date)

NOTE The committee or task group has the right to request additional technical information relating to the Patent or Patent Application in order to consider workarounds and other technical alternatives.

RETURN THE COMPLETED FORM BY MAIL TO
JEDEC, ATTENTION: LEGAL DEPARTMENT,
3103 NORTH 10TH STREET, SUITE 240-S, ARLINGTON, VA 22201-2107
OR BY E-MAIL TO JOHNK@JEDEC.ORG

A.4 Notice of Refusal to Offer Licenses on RAND Terms

Part 1

Date: _____

Committee Member: _____

Entity Mailing Address: _____

IPR Contact Name: _____

IPR Contact Telephone Number: _____

IPR Contact E-Mail Address: _____

Part 2

Identify the relevant JEDEC Standard:

Part 3

Pursuant to 8.2.3.1 of the JEDEC Patent Policy, the Committee Member hereby provides notice that the Committee Member is not willing to offer a license on RAND terms to the Potentially Essential Patents identified in A.4, Part 4.

Part 4

For Issued Patents (please attach additional pages as necessary):

1. Owner or assignee (if known at time of disclosure) name: _____
2. Patent Name or Title: _____
3. Patent Number: _____
4. Country of Issuance: _____
5. Portions of Proposed JEDEC Standard affected (please use reasonable efforts to identify which portions may be affected), e.g., clause, page, or any other identifiable information:

For Published Patent Applications (please attach additional pages as necessary):

1. Owner or assignee (if known at time of disclosure) name: _____
2. Published Patent Application Number: _____
3. Patent Application Name or Title: _____
4. Country of Filing: _____
5. Portions of Proposed JEDEC Standard affected (please use reasonable efforts to identify which portions may be affected), e.g., clause, page, or any other identifiable information:

JEDEC Manual No. 21T
Page 40

A.4 Notice of Refusal to Offer Licenses on RAND Terms (cont'd)

For Unpublished Patent Applications (please attach additional pages as necessary):

1. Owner or assignee (if known at time of disclosure) name: _____
2. General Subject Matter of Patent Application: _____

3. Portions of Proposed JEDEC Standard affected (please use reasonable efforts to identify which portions may be affected), e.g., clause, page, or any other identifiable information: _____

Part 5

By signing this Notice, the IPR Contact represents that they are authorized to complete and submit this Notice on behalf of the Committee Member. Pursuant to Section 8.2.3.1 of the JEDEC Patent Policy, the IPR Representative represents that the information provided in this Notice does not include a contribution that was contributed by the Committee Member. If the Committee Member becomes aware of any other Potentially Essential Patent not already covered by a Notice and the ballot period is open, then the Committee Member may submit additional Notices.

Signature of IPR Contact

Accepted and Agreed to:

JEDEC

Date Accepted: _____

*RETURN THE COMPLETED NOTICE BY MAIL TO
JEDEC, ATTENTION: LEGAL DEPARTMENT,
3103 NORTH 10TH STREET, SUITE 240-S, ARLINGTON, VA 22201-2107
OR BY E-MAIL TO JOHNK@JEDEC.ORG*

Annex B (informative) Patent policy application guidelines

The following points describe the application of the JEDEC patent policy:

- B.1** Committee discussion of pending or existing patents is a permissible activity and is encouraged when the committee feels that the patented item or process represents the best technical basis for a standard.
- B.2** Discussion of a pending or existing patent does not constitute an acknowledgment of the validity of the patent, because validity is based on prior art and determination of the company/individual that first made the invention or applied for the patent. The committee's concern is only with the technical merits of the patent and whether the technical proposal is a sound basis for standardization.
- B.3** By its terms, the JEDEC patent policy applies with equal force to situations involving:
 - a) the discovery of patents that may be required for use of a standard subsequent to its adoption, and
 - b) the initial issuance of a patent after the adoption of a standard.

Once disclosure is made, the holder is obligated to provide the same assurances to JEDEC as are required in situations where patents exist or are known prior to approval of a proposed standard.

Thus, if notice is given of a patent that may be required for use of an already approved JEDEC standard, a standards developer may wish to make it clear to other standards-making participants that the JEDEC procedures require the patent holder to provide the assurances contained in the patent policy or suffer the withdrawal of JEDEC's approval of the document as a JEDEC standard.

Annex C (normative) ANSI Standardization Procedures

For those documents that are intended for publication as an American National Standard (ANS), the provisions set forth in this annex apply.

C.1 Notification of American National Standards (ANS) development

Notification of American National Standards activity shall be announced in suitable media as appropriate to demonstrate provision of opportunity for participation by all directly and materially affected persons. At the initiation of a project to develop or revise a Standard, notification shall be transmitted to ANSI using the Project Initiation Notification System (PINS) form, or its equivalent, for listing in Standards Action. A PINS form may be submitted, but is not required, at the initiation of a project to reaffirm or withdraw a Standard. Prompt consideration shall be given to the written views and objections of all participants, including those commenting on the PINS announcement. Comments resulting from the filing of PINS will be handled in accordance with the ANSI Essential Requirements, 2.5.

C.2 Submission to ANSI for public review

Following approval of the standard, responsible JEDEC Standards staff shall submit the proposed ANS to ANSI to initiate public review. JEDEC Standards staff shall provide appropriate documentation to ANSI. Responsible JEDEC Standards staff shall receive, collect and acknowledge receipt of public review comments. In addition, responsible JEDEC Standards staff shall forward all public review comments to the formulating group for review and response. The response to each commenter shall include a notation of the JEDEC and ANSI appeals process.

NOTE A proposed ANS is subject to comment only (not for vote) during public review and by any and all materially or directly affected parties.

C.3 Public review comments

JEDEC Standards staff shall acknowledge receipt of a public review comment, regardless of its nature. The only comments normally addressed by a working group or formulating group are editorial or technical in nature. Technical comments should be accompanied with wording that would aid the formulating group in addressing the comments.

Other comments may be handled in the following manner on behalf of the formulating group:

- The JEDEC Legal Counsel shall respond to comments relating to legal concerns;
- JEDEC Standards staff shall respond to comments relating to concerns regarding the standards development process;
- The formulating group with input from JEDEC staff shall address scope concerns;
- The reason for non-acceptance shall be sent regarding editorial or technical concerns;
- Vague comments shall be returned to the respondent requesting expeditious clarification.

JEDEC Standards staff shall retain copies of response to comments.

Annex C (informative) ANSI Standardization Procedures (cont'd)

C.4 Public review commenter response

The commenter shall have 14 calendar days after response to acknowledge the acceptance or rejection of that reply. If no response is received, the responsible JEDEC Standards staff shall send a reminder notice to the commenter requesting response. If no response is received within an additional 5 calendar days, the responsible JEDEC Standards staff will send a registered letter to the commenter noting the lack of response. The commenter's comments are then recorded as unresolved and circulated to the formulating group per the provisions of C.3.

All correspondence showing the formulating group's effort to address the comments, noting a lack of response when appropriate, shall be included in any package of materials provided to ANSI.

C.5 Results scenarios

Results scenario #1: If no comments are received after the close of public review, the document shall be forwarded to ANSI with a request for approval.

Results scenario #2: If editorial comments are received by the close of public review, the document editor, editing group, or chair shall address the comments. The document, with edits, if any, shall be forwarded to ANSI for approval.

Results scenario #3: If substantive comments are received from public review, the working group or formulating group shall review these comments, and attempt to resolve them, and notify each commenter of said action (see C.4), along with a notification of the JEDEC and ANSI appeals process.

- a) If no substantive changes are made, the document is forwarded to ANSI for approval.
- b) If substantive changes are made, the formulating group shall conduct a recirculation vote/ballot and "second" public review (see C.4).
- c) If negative, substantive comments remain unresolved from public review, the formulating group shall conduct a recirculation vote/ballot, and, if necessary, a "second" public review (see C.4).

Annex C (informative) ANSI Standardization Procedures (cont'd)

C.6 Voting on an American National Standard

When a vote in a formulating group is for the purpose of approving a document for publication as an American National Standard, a Ballot Summary (Tally of votes) shall be prepared and included in the minutes. The Interest Category is normally self-selected by the voting member.

Action Required	Method (all calendar days)	Approval Criteria
Meeting Vote*	At the meeting	Two-thirds voting
Ballot	30 day ballot	Two-thirds voting
ANSI BSR8 (Public review)	Initial 60 day comment period [30 or 60 day subsequent comment period(s)] (or other ANSI designated time period where applicable)	Comments addressed by formulating group

* Voting members not present at the time of the vote shall be given the opportunity to vote on the standard after the meeting. The JEDEC staff shall notify the absent voting members via e-mail of any substantive changes made to the standard subsequent to the 21 day ballot period and advise that they have 7 days to cast their vote via the JEDEC Voting Machine. Any absent voting members that fail to respond shall be listed as No Response on the final submittal for approval to ANSI. Response to this opportunity to vote shall not be considered as attendance at the meeting for purposes of voting maintenance.

C.7 Legal approval

All ANSI/JEDEC standards shall receive JEDEC legal approval before publication.

C.8 ANSI approval

If ANSI/BSR is approved, responsible JEDEC Standards staff is notified and in turn publishes the document.

If ANSI/BSR does not approve, comment resolution issues will be dealt with by the formulating group and process-related issues will be dealt with by JEDEC Standards staff. If the ANSI/BSR comments are resolved and the document is approved, responsible JEDEC Standards staff publishes the document.

If the comments cannot be resolved, the document is returned to the formulating group for further resolution or cancellation. The formulating group may appeal the actions of the ANSI/BSR.

C.9 ANSI patent policy

If a JEDEC standard is intended for an American National Standard, then the standard must comply with the current ANSI patent policy.

Annex C (informative) ANSI Standardization Procedures (cont'd)

C.10 Appeals panel (for an appeal of an ANS)

In the case of an ANSI appeal, the appeals panel shall consist of three individuals who have not been directly involved in the matter in dispute, and who (knowingly in good faith) will not be materially or directly affected by any decision made or to be made in the dispute. At least two members shall be acceptable to the appellant and at least two members shall be acceptable to JEDEC. In the event the appellant does not wish to select a panelist or a third panelist cannot be agreed upon, JEDEC shall appoint this individual to the panel in order to hold a hearing.

C.11 ANSI Commercial Terms and Conditions

Provisions involving business relations between buyer and seller such as guarantees, warranties, and other commercial terms and conditions shall not be included in an ANS.

Annex D (informative) Differences between revisions

This annex briefly describes most of the changes made to entries that appear in this standard, JM21T, compared to its predecessor, JM21S (November 2015). If the change to a concept involves any words added or deleted (excluding deletion of accidentally repeated words), it is included. Some punctuation changes are not included.

Clause	Description of Change
1.4.4	4), Removed “within a calendar year” from 1 st sentence, removed “Participation via teleconference is allowed only once per calendar year” and replaced with “. It is the expectation that Directors attend all Board meetings, at minimum Directors must attend at least one face-to-face Board meeting, in person, each calendar year.” (JCB-19-27)
3.5	Quorum: Changed “20” to “12”. (JCB-19-04)
1.4.3.3	Changed to 1.4.3.4. (JCB-20-14)
1.4.3.3	New subclause added. (JCB-20-14)
8.2.1	New term and definition added for “Approval by the Committee”. (JCB-20-15)
8.2.2.2	4 th paragraph, last sentence, modified as follows: The sole remedy for any issues or concerns about deviations from approved scope shall be an appeal to the JEDEC Board of Directors made prior to ~~approval of the subject draft specification for resolution~~ Approval by the Committee , see 9.2. (JCB-20-15)
8.2.3	4 th paragraph, last sentence, modified as follows: Such disclosures or notices of known Potentially Essential Patents shall be delivered to the JEDEC Legal Department within thirty (30) calendar days of ~~draft specification completion~~ Approval by the Committee in order to be effective. (JCB-20-15)
8.2.3.1	1 st paragraph, modified as follows: ... reasonably possible but no later than thirty (30) calendar days after ~~completion of the draft specification~~ Approval by the Committee in the form of a Notice of Refusal to offer Licenses on RAND Terms, see Annex A.4. (JCB-20-15)

D.1 Differences between JM21S and JM21R (July 2015)

Clause	Description of Change
3.4	2 nd para., 1 st sentence modified as follows: A JEDEC staff member or a designated alternate must ~~be in attendance at~~ participate in every committee meeting, ~~either in person or by electronic means.~~



EXHIBIT 5

License Assurance/Disclosure Form

PART 1

Date: 11/22/10

Entity Name: Netlist, Inc.

Address: 51 Discovery, Suite 150
Irving, TX 75038

IPR Contact: Mario Martinez

Phone: 949 679 0159

E-Mail: mmartinez@netlist.com

PART 2

Identify the relevant JEDEC Standard:

JC42.3 3DS (#1655.19, #1644.96, #1655.1)

JC45 LRDIMM (#2192.62, #2192.43, #2192.58, #2192.57)

PART 3

For Issued Patents:

Patent No.: _____

Patent Name or Title: _____

Country of Issuance: _____

For Published Patent Applications:

Published Patent Application No.: 12/422,853

Patent Application Name or Title: Circuit Providing Load Isolation & Noise Reduction

Country of Filing: United States

For Unpublished Patent Applications:

Subject Matter of Patent Application:

License Assurance/Disclosure Form

You must complete Part 4 if the entity holds a Patent or has applied for a Patent on an invention the use of which is or may be required to comply with a Standard that may result from the JEDEC Standard Activity.

PART 4

For any Essential Patent Claims held or controlled by the entity, pending or anticipated to be filed, the entity states:

- ☐ (i) A license will be offered, without compensation, under reasonable terms and conditions that are demonstrably free of any unfair discrimination to applicants desiring to utilize the license for the purpose of implementing the JEDEC Standard; *or*
- ☒ (ii) A license will be offered to applicants desiring to utilize the license for the purpose of implementing the JEDEC Standard under reasonable terms and conditions that are demonstrably free of any unfair discrimination.

PART 5

For any disclosed Patent or Patent Application that contains Essential Patent Claims which, if licensed, would require a payment of royalties or other material consideration to an unaffiliated third party, provide the following information:

Name of Third Party: _____


Address: _____

IPR Contact: _____

Phone: _____

E-Mail: _____

AGREED, ON BEHALF OF THE ENTITY:


(Signature)

11/30/2010
(Date)

MARIO MARTINEZ
(Name printed)

NOTE: The committee or task group has the right to request additional technical information relating to the Patent or Patent Application in order to consider workarounds and other technical alternatives.

RETURN THE COMPLETED FORM BY MAIL TO
JEDEC, ATTENTION: LEGAL DEPARTMENT,
3103 NORTH 10TH STREET, SUITE 240-S, ARLINGTON, VA 22201-2107
OR BY E-MAIL TO JOHNK@JEDEC.ORG

28/

EXHIBIT 6

License Assurance/Disclosure Form**PART 1**

Date: 9/7/11

Entity Name: Netlist, Inc.

Address: 51 Discovery, Suite 150
Irvine, CA 92618

IPR Contact: Jayesh Bhakta

Phone: 949 435 0025

E-Mail: jbhakta@netlist.com

PART 2

Identify the relevant JEDEC Standard:

JC40

PART 3**For Issued Patents:**

Patent No.: _____

Patent Name or Title: _____

Country of Issuance: _____

For Published Patent Applications:

Published Patent Application No.: 12/422,853

Patent Application Name or Title: circuit providing Load Isolation &

Country of Filing: USA Noise Reduction

For Unpublished Patent Applications:

Subject Matter of Patent Application:

License Assurance/Disclosure Form

You must complete Part 4 if the entity holds a Patent or has applied for a Patent on an invention the use of which is or may be required to comply with a Standard that may result from the JEDEC Standard Activity.

PART 4

For any Essential Patent Claims held or controlled by the entity, pending or anticipated to be filed, the entity states:

- ☐ (i) A license will be offered, without compensation, under reasonable terms and conditions that are demonstrably free of any unfair discrimination to applicants desiring to utilize the license for the purpose of implementing the JEDEC Standard; *or*
- ☒ (ii) A license will be offered to applicants desiring to utilize the license for the purpose of implementing the JEDEC Standard under reasonable terms and conditions that are demonstrably free of any unfair discrimination.

PART 5

For any disclosed Patent or Patent Application that contains Essential Patent Claims which, if licensed, would require a payment of royalties or other material consideration to an unaffiliated third party, provide the following information:

Name of Third Party: _____

Address: _____

IPR Contact: _____

Phone: _____

E-Mail: _____

AGREED, ON BEHALF OF THE ENTITY:

Jayash Bhargava
(Signature)

9/7/2011
(Date)

Chief Module Architect.
(Name printed)

NOTE: The committee or task group has the right to request additional technical information relating to the Patent or Patent Application in order to consider workarounds and other technical alternatives.

RETURN THE COMPLETED FORM BY MAIL TO
JEDEC, ATTENTION: LEGAL DEPARTMENT,
3103 NORTH 10TH STREET, SUITE 240-S, ARLINGTON, VA 22201-2107
OR BY E-MAIL TO JOHNK@JEDEC.ORG

29