

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

SK HYNIX INC., SK HYNIX AMERICA INC., and SK HYNIX MEMORY
SOLUTIONS INC.,

Petitioners,

v.

NETLIST, INC.

Patent Owner

Patent No. 8,874,831

Issued: October 28, 2014

Filed: July 26, 2012

Inventors: Hyun Lee, Chi-She Chen, Jeffrey C. Solomon,
Scott Milton, Jayesh Bhakta

Title: Flash-DRAM Hybrid Memory Module

Inter Partes Review No. IPR2017-00692

**PETITION FOR *INTER PARTES* REVIEW OF U.S. PATENT NO. 8,874,831
UNDER 35 U.S.C. §§ 311-319 AND 37 C.F.R. § 42.1-.80 & 42.100-.123**

TABLE OF CONTENTS

I. COMPLIANCE WITH REQUIREMENTS FOR A PETITION FOR INTER PARTES REVIEW.....1

A. Certification the 831 Patent May Be Contested by Petitioners..... 1

B. Fee for Inter Partes Review (§ 42.15(a))..... 1

C. Mandatory Notices (37 CFR § 42.8(b)) 1

D. Proof of Service (§§ 42.6(e) and 42.105(a)) 2

II. IDENTIFICATION OF CLAIMS BEING CHALLENGED.....3

III. RELEVANT INFORMATION CONCERNING THE CONTESTED PATENT4

A. Effective Filing Date of the 831 Patent..... 4

B. Person of Ordinary Skill in the Art 7

C. The 831 Patent..... 7

 1. Technical Overview 7

 2. Prosecution History 10

D. Construction of Terms Used in the Claims 10

 1. “Bi-Directional Data Transfer Fabric” (All Claims) 11

 2. “Set of Data Ports” (All Claims)..... 12

 3. “Format Data” (Claims 1-6)..... 14

 4. “Operable at a ... Clock Frequency” (Claim 15)..... 14

IV. OVERVIEW OF THE PRINCIPAL PRIOR ART15

A. U.S. Patent Application Publication No. 2010/0110748 to Best (Ex. 1006)..... 15

 1. Prior Art Status..... 15

 2. Overview of Best..... 17

V. PRECISE REASONS FOR RELIEF REQUESTED.....20

A.	Claims 1-14 Are Anticipated by Best	20
1.	Claim 1	20
	a) Preamble	20
	b) “A Non-Volatile Memory Subsystem”	21
	c) “A Data Manager”	21
	d) “A Volatile Memory Subsystem”	22
	e) “A Controller”	23
	f) “At Least One of the Volatile and Non-Volatile Memory Subsystems Comprises One or More Memory Segments”	24
	g) “The Data Manager Is Configured as a Bi-Directional Data Transfer Fabric”	25
	h) “The Data Manager Further Including a Data Buffer”	27
	i) “A Data Format Module”	28
2.	Claim 2	30
3.	Claim 3	32
4.	Claim 4	33
5.	Claim 5	34
6.	Claim 6	37
7.	Claim 7	38
	a) Preamble	38
	b) “Receiving Control Information”	38
	c) “Identifying a Data Path”	39
	d) “Using a Data Manager and Controller”	40
	e) “Operating the Data Manager as a Bi-Directional Data Transfer Fabric”	41
	f) “Operating the Two or More Sets of Data Ports to Transfer Data”	41
	g) “Using the Controller of the Memory Module to Perform One or More of”	42
8.	Claim 8	44

9. Claim 9	44
10. Claim 10	45
11. Claim 11	45
12. Claim 12	46
13. Claim 13	47
14. Claim 14	48
B. Claims 1-14 Are Obvious Over Best in view of Roy	49
C. Claims 2 and 8 Are Obvious Over Best in view of Tsunoda, With or Without Roy.....	53
D. Claims 5 and 12-14 Are Obvious Over Best in view of Roothparvar, With or Without Roy.....	55
1. Claims 5 and 12-14	55
2. Claim 14	57
E. Claim 15 Is Obvious Over Best in view of Bonella, With or Without Roy	57
1. “Operating the Volatile Memory Subsystem at a First Clock Frequency”	57
2. “Operating the Non-Volatile Memory Subsystem at a Second Clock Frequency”	60
3. “Operating the Volatile Memory Subsystem at a Third Clock Frequency”	64
F. Claim 15 Is Obvious Over Best in view of Bonella and Ashmore, With or Without Roy	68
VI. CONCLUSION	69

Attachment A. Proof of Service of the Petition

Attachment B. List of Evidence and Exhibits Relied Upon in Petition

I. COMPLIANCE WITH REQUIREMENTS FOR A PETITION FOR INTER PARTES REVIEW

A. Certification the 831 Patent May Be Contested by Petitioners

Petitioners certify they are not barred or estopped from requesting *inter partes* review of U.S. Patent No. 8,874,831 (“831 Patent”) (Ex. 1001). No Petitioner, nor any party in privity with a Petitioner, has filed a civil action challenging the validity of any claim of the 831 Patent. The 831 Patent has not been the subject of a prior *inter partes* review by any Petitioner or a privy of a Petitioner.

Petitioners also certify this petition for *inter partes* review is filed within one year of the date of service of a complaint alleging infringement of a patent – no complaint alleging infringement of the 831 Patent has been served on any Petitioner. Petitioners therefore certify this patent is available for *inter partes* review.

B. Fee for Inter Partes Review (§ 42.15(a))

The Director is authorized to charge the fee specified by 37 CFR § 42.15(a) to Deposit Account No. 50-1597.

C. Mandatory Notices (37 CFR § 42.8(b))

The real parties of interest of this petition are the Petitioners: SK hynix Inc., SK hynix America Inc. and SK hynix memory solutions Inc.

Petition for *Inter Partes* Review of U.S. Patent No. 8,874,831

The 831 Patent has not been involved in a prior legal proceeding. Claim 15 of the 831 Patent, however, recites similar limitations to the independent claims of the 831 Patent's parent patent, U.S. Patent No. 8,301,833 ("833 Patent"). The 833 Patent is involved in the following legal proceedings: *Netlist, Inc. v. SMART Modular Technologies, Inc.*, Case No. 8-13-cv-00996 (C.D. Cal.); *Smart Modular Technologies, Inc. v. Netlist, Inc.*, Case No. 4-13-cv-03916 (N.D. Cal.); *Diablo Technologies, Inc. v. Netlist, Inc.*, Case No. 4-13-cv-03901 (N.D. Cal.); *Netlist, Inc. v. Smart Modular Technologies, Inc.*, 4-13-cv-05889 (N.D. Cal.); *SanDisk Corp. v. Netlist, Inc.*, IPR2014-00994 (institution denied); and *SMART Modular Technologies Inc. v. Netlist, Inc.*, IPR2014-01370 (institution denied).

Lead Counsel is Joseph A. Micallef (Reg. No. 39,772), Sidley-SKH-IPR@sidley.com, (202) 736-8492. Backup Lead Counsel is Samuel A. Dillon (Reg. No. 65,197), Sidley-SKH-IPR@sidley.com, (202) 736-8298.

Service on Petitioner may be made by e-mail (Sidley-SKH-IPR@sidley.com), mail, or hand delivery to: Sidley Austin LLP, 1501 K Street, N.W., Washington, D.C. 20005. The fax number for lead and backup counsel is (202) 736-8711.

D. Proof of Service (§§ 42.6(e) and 42.105(a))

Proof of service is provided in **Attachment A**.

II. IDENTIFICATION OF CLAIMS BEING CHALLENGED

Petitioners propose several grounds for trial as set forth below, none of which is redundant. Each ground is based primarily on U.S. Patent Application Publication No. 2010/0110748 to Best (“Best”) (Ex. 1006). However, Petitioners also address several arguments that Patent Owner may raise by proposing grounds that more closely satisfy certain claim limitations. Such additional grounds are not redundant because they are “rational, narrowly targeted, and not burdensome.”

Great W. Casualty Co. v. Transpacific IP I Ltd., IPR2015-01912, Paper 10 at 17-18 (Mar. 22, 2016). Petitioners therefore respectfully request that trial be instituted on all grounds and arguments advanced herein. Specifically, this Petition seeks a finding that claims 1-15 of the 831 Patent are unpatentable as follows:

- (i) Claims 1-14 are anticipated under 35 U.S.C. § 102(e) by Best (Ex. 1006);
- (ii) Claims 1-14 are obvious under 35 U.S.C. § 103(a) over Best in view of Roy (Ex. 1008);
- (iii) Claims 2 and 8 are obvious under § 103(a) over Best in view of Tsunoda (Ex. 1009), with or without Roy;
- (iv) Claims 5 and 12-14 are obvious under § 103(a) over Best in view of Roohparvar (Ex. 1019), with or without Roy;
- (v) Claim 15 is obvious § 103(a) over Best in view of Bonella (Ex. 1013), with or without Roy;
- (vi) Claim 15 is obvious § 103(a) over Best in view of Bonella and Ashmore (Ex. 1011), with or without Roy.

Petitioner's proposed claim constructions, evidence relied upon, and precise reasons why the claims are unpatentable are provided in §§ **III-V**, below. The evidence relied upon in this petition is listed in **Attachment B**.

III. RELEVANT INFORMATION CONCERNING THE CONTESTED PATENT

A. Effective Filing Date of the 831 Patent

The 831 Patent resulted from Application No. 13/559,476, filed July 26, 2012, a continuation-in-part of Application No. 12/240,916 (now U.S. Patent No. 8,301,833) filed on September 29, 2008, which is a continuation of Application No. 12/131,873, filed on June 2, 2008 (abandoned). The 831 Patent claims priority to Provisional Application Nos. 60/941,586 ("586 Application") (Ex. 1005), filed on June 1, 2007, and 61/512,871, filed July 28, 2011. The only claim of priority relevant to this proceeding is the claim of priority to the 586 Application.

The 831 Patent is not entitled to its claim of priority to the 586 Application. The Examiner determined as much during prosecution, and the Patent Owner apparently agreed. The Examiner initially rejected all of the 831 Patent's claims for priority. Ex. 1002, 126-35. After a telephone interview, "[t]he Examiner and Applicant's representative reached agreement that the present application is entitled to the benefit" of "Applicant's claim for priority based on prior provisional application 61/512,871, filed 7/28/2011, and on prior application 12/240,916, filed

9/29/2008.” Ex. 1002, 97-100. Notably, there is no mention about the Patent Owner arguing or the Examiner agreeing that the 831 Patent was entitled to its claim of priority to the 586 Application. Indeed, Patent Owner subsequently filed a document attempting to clarify the Examiner’s reasons for allowance, yet made no mention of any disagreement with the Examiner’s priority date determination. *See* Ex. 1002, 81-82. The Patent Owner thus did not contest the Examiner’s determination that the 831 Patent was not entitled to a priority date earlier than September 29, 2008.

None of the claims of the 831 Patent are entitled to priority to the 586 Application. For example, independent claim 1 recites “*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.*” This limitation is not supported by the 586 Application, which only discloses “signal level translation” and “address decoding.” Ex. 1005, ¶16. Neither “signal level translation” or “address decoding” involve “*format[ing] data to be transferred.*” Decl. of Ron Maltiel (Ex. 1003), ¶¶48-50. The 586 Application briefly mentions different bus widths, but there is no discussion of “*format[ing]*” or re-“*format[ing]*” the data between these various bus widths. Ex. 1005, ¶¶7, 10-11. The remaining written description does not describe anything that one of ordinary

skill in the art could reasonably characterize as “*format[ing] data to be transferred.*” Ex. 1003, ¶¶50-51. Thus, the 586 Application does not provide written description support for claim 1.

Independent claim 7 recites “*using the controller of the memory module 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 any two or more of the memory controller, the volatile memory subsystem, and the non-volatile memory subsystem.*” This limitation is not supported by the written description of the 586 Application, which does not disclose any examples of “*memory address mapping,*” “*address domain conversion,*” and “*data width modulation.*” Ex. 1003, ¶52. For example, the 586 Application does not disclose performing “*data width modulation,*” at most disclosing “*signal level translation,*” “*address decoding,*” and briefly mentioning different bus widths, but none of these involve “*data width modulation.*” Ex. 1005, ¶¶7, 10-11, 16. The remaining portions of the written description do not describe anything that one of ordinary skill in the art could reasonably characterize as “*data width modulation.*” Ex. 1003, ¶53. Further, address decoding is not the same thing as “*address domain conversion*”—one of ordinary skill in the art would not consider the 831 Patent’s brief reference to “*address decoding*” to provide written

description support for “*address domain conversion.*” Ex. 1003, ¶53. Thus, the 586 Application does not provide written description support for claim 7.

Claims 2-6 and 8-15 depend from claims 1 and 7 and lack written description support in the 586 Application for these same reasons. Claims 1-15 are therefore entitled to an earliest effective filing date of no earlier than **June 2, 2008.**¹

B. Person of Ordinary Skill in the Art

One of ordinary skill in connection with the subject matter of the 831 Patent in either the 2007 or 2008 time frames would be a person with a Bachelor’s degree in materials science, electrical engineering, computer engineering, computer science, or in a related field and at least one year of experience with the design or development of semiconductor non-volatile memory circuitry or systems. Ex. 1003, ¶¶55-56.

C. The 831 Patent

1. Technical Overview

The 831 Patent discloses a memory module couplable to a memory controller hub (MCH) of a host system including a non-volatile memory

¹ As explained below, the references at issue are prior art even if the 831 Patent is entitled to its claim of priority to June 1, 2007.

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 read/write commands from the MCH and to direct transfer of data between any two or more of the MCH, the volatile memory subsystem, and the non-volatile memory subsystem based on the commands. Ex. 1001, Abstract. Figure 5A shows a memory module in accordance with certain embodiments of the 831 Patent:

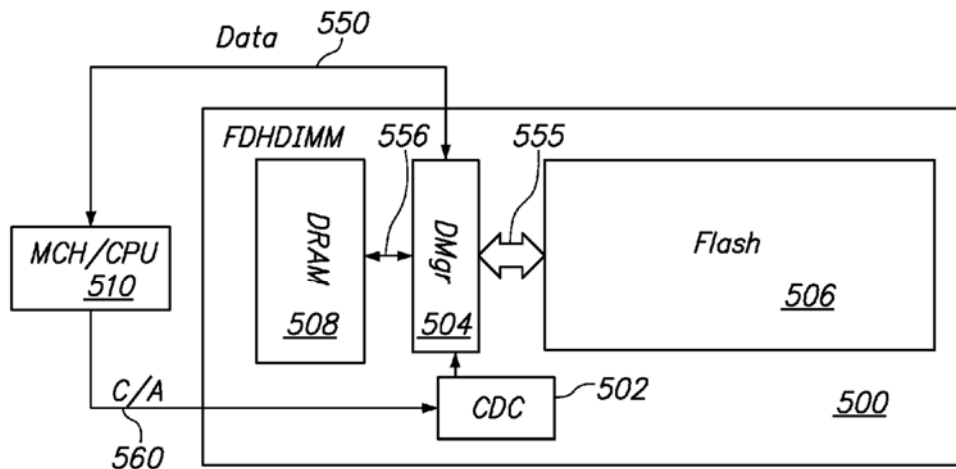


FIG. 5A

Id., 7:7-8, Fig. 5A; Ex. 1003, ¶57.

As shown in Figure 5A, the 831 Patent discloses an on-DIMM data manager (DMgr) 504. Ex. 1001, 10:43-46. The 831 Patent states that in 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. *Id.*, 11:8-13. The DMgr controls the data path routing amongst DRAMs, Flash and MCH. *Id.*, 11:13-15. An exemplary role of DMgr 504 is described with reference to Figure 6:




FIG. 6

Id., Fig. 6; Ex. 1003, ¶58. “DMgr 504 receives the data transfer size, formatting information, direction of data flow (via one or more multiplexers [611-622]), and the starting time of the actual data transfer from CDC 502.” Ex. 1001, 11:56-12:1; Ex. 1003, ¶59. The 831 Patent also states that “[i]n certain embodiments,

DMgr 504 also functions as a bi-directional data transfer fabric,” and discusses one embodiment of that functionality. Ex. 1001, 12:1-19; Ex. 1003, ¶60.

2. Prosecution History

The 831 Patent’s application was filed on July 26, 2012. After a preliminary amendment, Ex. 1002, 163-69, the Examiner issued a notice of allowance, *id.*, 126-35. The Examiner stated that “[t]he priority date granted to the examination of the present application is 7/26/2012,” *id.*, 132, and then indicated claims 1 and 13 (now claims 1 and 7) were allowable, *id.*, 133-34. On May 16, 2014, the Examiner and the Patent Owner conducted a telephone interview in which they agreed the 831 Patent was entitled to a September 29, 2008 priority date at earliest. *Id.*, 96-100.

D. Construction of Terms Used in the Claims

In this proceeding, claims must be given their broadest reasonable construction in light of the specification. 37 CFR § 42.100(b). If Patent Owner contends certain claims terms should have a special meaning, Patent Owner must amend the claims compliant with 35 U.S.C. § 112 to make them expressly correspond to those contentions. *See* 77 Fed. Reg. 48764 at II.B.6 (Aug. 14, 2012); *cf. In re Youman*, 679 F.3d 1335, 1343 (Fed. Cir. 2012).

1. “Bi-Directional Data Transfer Fabric” (All Claims)

The broadest reasonable interpretation of “*bi-directional data transfer fabric*” as recited in claims 1 and 7 includes a data transfer interconnect operable in two directions.

The 831 Patent does not define the phrase, and uses it or a similar phrase only four times in the specification, three of which are in the summary of the invention. *E.g.*, Ex. 1001, 3:52-55, 5:12-15, 5:32-36. There is only one use of “*bi-directional data transfer fabric*” or a similar phrase in the detailed description:

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).

Ex. 1001, 12:1-11 (emphasis added).

This paragraph relates DMgr 504’s function as a “bi-directional data transfer fabric” to the operation of multiplexors 611, 612, 621, and 622, shown as part of DMgr 504 in Figure 6. Ex. 1001, 12:1-5. These multiplexers provide “controllable data paths” in both directions between the two DRAMs, the host, and

the Flash. The paragraph as a whole therefore equates DMgr 504's function as a "bi-directional data transfer fabric" with its ability to provide "controllable data paths" operable in two directions between each of the two DRAMs, Flash, and the memory controller hub. Ex. 1003, ¶75. This understanding is supported by the ordinary meaning of "bidirectional" ("operating in two directions," *see, e.g.*, Microsoft Computer Dictionary (5th Ed.) (2002) (Ex. 1015), 57) and the 831 Patent's use of "fabric" as synonymous with "interconnect." *See, e.g.*, Ex. 1001, 11:56-12:1, 12:17, 13:55; Ex. 1003, ¶76.

The broadest reasonable interpretation of "*bi-directional data transfer fabric*" thus includes a data transfer interconnect operable in two directions.

2. "Set of Data Ports" (All Claims)

The broadest reasonable interpretation of "*set of data ports*" as recited in claims 1 and 7 includes an interface for a data bus comprising a set of interface lines. The 831 Patent does not define the term, but references the term in the context of Figure 6: "DMgr 504 may have more than 2 sets of data ports facing the Flash 506 and the DRAM 508." Ex. 1001, 12:3-5. These "more than 2 sets of data ports" are the three respective interfaces to data buses 608(1), 608(2), and 606 (shown in Fig. 6) operating between and DMgr 504 and Flash 506, DRAM 508(1), and DRAM 508(2), respectively. *Id.*, 12:5-11. One of ordinary skill in the art would understand that DMgr 504 having "more than 2 sets of data ports" refers to

DMgr 504 having more than two memory-facing interfaces to data buses, such as the three memory-facing data buses shown in Figure 6:

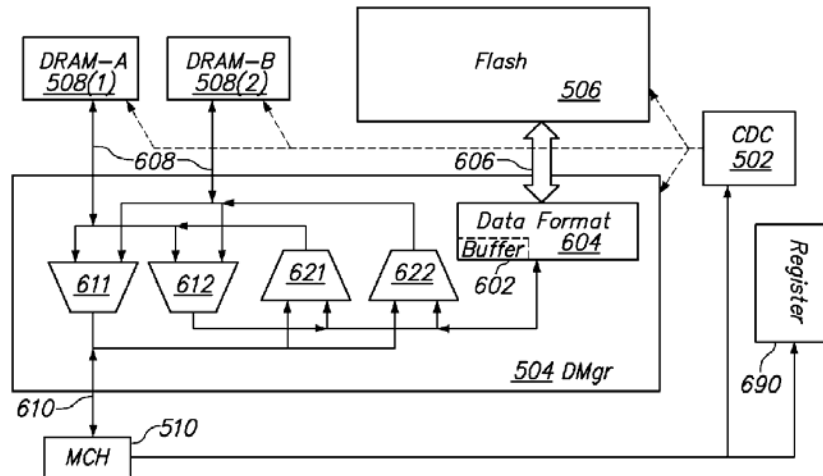


FIG. 6

Id., Fig. 6, 12:3-5; Ex. 1003, ¶¶79-80. One of ordinary skill in the art would understand there is necessarily an interface between DMgr 504 and each of these data buses. Ex. 1003, ¶¶79-80. Elements 606 and 608 are “wide data bus 606” and “relatively smaller width data bus 608.” Ex. 1001, 12:53-64. One of ordinary skill in the art would understand that this disclosure explains that data buses can have different widths, *i.e.*, a different number of hardware lines that make up the bus. Ex. 1003, ¶81; *see, e.g.*, Ex. 1015, 77 (defining “bus”), 412 (defining “port”: “[a]n interface through which data is transferred”).

Therefore, the broadest reasonable interpretation of the phrase “*set of data ports*” includes an interface for a data bus comprising a set of interface lines.

3. “Format Data” (Claims 1-6)

The broadest reasonable interpretation of “*format data*” includes adjusting the data’s width, *e.g.*, changing data from one width to another width.

The 831 Patent discloses “data format module 604” controlling data transfer from the Flash memory so as to “match[] the data rate and/or data format” through, for example, adjusting the width of data when traversing dissimilarly sized buses. Ex. 1001, 12:36-42, 12:53-64. One of ordinary skill in the art would understand that the phrase “*format data*” would at least encompass the ways the 831 Patent itself “*format[s] data*,” such as adjusting the data’s width from a first width to a second width. Ex. 1003, ¶¶84-85.

Therefore, the broadest reasonable interpretation of “*format data*” includes adjusting the data’s width.

4. “Operable at a ... Clock Frequency” (Claim 15)

The Board has previously construed the claim term “*clock frequency*” in a proceeding involving related the 833 Patent (Ex. 1014) to require “identification of a clock running at a particular frequency.” IPR2014-00994, Paper 8 at 6. This interpretation is consistent with the 831 Patent’s specification, which incorporates the 833 Patent by reference. Ex. 1001, 1:5-13; *see, e.g.*, Ex. 1014, 17:25-18:13; Ex. 1003, ¶40. Petitioners have applied this interpretation below.

IV. OVERVIEW OF THE PRINCIPAL PRIOR ART

A. U.S. Patent Application Publication No. 2010/0110748 to Best (Ex. 1006)

1. Prior Art Status

U.S. Patent Application Publication No. 2010/0110748 to Best (“Best”) (Ex. 1006) was filed on October 15, 2009 as a national stage application of PCT/US08/60566, filed April 17, 2008, and claimed priority to a provisional application (60/912,321) (“321 Application”) (Ex. 1007) filed on April 17, 2007. Best is prior art under 35 U.S.C. § 102(e) (pre-AIA) for two reasons. First, as explained above in Section III.A, the earliest effective filing date of the 831 Patent’s claims is June 2, 2008. Second, Best is entitled to the priority date of the 321 Application because it provides written description support for Best’s claims, therefore entitling Best to the filing date of the 321 Application. *See Dynamic Drinkware, LLC v. Nat’l Graphics, Inc.*, 800 F. 3d 1375, 1381-82 (Fed. Cir. 2015). Best is therefore prior art under Section 102(e) regardless of whether the 831 Patent is entitled to its June 1, 2007 priority date.

Best and the 321 Application contain essentially identical written descriptions, as can be seen in Appendix A to Mr. Maltiel’s report (Ex. 1003).

Best's paragraphs correspond to the 321 Application as follows (Ex. 1003, ¶90):

Best (Ex. 1006)	321 Application (Ex. 1007)
¶2	¶2
¶3-11	¶3
¶¶12-31	¶¶4-23
¶¶32-33	¶24
¶34	¶25

The 321 Application was also filed with the same 40 claims filed in Best, explicitly providing written description support for each claim, and includes the same set of figures. *Compare* Ex. 1006, claims 1-40, Figs. 1-7 *with* Ex. 1007, 27-29 (claims 1-40), 37-38 (Figs. 1-7).

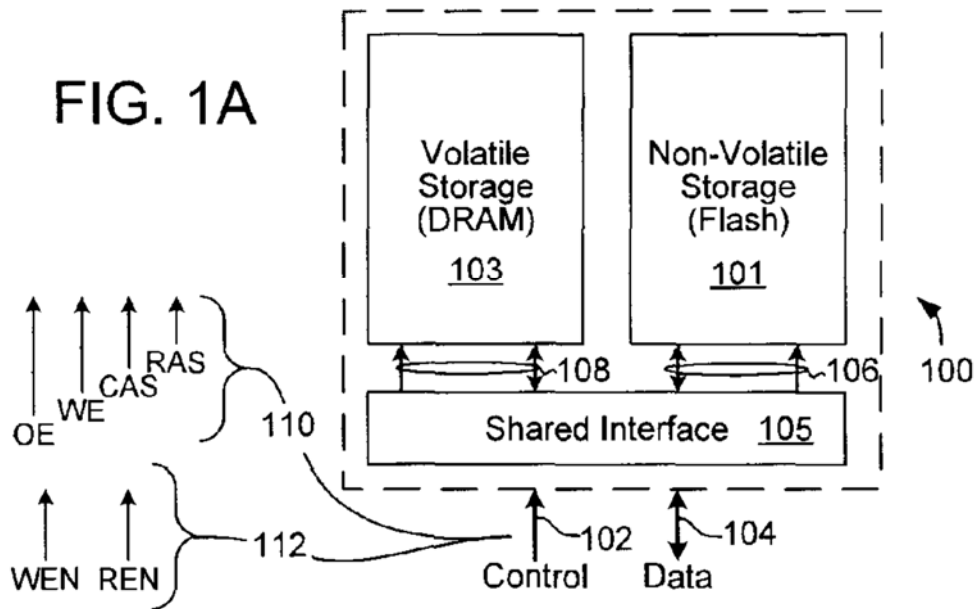
The 321 Application provides written description support for each of Best's claims. Each element of Best's claim 1, for example, has written description support in the 321 Application. The 321 Application discloses a “*memory device disposed within an integrated circuit (IC) package,*” *see, e.g.,* Ex. 1007, ¶8 (“[T]he hybrid memory device may ... form an integrated-circuit (IC) package.”); “*a first storage die having an array of volatile storage cells,*” *see, e.g., id.,* 37 (Fig. 2) (“DRAM Memory Array Die 103”); “*a second storage die having an array of non-volatile storage cells,*” *see, e.g., id.,* 37 (Fig. 2) (“NV Memory Array Die 101”); “*a shared interface circuit to receive information associated with a memory access operation to be performed within the memory device and to select, according to the information, either the first storage die or the second storage die to be accessed in*

the memory access operation,” see, e.g., id., ¶7 (“[T]he shared interface 105 includes circuitry to forward control and data signals to the appropriate memory die”), ¶9; Ex. 1003, ¶93.

As Dr. Maltiel explains in greater detail in his report, each of Best’s other claims are similarly supported in the 321 Application to the same extent and same manner as in Best. Ex. 1003, ¶94-133. Best is therefore prior art to the 831 Patent regardless of whether the patent is entitled to a June 1, 2007 date.

2. Overview of Best

Best discloses a composite, hybrid memory device including a first volatile storage die and a second non-volatile storage die disposed within an integrated circuit package. Ex. 1006, Abstract. The device includes a shared interface circuit to receive memory access commands directed to the first storage die and the second storage die and to convey read and write data between an external data path and the first and second storage dice. Ex. 1006, Abstract. Figure 1A illustrates an embodiment of this hybrid memory device:



Ex. 1006, Fig. 1A, ¶4; Ex. 1003, ¶134.

As shown in Figure 1A, the non-volatile storage IC 101 is implemented by a Flash memory die, and the volatile storage IC 103 is implemented by a DRAM die. Ex. 1006, ¶13. This embodiment is further described by Figures 2 and 3. Ex. 1003, ¶135. Best teaches that Figure 3 further describes the embodiment of Figure 2, which itself further describes the embodiment of Figure 1A, such that Figures 1A, 2, and 3 describe a single embodiment. *Id.*, ¶136. For example, Figure 1A illustrates “an embodiment of a hybrid, composite memory device 100 having ... shared-interface IC 105.” Ex. 1006, ¶13. Figure 1B illustrates an alternative embodiment that puts the location of shared-interface 105 inside the Flash memory, but otherwise the shared-interface’s functionality is the same. Ex. 1006,

¶16. Figure 2 “illustrates an embodiment ... with *the* shared interface circuitry shown in greater detail,” *id.*, ¶17 (emphasis added), and thus further describes the shared interface described as part of Figure 1A. Figure 3 “illustrates an embodiment of a data control/steering circuit 150 that may be used to *implement* the data control/steering circuit 131 of FIG. 2.” *Id.*, ¶21 (emphasis added). Figure 3 (and its associated description) thus further describes the functionality of Figure 2 (and its associated description), which itself further describes the functionality of Figure 1A (and its associated description). Ex. 1003, ¶136.

Best discloses two alternative ways that memory addresses are mapped to the volatile and non-volatile storage dies. In the Figure 4 hybrid storage embodiment, “non-overlapping address ranges apply to each of the storage dice 101 and 103 to form” a contiguous address space. Ex. 1006, ¶17. In the alternative Figure 7 “Shadow Operation” embodiment, “some or all of the volatile memory address range may overlap with the non-volatile memory address range to enable an operation referred to herein as memory shadowing.” Ex. 1006, ¶24.

One of ordinary skill in the art would understand that Best’s functionality in Figures 1-3 would operate the same in the “Shadow Operation” embodiment except that the address ranges may overlap so as to enable memory shadowing as described. Ex. 1003, ¶137.

V. PRECISE REASONS FOR RELIEF REQUESTED

A. Claims 1-14 Are Anticipated by Best

1. Claim 1

a) Preamble


Claim 1 recites “[a] memory module couplable to a memory controller of a host system.” Best discloses a “hybrid, composite memory device having non-volatile and volatile memories implemented in distinct integrated circuit (IC) dice that are packaged together and accessed through a shared interface.” Ex. 1006, ¶¶12, 17, Fig. 2; Ex. 1003, ¶¶139-140. The memory device includes DRAM die 103, Flash memory die 101, and shared interface circuitry including, *e.g.*, “an external request interface 125, external data interface 133, command decoder 122, address queue 135, DRAM control circuit 129, Flash control circuit 137, and data control/steering circuit 131.” Ex. 1006, ¶17. “[I]ncoming control signals and addresses ... are received in the external request interface 125 via control/address (CA) path 126 and then forwarded to the command decoder 122.” *Id.*, ¶17. This memory device (“*memory module*”) includes an interface that “receive[s] commands from a controller device (not shown in FIG. 1A),” (“*couplable to a memory controller of a host system*”). *Id.*, ¶14; Ex. 1003, ¶141. Therefore, Best discloses this claim element.

b) “A Non-Volatile Memory Subsystem”

Claim 1 recites “a non-volatile memory subsystem.” Best’s memory device includes a “non-volatile memory subsystem” in the form of a Flash memory. Ex. 1006, ¶17. The Flash memory can be “NAND-Flash or NOR-Flash” or “any other electrically-erasable or electrically-alterable storage technology.” *Id.*, ¶13; Ex. 1003, ¶¶139-141, 144. Therefore, Best discloses this claim element.

c) “A Data Manager”

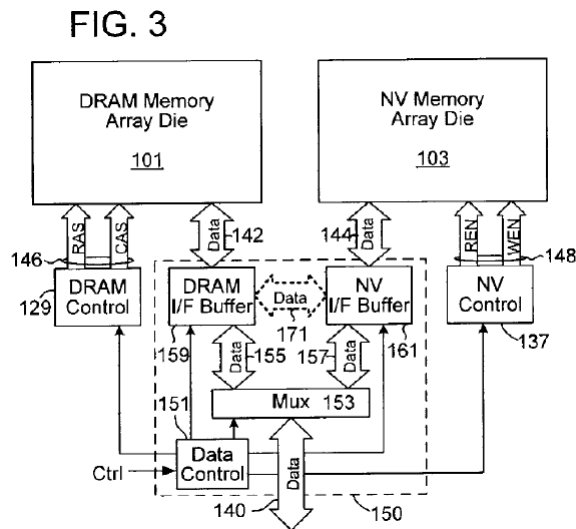
Claim 1 recites “a data manager coupled to the non-volatile memory subsystem.” Best discloses a “data manager” in the form of a data control/steering circuit in combination with an external interface, as annotated on Figure 2:



Ex. 1006, Fig. 2 (annotated); Ex. 1003, ¶147. The data control/steering circuit controls “the transfer of data between a shared internal data bus and dedicated

internal data buses associated with the volatile and non-volatile storage dice, respectively.” Ex. 1006, ¶20.

Figure 3 shows one embodiment of the data control/steering circuit with “exemplary interconnections with the DRAM and NV control circuits 129, 137 and also the volatile and non-volatile memory dice” (“coupled to the non-volatile memory subsystem”):



Id., ¶21; Ex. 1003, ¶148. Therefore, Best discloses this claim element.

d) “A Volatile Memory Subsystem”

Claim 1 recites “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.” Best discloses a “volatile memory subsystem” in the form of a DRAM. Ex. 1006, ¶21. The data controller/steering circuit (part of the “data manager”) is interconnected with (“coupled to”) the DRAM control circuit and the

DRAM die. *Id.* Data may be “*exchange[d]*” in either direction between the DRAM and Flash memory along inter-die data path 171 (“... *operable to exchange data with the non-volatile memory subsystem*”). *Id.* As shown in Figure 3, data control/steering circuit 150 (part of the “*data manager*”) includes inter-die data path 171, so transfers along that path occur “*by way of the data manager.*” *Id.*; Ex. 1003, ¶151. Therefore, Best discloses this claim element.

e) “A Controller”

Claim 1 recites “*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.*”

Best discloses a command decoder (“*controller operable to receive commands from the memory controller*”) which is forwarded “incoming control signals and addresses” (Ex. 1006, ¶17) and “outputs ... an enable signal and corresponding memory access control signals to the DRAM control circuit ... and NV control circuit ...” (*id.*, ¶18). This output enable signal is generated based on the incoming control signals and addresses (“*based on at least one received command from the memory controller*”) and directs both “*(i) operation of the non-*

volatile memory subsystem” and “(ii) operation of the volatile memory subsystem”. *Id.*, ¶19, 29; Ex. 1003, ¶154. The command decoder also controls the transfer of data between these memories and the external “*memory controller*” in each possible direction “*based on at least one command from the memory controller*”. Ex. 1006, ¶21; *see id.*, ¶17 (describing how the output enable is chosen in Best’s Figure 4 contiguous addressing mode), ¶29 (describing how the output enable is chosen in Best’s Figure 7 shadow address mode); Ex. 1003, ¶¶155-157. Therefore, Best discloses this claim element.

f) “At Least One of the Volatile and Non-Volatile Memory Subsystems Comprises One or More Memory Segments”

Claim 1 recites “*at least one of the volatile and non-volatile memory subsystems comprises one or more memory segments, each memory segment comprising at least one memory circuit, memory device, or memory die.*”

Best discloses “the volatile and non-volatile memories being implemented by a DRAM die 103 and Flash memory die 101, respectively,” each of which is a “*memory segment.*” Ex. 1006, ¶17. Best further discloses that “while only two storage dice are shown, multiple non-volatile storage dice and/or multiple volatile storage dice may be provided and selected by the shared interface circuitry based on incoming address and/or control signals.” *Id.*, ¶15. One of ordinary skill in the art would understand that each of these storage dice (*i.e.*, a “*memory die*”) is also a

“*memory circuit*” and a “*memory device.*” Ex. 1003, ¶160. Therefore, Best discloses this claim element.

g) “*The Data Manager Is Configured as a Bi-Directional Data Transfer Fabric*”

Claim 1 recites “*the data manager is configured as a bi-directional data transfer fabric having two or more sets of data ports, a first set of data ports of the two or more sets of data ports is coupled to the volatile memory subsystem, a second set of data ports of the two or more sets of data ports is coupled to the non-volatile memory subsystem, the two or more sets of data ports being operable by the data manager to transfer data to or from one or more memory segments of the volatile or non-volatile memory subsystems.*”

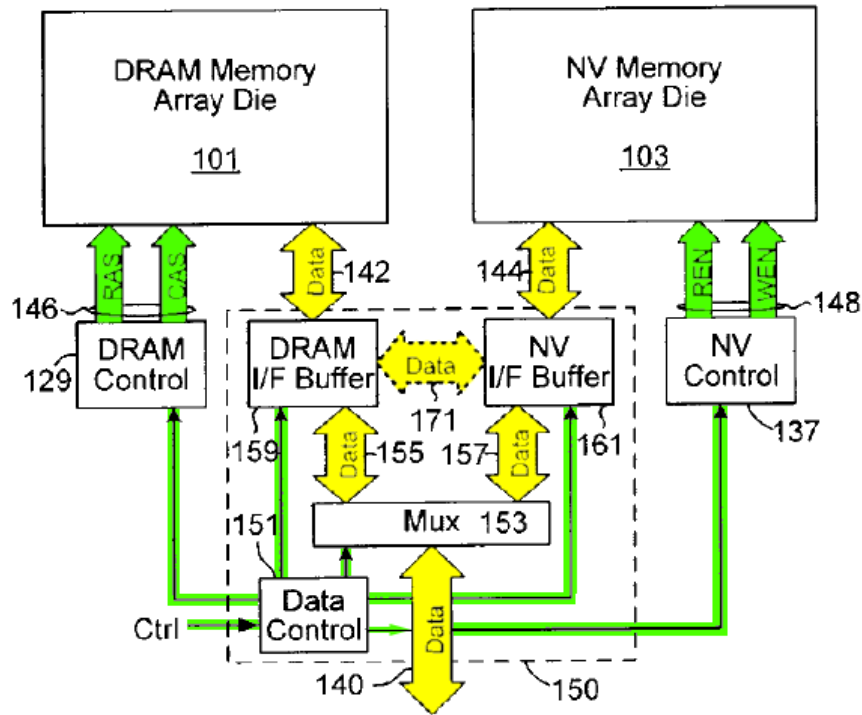
Best’s data control/steering circuit contains a data transfer interconnect where each connection can operate in two directions (“*bi-directional data transfer fabric*”), allowing data to flow from the outside memory controller, DRAM, and Flash memory to any other one of those same components. Ex. 1006, ¶21; *id.*, ¶20 (“The data control/steering circuit 131 is used to control the transfer of data between the shared internal data bus and dedicated internal data buses associated with the volatile and non-volatile storage dice, respectively.”). One of ordinary skill in the art would understand that this data control/steering circuit would be a

data transfer interconnect operable in two directions, *i.e.*, either to or from any of the connected components. Ex. 1003, ¶163.

The data control/steering circuit has “*two or more sets of data ports.*” The “*first set of data ports ... coupled to the volatile memory subsystem*” is the interface to the primary volatile data path 142 between data control/steering circuit 150 and DRAM 101. Ex. 1006, ¶21, Fig. 3. The “*second set of data ports ... coupled to the non-volatile memory subsystem*” is the interface to primary non-volatile data path 144 between data control/steering circuit 150 and NV memory 103. *Id.*, ¶21, Fig. 3. Both of these interfaces are “*sets of data ports*” because each is a data bus interface comprising a set of signal lines. Ex. 1006, ¶¶15, 32; Ex. 1003, ¶164; *see* §III.D.1 above.

These sets of data ports are also “*operable by the data manager to transfer data to or from one or more memory segments of the volatile or non-volatile memory subsystems*” as explained above regarding the claimed “*controller,*” and as shown in Figure 3. Each set of interconnections between components is “*bi-directional*” in that data can flow in either direction, as exemplified by the two sided-arrow between each component annotated in yellow (compared to the single-sided arrows annotated in green, *e.g.*, the control signals):

FIG. 3



Ex. 1006, Fig. 3 (annotated); Ex. 1003, ¶165. The 831 Patent uses similar arrows to describe equivalent functionality. *See, e.g.*, Ex. 1001, Fig. 6. Therefore, Best discloses this claim element.

h) “The Data Manager Further Including a Data Buffer”

Claim 1 recites “*the data manager further including a data buffer for buffering data delivered to or from the non-volatile memory subsystem.*” Best discloses that “non-volatile-storage-die interface buffer 161” (“*a data buffer*”) buffers data “*delivered to*” the Flash memory: “data read out of the volatile storage die may be transferred ... to the non-volatile-die interface buffer 161 via the inter-die data path 171.” Ex. 1006, ¶21. This buffer also buffers data “*delivered from*”

the Flash memory: “Data may similarly be transferred in the opposite direction from the non-volatile die 101 to the volatile storage die 103 ... (i.e., transferring the data from buffer 161 ...).” *Id.*, ¶21. Ex. 1003, ¶180. Therefore, Best discloses this claim element.

i) “A Data Format Module”

Claim 1 recites “*the data manager further including ... 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.*”

Best discloses “*a data format module*” in the form of the underlying logic for a serializing/deserializing function contained within the data steering/control circuit and the external data interface. This serializing functionality is first described in relation to external data interface 133: “converting a sequence of relatively narrow data words received at a relatively high frequency via the external data path 128, to a lower frequency, wider-data-word sequence on the primary internal data path 140, and performing the reverse operation (serializing) for [the reverse] data flow.” Ex. 1006, ¶20. One of ordinary skill in the art would understand that “serializing” involves changing from a parallel transmission with a wider data width to a serial transmission with a narrower data width, and “deserializing” involves changing from a serial transmission to a parallel

transmission. Ex. 1003, ¶183; *see, e.g.*, Ex. 1015, 153 (defining “deserialize”), 473 (defining “serialize”). The data steering/control circuit also “perform[s] a serializing/ deserializing function for data transferred between the shared internal data path 140 and the volatile-die data path 142 and/or between the shared internal data path 140 and the non-volatile-die data path 144.” Ex. 1006, ¶20. One of ordinary skill would recognize that this functionality would necessarily be implemented in digital logic in some form, *i.e.*, a “*module*.” Ex. 1003, ¶184.

This serializing/deserializing functionality “*formats data to be transferred*” between the external interface, DRAM, and Flash memories by converting data from one format (a first data word width) into another format (a second data word width). Ex. 1003, ¶185.² Further, “any external protocol may be used ... and any internal protocol or set of protocols may be used to control the volatile and non-volatile storage dice.” Ex. 1006, ¶15. This is accomplished by “converting signals from the external data access protocol to an internal data access protocol” such as “convert[ing] from DRAM format to Flash memory format.” Ex. 1006, ¶15. This functionality is activated based on the external commands received by the memory

² The 831 Patent similarly describes adjusting for data width, frequency, and rate as part of the “data format module.” Ex. 1001, 12:53-64; Ex. 1003, ¶186.

device (“*based on control information received from the controller*”), as explained above regarding the claimed “*controller.*” Ex. 1003, ¶185.

Best teaches that this functionality includes other types of coordination during transfer operations, such as minimum transfer size parameters set programmatically. Ex. 1006, ¶23. This type of adjustment of the data is also “*format[ing] data*” in that it reformats the data’s block size between memories. For example, data that is written from DRAM to Flash would need to be reformatted from, for example, 32 bits per transaction to 32,768 bits transaction. *Id.*, ¶23. This type of functionality is inherently required when transferring data between memories with different transfer sizes. Ex. 1003, ¶186. Therefore, Best discloses this claim element.

2. Claim 2

Claim 2 recites “*[t]he memory module of claim 1, 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.*”

Best discloses that the data control/steering circuit in combination with an external interface (collectively the “*data manager*”) is “*operable to control one or more of data flow rate*” and “*data transfer size ... in response to receiving at least one of a control signal*” by “control[ing] the transfer of data between a shared

internal data bus and dedicated internal data buses associated with the volatile and non-volatile storage dice.” Ex. 1006, ¶20; Ex. 1003, ¶189. Commands direct the memory device to, *e.g.*, write data from the external interface to either the DRAM and Flash memory, or between the DRAM and Flash memory. Ex. 1006, ¶21.

These different components have different data word widths, Ex. 1003, ¶¶183-186, so by selecting a given transfer the “*data manager*” causes and “*controls*” the formatting operations performed, Ex. 1006, ¶20; Ex. 1003, ¶190.

These formatting operations affect the “*data flow rate*” and the “*data transfer size*.” For example, as explained above regarding claim 1, Best teaches a serializing/deserializing function that converts between relatively narrow data words to relatively wider data words, or vice versa. Ex. 1006, ¶20; Ex. 1003, ¶¶183-184. Selecting between the Flash and DRAM necessarily affects the “*data flow rate*” due to the different access speeds associated with DRAM compared to Flash — writing to Flash is much slower than writing to DRAM. Ex. 1003, ¶191 (discussing, *e.g.*, Ex. 1009, ¶97, Ex. 1016, 25, Ex. 1017, 67). The data manager will have to therefore set the data flow rate based on whether the data is written to DRAM or to Flash memory. Ex. 1003, ¶191.

As also explained above regarding claim 1, Best discloses setting a “minimum transfer size parameter” (“*data transfer size*”) that controls the

minimum amount of information read or written during a given transaction. Ex. 1006, ¶23. Ex. 1003, ¶¶186, 192. Therefore, Best discloses this claim element.

3. Claim 3

Claim 3 recites “[t]he memory module of claim 1, 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.”

Best teaches that external commands can be received that direct the memory device to write data from the external interface to either the DRAM and Flash memory, or between the DRAM and Flash memory. Ex. 1006, ¶21. This control is affected through the interconnections made by the data control/steering circuit (part of the “data manager”). *Id.*; Ex. 1006, ¶20; Ex. 1003, ¶198. Data is transferred “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” (e.g., from the command decoder). Ex. 1006, ¶21. This functionality is further explained above regarding claim 1’s “controller” and “data manager.” See Ex. 1003, ¶¶147-148, 154-155, 199. Therefore, Best discloses this claim element.

4. Claim 4

Claim 4 recites “[t]he memory module of claim 3, 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.”

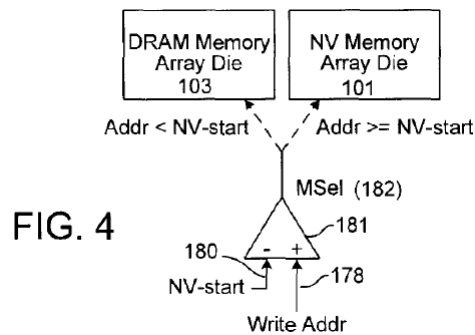
As explained above regarding claim 2, Best discloses a “data manager” in the form of a data control/steering circuit in combination with an external interface that collectively control the “data flow rate” and “data transfer size.” Ex. 1006, ¶¶20, 23. This same functionality also discloses controlling “data transfer bit width” due to the serializing/de-serializing functionality that converts between different data widths, *i.e.*, data width in bits. Ex. 1006, ¶20. Ex. 1003, ¶¶189-191, 202. As explained above regarding claim 1’s “controller,” Best teaches the “data manager” controlling the “direction of data flow” through directing the memory device to write data from the external interface to either the DRAM and Flash memory, or between the DRAM and Flash memory. Ex. 1006, ¶21; Ex. 1003, ¶¶154-155. These different components have different data word widths and minimum data transfer sizes (“data transfer size”), Ex. 1003, ¶¶183-186, and selecting a given transfer will affect the “data flow rate” because of the differences between DRAM and Flash access rates. Ex. 1006, ¶20; Ex. 1003, ¶¶189-191, 203. Finally, by directing a write along one of these paths, Best’s data control/steering

circuit controls the “*starting time of data transfer*” by indicating that the data transfer should start immediately. Ex. 1003, ¶204. Therefore, Best discloses this claim element.

5. Claim 5

Claim 5 recites “[t]he memory module of claim 1, 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.” Best teaches this feature in two ways.

First, Best’s Figure 4 illustrates “an address comparator circuit that may be implemented within the command decoder of” Figure 2:




Ex. 1006, ¶8, Fig. 4; Ex. 1003, ¶207. Best explains that, for example, the memory device can include “a predetermined (or programmatically established) memory address, NV-start 180, that marks the start of the non-volatile memory address

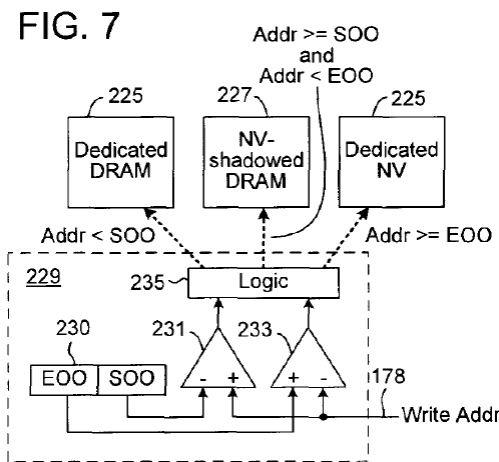
range.” Ex. 1006, ¶17. This starting address demarcates “*a first memory address space of the volatile memory subsystem*” and “*a second memory address space of the non-volatile memory subsystem*” because addresses equal to or greater than it are assigned to the non-volatile memory and all addresses less than it are assigned to the volatile memory. *Id.*; Ex. 1003, ¶208. Both “*memory spaces*” are, in Best’s “predetermined” example, “*configure[d] ... in response to ... memory address space initialization information of the memory module.*” Ex. 1006, ¶17. Best’s “predetermined” value for NV-start 180 is “information” used to initialize the memory address space for the memory device, so is “*memory address initialization information of the memory module.*” *Id.*; Ex. 1003, ¶209.

Additionally, when “programmatically established” the “*memory spaces*” are “*configure[d] ... in response to ... a received command from the memory controller,*” which is the disclosed component that externally controls the memory device. *See* Ex. 1006, ¶14. One of ordinary skill in the art would understand Best’s “programmatically established” embodiment would involve commands received externally to the memory device. Ex. 1006, ¶¶14-15; Ex. 1003, ¶210. Best’s “programmatically establish[ing]” functionality would thus occur in response to a “*received command from the memory controller.*” Ex. 1003, ¶210.

Second, Best discloses a “Shadow Operation” addressing mode in Figure 6:



Ex. 1006, Fig. 6. Ex. 1003, ¶211. Here, data is initially “stored in a fast-access volatile storage die.” Ex. 1006, ¶25. “Some time after one or more write operations have been performed, a write-back trigger is detected” and “one or more internal data transfer operations are performed to transfer data ... from the DRAM to corresponding locations ... within the NV memory.” Ex. 1006, ¶25. Ex. 1003, ¶212. Figure 7 illustrates an address comparator circuit in such a mode:



Ex. 1006, ¶29, Fig. 7; Ex. 1003, ¶213.

The “start-of-overlap and/or end-of-overlap values may be reprogrammed during the course of normal operation of the hybrid memory device (i.e.,

configuration of the device need not limited to a one-time initialization),” Ex. 1006, ¶28, meaning that “*controller configures*” the address spaces assigned to DRAM (“*a first memory address space*”), NV-shadowed DRAM, and dedicated NV (“*a second memory address space*”). The NV-shadowed DRAM portion could be considered as part of both “*first*” and “*second*” address spaces, or as separate from both. Ex. 1003, ¶¶214-215. The configuration can also be done as a “one-time initialization” (“*in response to ... memory address space initialization information on the memory module*”) or “reprogrammed during the course of normal operation of the hybrid memory device” (“*in response to at least one of a received command from the memory controller*”). One of ordinary skill in the art would understand that this “reprogram[ing]” would necessarily occur “*in response to ... a received command from the memory controller.*” Ex. 1003, ¶¶210, 214-215. Therefore, and under either mapping, Best discloses this claim element.

6. Claim 6

Claim 6 recites “*the memory module of claim 1, wherein the volatile memory subsystem comprises DRAM memory.*” Best discloses that the volatile memory can be “a DRAM memory.” Ex. 1006, ¶12; Ex. 1003, ¶222. Therefore, Best discloses this claim element.

7. Claim 7

a) Preamble

Claim 7 recites “[a] method for managing a memory module by a memory controller, the memory module including volatile and non-volatile memory subsystems.” The broadest reasonable interpretation of “subsystems” requires at least one “volatile ... memory subsystem” and at least one “non-volatile memory subsystem.” As explained above regarding claim 1’s preamble, Best discloses a “method for managing a memory module by a memory controller.” Ex. 1003, ¶139; Ex. 1006, ¶12. Best’s “method of managing” this memory device is the functionality described in Best’s disclosure. Ex. 1003, ¶225. As explained above regarding claim 1, Best discloses “the memory module including volatile and non-volatile memory subsystems:” a DRAM (Ex. 1006, ¶14), and a Flash memory (*id.*, ¶14); Ex. 1003, ¶¶144, 151, 226. Therefore, Best discloses this claim element.

b) “Receiving Control Information”

Claim 7 recites “receiving control information from the memory controller, wherein the control information is received using a protocol of the volatile memory subsystem.” As explained above regarding claim 1’s “controller,” Best discloses “receiving control information from the memory controller.” Ex. 1003, ¶¶154-155. A command decoder is forwarded “incoming control signals and addresses” (Ex. 1006, ¶17) and outputs “an enable signal and corresponding memory access

control signals to the DRAM control circuit 129 and NV control circuit 137,” (*id.*, ¶18). Ex. 1003, ¶229. Best discloses “an embodiment having a DRAM external interface” such that “control and data signals may be forwarded without change to a DRAM die” (“*the control information is received using a protocol of the volatile memory subsystem*”). Ex. 1006, ¶15; Ex. 1003, ¶230. Therefore, Best discloses this claim element.

c) “Identifying a Data Path”

Claim 7 recites “*identifying a data path to be used for transferring data to or from the memory module using the received control information.*” As explained above regarding claim 1’s “*controller,*” Best’s command decoder controls the transfer of data between these memories and the external interface in each possible direction based on the “incoming control signals and addresses.” Ex. 1006, ¶¶17-18, 21. These control signals “indicate the direction” and target “of data flow during a memory access operation (read or write).” Ex. 1006, ¶21. Each is thus a “*data path to be used for transferring data to or from the memory module*” identified based on the “control signals and addresses” (“*using the received control information*”). Ex. 1006, ¶17; Ex. 1003, ¶233. Therefore, Best discloses this claim element.

d) “Using a Data Manager and Controller”

Claim 7 recites “*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.*”

As explained above regarding claim 1, Best discloses a “*data manager*” in the form of a data control/steering circuit in combination with an external interface. Ex. 1006, Fig. 2, ¶¶20-21; Ex. 1003, ¶¶147-148. Best also discloses a command decoder (“*controller of the memory module*”) which is forwarded “incoming control signals and addresses” (Ex. 1006, ¶17) and outputs “an enable signal and corresponding memory access control signals to the DRAM control circuit” and “NV control circuit,” (*id.*, ¶18). Ex. 1003, ¶¶154-155. These two components are used to “*transfer data between any two or more of the memory controller, the volatile memory subsystem, and the non-volatile subsystem,*” as explained above regarding claim 7’s “*identifying a data path*” step. Ex. 1003, ¶233. This transfer is performed “*based on*” both the “incoming control signals and addresses” (“*received control information*”) (Ex. 1006, ¶17) and the output of the command decoder that “indicate[s] the direction” and targets of the “data flow” (“*the identified data path*”) (*id.*, ¶21). Ex. 1003, ¶¶233, 236. Therefore, Best discloses this claim element.

e) “Operating the Data Manager as a Bi-Directional Data Transfer Fabric”

Claim 7 recites “*operating the data manager as a bi-directional data transfer fabric with two or more sets of data ports, wherein a first set of data ports of the two or more sets of data ports is coupled to the volatile memory subsystems, and a second set of data ports of the two or more sets of data ports is coupled to the non-volatile memory subsystem.*” Best discloses these features for the same reasons as it discloses claim 1’s “*bi-directional data transfer fabric*” limitation. Ex. 1003, ¶¶163-165, 239.

This clause refers to “*the volatile memory subsystems*” (*i.e.*, plural), but the claim does not introduce multiple “*volatile memory subsystems.*” See §V.A.7.a). To the extent this is not a typographical error, Best discloses embodiments with multiple volatile storage dice (plural “*volatile memory subsystems*”) operable using the same functionality described regarding single volatile storage dice systems. Ex. 1006, ¶15; Ex. 1003, ¶241. Therefore, Best discloses this claim element.

f) “Operating the Two or More Sets of Data Ports to Transfer Data”

Claim 7 recites “*operating the two or more sets of data ports to transfer data to or from one or more memory segments of the volatile or non-volatile memory subsystems based on control information received from the controller of the memory module.*” As explained above regarding claim 1’s “*bi-directional transfer*

fabric,” Best discloses these features. Ex. 1003, ¶¶163-165. For example, Best explains that “data control circuit 151 receives control signals from the command decoder that indicate the direction of data flow during a memory access operation (read or write).” Ex. 1006, ¶21. Each transfer is determined based on the “incoming control signals and addresses” (“*based on control information received from the controller of the memory module*”). *Id.*, ¶17. Ex. 1003, ¶244. Therefore, Best discloses this claim element.

g) “Using the Controller of the Memory Module to Perform One or More of”

Claim 7 recites “*using the controller of the memory module 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 any two or more of the memory controller, the volatile memory subsystem, and the non-volatile memory subsystem.*”

As explained above with respect to claim 7’s “*receiving control information*” step, Best discloses that two levels of conversion can occur: (1) a protocol conversion (*e.g.*, from DRAM to Flash), and (2) a “data serialization or deserialization” conversion that adjusts the data from one data path width to another. Ex. 1006, ¶15; Ex. 1003, ¶229-230, 247. As explained above regarding claim 1’s “*data format module*,” Best’s serializing/de-serializing functionality

involves changing the width of data words from wide to narrow, or vice versa. Ex. 1006, ¶20; Ex. 1003, ¶¶183-186. This serializing/deserializing functionality constitutes “*data width modulation between any two or more of the memory controller*” (controller device (Ex. 1006, ¶14)), “*the volatile memory subsystem*” (volatile-die), “*and the non-volatile memory*” (non-volatile-die). Ex. 1003, ¶248.

Best’s command decoder performs “*memory address translation, memory address mapping,*” and “*address domain conversion.*” As explained above regarding claim 5, Best discloses two alternative ways of mapping the address space of the hybrid device (“*memory address mapping*”): Figure 4’s no-overlap address space, Ex. 1006, ¶17, and Figure 7’s overlapping address space where “shadow range” data is initially written to DRAM and eventually written to non-volatile memory, Ex. 1006, ¶29; Ex. 1003, ¶¶207-214. In either embodiment, Best performs “*memory address mapping*” of the incoming address to the appropriate address on the respective memories. This is also “*address domain conversion*” because it converts addresses from the hybrid memory device’s domain to the domain of the volatile of non-volatile memory. Ex. 1003, ¶249.

In the shadow operation embodiment, “address translation may be performed to correlate entries within the non-volatile storage die to counterpart entries within the volatile storage die.” Ex. 1006, ¶25. This is “*memory address translation*” explicitly, and also “*memory address mapping*” in that a volatile

memory address is mapped to a non-volatile storage address. Ex. 1003, ¶250.

Therefore, Best discloses this claim element.

8. Claim 8

Claim 8 recites “[t]he method of claim 7, 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.” As explained above regarding claim 2, Best’s data control/steering circuit (part of the “data manager”) controlling “one or more of data flow rate, data transfer size, ... [and] data error monitoring.” Ex. 1003, ¶¶189-195. Best’s “data manager” also controls the “data width size” through the serializing/de-serializing functionality that adjusts the data width. Ex. 1006, ¶20; Ex. 1003, ¶¶183-186, 253. As explained above regarding claim 4, Best’s data control/steering circuit controls the “starting time of the transfer of data” by causing the data transfer to start immediately. Ex. 1003, ¶¶204, 254. Therefore, Best discloses this claim element.

9. Claim 9

Claim 9 recites “[t]he memory module of claim 7 [sic], wherein at least one of the volatile and non-volatile memory subsystems comprises one or more memory segments.” Best discloses this element for the same reasons as claim 1’s “memory segments” limitation. Ex. 1003, ¶¶160, 257.

10. Claim 10

Claim 10 recites “[t]he method of claim 7, 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.” As explained above regarding claim 1’s “volatile memory subsystem” limitation, Best discloses transferring from volatile to non-volatile memory and vice versa (“directing transfer of data bi-directionally between the volatile and non-volatile memory subsystems”). Ex. 1006, ¶21; Ex. 1003, ¶¶151, 260. As shown in Figure 3, data control/steering circuit 150 includes “inter-die data path 171” (“using the data manager”), (Ex. 1006, ¶21) and is based on “incoming control signals and addresses” from the outside controller (*id.*, ¶¶17-18) (“and in response to memory access commands received by the controller from the memory controller”). Ex. 1003, ¶261. Therefore, Best discloses this claim element.

11. Claim 11

Claim 11 recites “[t]he method of claim 10, further comprising buffering the data transferred between the memory controller and non-volatile memory subsystem using the volatile memory subsystem.” As explained above regarding claim 5, Best discloses a shadow operation mode where initially “data is stored in a fast-access volatile storage die” and a “write-back table within the volatile storage

die is updated,” and later, “after one or more write operations have been performed, a write-back trigger is detected” and “transfer operations are performed to transfer data ... from the DRAM to corresponding locations within the NV memory.” Ex. 1006, ¶25; Ex. 1003, ¶¶207-214. This operation in effect “*buffer[s] the data transferred between the memory controller and the non-volatile memory subsystem*” by using the DRAM (“*the volatile memory subsystem*”) as a temporary storage location until the write-back trigger is detected. Ex. 1006, ¶25; Ex. 1003, ¶264. Therefore, Best discloses this claim element.

12. Claim 12

Claim 12 recites “*[t]he method of claim 7, 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.*”

As explained above regarding claim 5, Best discloses two alternative ways of address space mapping : (1) Figure 4’s address space with no overlap, Ex. 1006, ¶17, and (2) Figure 7’s address space with an overlapping “shadow range” where data is initially written to DRAM and eventually written to non-volatile memory, *id.*, ¶29; Ex. 1003, ¶¶207-214, 267. In both embodiments, the memory space is

configured based on either “*a programmable value written into a register*” or “*a command received from the memory controller.*” Ex. 1003, ¶268.

Additionally, in the Figure 4 embodiment the memory device includes “a predetermined (or programmatically established) memory address, NV-start 180, that marks the start of the non-volatile memory address range.” Ex. 1006, ¶17. In the Figure 7 embodiment, the “start-of-overlap and/or end-of-overlap values may be reprogrammed during the course of normal operation of the hybrid memory device (i.e., configuration of the device need not limited to a one-time initialization).” Ex. 1006, ¶¶28, 33. These values also are “*value[s] corresponding to a first portion of the volatile memory subsystem*” or “*first portion of the non-volatile memory subsystem*” because they demarcate either the entirety or a portion of the address space of both memories. Ex. 1006, ¶¶17, 28; Ex. 1003, ¶269. Therefore, Best discloses this claim element.

13. Claim 13

Claim 13 recites “[*t*]he method of claim 12, 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.”

As explained above regarding claim 12, Best discloses “*configur[ing] 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*” through either Figure 4’s contiguous address space or Figure 7’s “shadow” or overlapping address space. Ex. 1006, ¶¶17, 28; Ex. 1003, ¶¶267-269. In both cases, setting the NV-start or start-of-overlap/end-of-overlap values “*configures*” the portions of the address space mapped to the volatile or non-volatile memories by dividing the memory space into respective portions. Ex. 1003, ¶273. Both embodiments “*present[] a unified memory space to the memory controller*” because they combine (with or without an overlap) the address spaces of the DRAM and non-volatile memory into a single address space accessible to the controller. Ex. 1006, Figs. 4, 7; Ex. 1003, ¶273. The 831 Patent similarly refers to a combined (with or without an overlap) address space when briefly discussing a “unified address space.” Ex. 1001, 14:13-52; Ex. 1003, ¶274. Therefore, Best discloses this claim element.

14. Claim 14

Claim 14 recites “[*t*]he method of claim 12, wherein the controller configures the memory space in the memory module using partitioning instructions that are application-specific.” As explained above regarding claim 12, Best discloses “programmatically establish[ing]” or “reprogramm[ing]” the memory

spaces “during the course of normal operation of the hybrid memory device (i.e., configuration of the device need not be limited to a one-time initialization).” Ex. 1006, ¶¶17, 28; Ex. 1003, ¶¶267-269.

These actions “*configure[] the memory space in the memory module using partitioning instructions*” by allocating portions of the memory space to the volatile and non-volatile memories. An ordinary artisan would understand that these memory spaces (and by extension the instructions configuring that space themselves) are “*application-specific*” because they are “programmatically established” and “reprogrammed” during the use of the memory device for future uses of the device. Ex. 1003, ¶¶277-278. This is consistent with the 831 Patent, which states that the memory module are “configurable for optimal use with a particular application,” Ex. 1001, 8:23-36, 11: 28-48, 14:64-66, 15:6-9, but does not disclose a specific *type* of partitioning instruction for use only with a specific application. Ex. 1003, ¶278. Therefore, Best discloses this claim element.

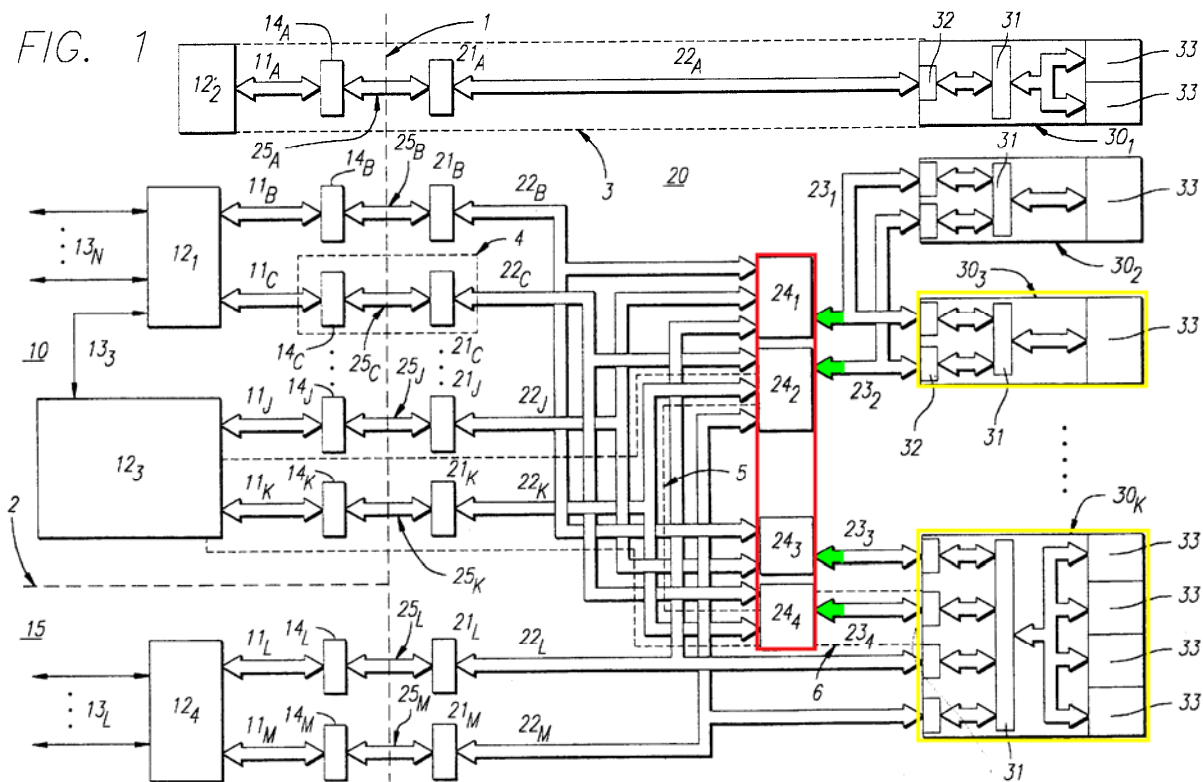
B. Claims 1-14 Are Obvious Over Best in view of Roy

Claims 1 and 7 recite “*a bi-directional data transfer fabric with two or more sets of data ports ...*” To the extent the claims require “*two or more*” independent read or write paths to the “*volatile memory subsystem*” and “*non-volatile memory subsystem,*” such a feature would have been obvious by the 831 Patent’s priority

date. Providing multiple independent write paths to volatile or non-volatile memories was well-known in the art by that time.

For example, U.S. Patent No. 6,065,092 to Roy (“Roy”) (issued May 16, 2000) (Ex. 1008) discloses a multichannel memory architecture including two independent channels between a master device and one or more memory clusters. Ex. 1008, 7:17-49; Ex. 1003, ¶¶166-167. Specifically, Roy discloses “two or more” independent read or write paths to multiple “memory subsystems.”

Annotated Figure 1 is a block diagram of a multichannel memory architecture:



Ex. 1008, Fig. 1 (annotated), 7:59-60. Ex. 1003, ¶168.

Roy discloses a “*bi-directional data transfer fabric*” (red set of multiplexer units 24, Ex. 1008, Fig. 1, 11:52-12:4) “*having two or more sets of data ports*” (green interfaces between multiplexer units 24₁-24₄ and bus segments 23₁-23₄, *id.*), “*a first set of data ports of the two or more sets of data ports is coupled to*” memory cluster 30₃ (top two green interfaces between multiplexer units 24₁ and 24₂ and bus segments 23₁ and 23₂, *id.*), “*a second set of data ports of the two or more sets of data ports is coupled to*” memory cluster 30K (bottom two green interfaces between multiplexer units 24₃ and 24₄ and bus segments 23₃ and 23₄, *id.*). Ex. 1003, ¶169.

This architecture is intended to be used with DRAM (a “*volatile memory subsystem*”) and Flash memory (a “*non-volatile memory subsystem*”). Ex. 1008, 9:43-49. Ex. 1003, ¶170. The master devices can be CPU or controller circuits, or specific circuitry dedicated solely to a memory interface. Ex. 1008, 10:50-67; Ex. 1003, ¶171. Best and Roy are both analogous art to the 831 Patent because each is from the field of memory systems. Ex. 1001, 1:20-23; Ex. 1006, ¶1; Ex. 1008, 1:21-26. Ex. 1003, ¶172.

It would have been obvious at the time of the 831 Patent’s priority date to modify Best’s shared interface to use a multichannel memory architecture, as taught by Roy. One of ordinary skill in the art would have been motivated to implement this architecture for all the reasons Roy describes, including allowing

independent and simultaneous transactions, Ex. 1008, 7:37-40, and increased performance by providing a wide effective channel, *id.*, 7:45-49; Ex. 1003, ¶173. Roy also teaches that a multichannel architecture provides substantial flexibility. Ex. 1008, 9:30-42; Ex. 1003, ¶173.

Roy discloses that “nearly identical address and control information” can be applied to each channel such that “[s]ubsequent transfer[s] of data on each of these channels can be synchronized to provide an effectively wider channel.” Ex. 1008, 10:28-32. This provides particular motivation to combine with Best in light of Best’s disclosure that “multiple non-volatile storage dice and/or multiple volatile storage dice may be ... selected ... based on incoming address and/or control signals.” Ex. 1006, ¶15; Ex. 1003, ¶174.

Best suggests such a modification through his disclosure of overlapping and pipelined memory operations. Ex. 1006, ¶18. One of ordinary skill would understand that multiple channels allow for further overlapping or pipelining of operations, such as allowing Best to write data from volatile to non-volatile memory as part of the “Shadow Mode” operation while allowing the host to simultaneously write data to volatile memory, thus improving the operation and responsiveness of the system. Ex. 1003, ¶175.

Modifying Best to use a multichannel architecture such as Roy’s would have been an arrangement of old elements (Best’s hybrid memory, Roy’s multichannel

architecture) with each performing the same function it had been known to perform and yielding no more than what one would expect from such an arrangement, *i.e.*, Best's system with a multichannel architecture. Ex. 1003, ¶176. Multichannel architectures were known in the art, and using one in Best would have involved only routine skill to implement the functionality described by Roy. *Id.* Such a modification would have therefore been obvious. *Id.*, ¶¶176, 240.

To the extent “*bi-directional data transfer fabric*” requires the ability for simultaneous data transfer, such a feature would be taught by the combination with Roy. Roy's multichannel architecture provides “bi-directional data transfers” that “can be operated simultaneously.” Ex. 1008, 9:11-22, 10:23-32. Best as combined with Roy would thus disclose a “*bi-directional data transfer fabric*” that allows for simultaneously data transfer. Ex. 1003, ¶177.

C. Claims 2 and 8 Are Obvious Over Best in view of Tsunoda, With or Without Roy

To the extent Best does not disclose claims 2 or 8, it would have been obvious to modify Best's data control/steering circuit in combination with an external interface (collectively the “*data manager*”) “*to control ... data error monitoring*” and “*data error correction in response to receiving at least one of a control signal and control information from the controller*” (claim 2) and “*data error correction*” (claim 8). Performing data error monitoring and correction in a

data transfer interconnect was well-known by the 831 Patent's priority date. For example, U.S. Patent Application Publication No. 2003/0028733 A1 to Tsunoda ("Tsunoda") (Feb. 6, 2003) (Ex. 1009) discloses controlling data transfer between volatile and nonvolatile memory by performing error correction whenever a data is read from Flash ("*data error monitoring*" and "*data error correction*"), which occurs "*in response to receiving at least one of a control signal and control information from the controller,*" *i.e.*, in response to receiving a command that reads data from Flash. Ex. 1009, ¶¶4, 97, 140, Figs. 18-20; Ex. 1003, ¶193.

Best and Tsunoda are analogous art to the 831 Patent because each is from the field of memory systems. Ex. 1001, 1:20-23; Ex. 1006, ¶1; Ex. 1009, 1:21-26; Ex. 1003, ¶194. It would have been obvious to one of ordinary skill to modify Best to perform error monitoring and correction when data is read from Flash, per Tsunoda. This would have been an arrangement of old elements (Best's memory, Tsunoda's ECC) with each performing the same function it had been known to perform and yielding no more than what one would expect from such an arrangement. Ex. 1003, ¶¶195, 253. One of ordinary skill would also have been motivated to include ECC functionality in Best to improve data security and protect against data loss. *Id.*; see Ex. 1009, ¶¶97, 140. Such a modification would have therefore been obvious.

D. Claims 5 and 12-14 Are Obvious Over Best in view of Roohparvar, With or Without Roy

1. Claims 5 and 12-14

To the extent Best does not disclose configuring an address space in response to / based on “*at least one of a received command from the memory controller*” as recited in claims 5 and 12 (upon which claims 13 and 14 depend), it would have been obvious to include that functionality in the system of Best. Configuring memory address spaces in response to external commands was well known in the art by the 831 Patent’s priority date. Ex. 1003, ¶216.

For example, U.S. Patent App. Pub. No. 2005/0273548 A1 to Roohparvar (“Roohparvar”) (Dec. 8, 2005) (Ex. 1019) discloses a memory system with user configurable density and performance options. Roohparvar teaches that Flash memory users typically have to choose between single bit cells (SBC) or multilevel cells (MLC) memories, and that both have benefits in different use cases. Ex. 1019, ¶¶2-7, 19-21. Roohparvar discloses a memory system that has user selectable MLC and SBC options, where the user can assign different memory densities to different memory blocks. *Id.*, ¶¶8-12. A user, for example, may input that a high priority performance parameter to set the flash memories to SBC in order to get the highest reliability possible for a given application. *Id.*, ¶¶35-36. “[T]he user determines the configuration of each block ... and stores this data into

the configuration register,” *id.*, ¶48, which is then used to configure the memory settings, *id.*, ¶¶42-49; Ex. 1003, ¶217. Best and Roohparvar are analogous art to the 831 Patent because each is from the field of memory systems. Ex. 1001, 1:20-23; Ex. 1006, ¶1; Ex. 1019, ¶1; Ex. 1003, ¶218.

It would have been obvious at the time of the 831 Patent’s priority to modify Best to allow for user configurable memory settings, as taught by Roohparvar. Best already discloses reprogramming or programmatically establishing the memory spaces. Ex. 1006, ¶¶14, 28. Roohparvar teaches the functionality that allows a user to configure settings stored on a memory device. Ex. 1019, ¶¶42-49. One of ordinary skill would have been motivated to allow user configuration of Best’s memory spaces for the same reasons Roohparvar identifies: not every application has the same requirements, and the optimal setting for one application is not the optimal setting for another. Ex. 1019, ¶36. Such a combination would also have been an arrangement of old elements (Best’s hybrid memory, Roohparvar’s configurable memory settings) with each performing the same function it had been known to perform and yielding no more than what one would expect from such an arrangement, *i.e.*, Best’s system with user configurable memory settings. Ex. 1003, ¶219. Such a modification would have therefore been obvious.

2. Claim 14

To the extent Best does not disclose “*configur[ing] the memory space... using partitioning instructions that are application-specific,*” such a feature would have been obvious. As explained above regarding claim 5, customizing memory settings for specific applications was well-known in the art. Ex. 1003, ¶¶216-217. Roohparvar discloses a user selecting a memory configuration based on the application she is using, *e.g.*, photos require more density so the user would choose a high density setting, or code requires more reliability so the user would choose a lower density but higher reliability setting. Ex. 1019, ¶36. As explained above, it would have been obvious to modify Best to incorporate user configurable memory settings such as application specific memory settings, as taught by Roohparvar, thus disclosing this feature. Ex. 1003, ¶¶216-219, 279.

E. Claim 15 Is Obvious Over Best in view of Bonella, With or Without Roy

1. “Operating the Volatile Memory Subsystem at a First Clock Frequency”

Claim 15 recites “[*t*]he method of claim 7, further comprising: operating the volatile memory subsystem at a first clock frequency when the memory module is in a first mode of operation in which data is communicated between the volatile memory subsystem and the memory controller.”

Best discloses “*operating the volatile memory subsystem at a first clock frequency.*” Best explains that the volatile memory die (“*volatile memory subsystem*”) can be synchronous or asynchronous, and explains that a conventional synchronous DRAM interface relies on “clock and clock-enable signals.” Ex. 1006, ¶14. One of ordinary skill in the art would have understood this disclosure to mean that a synchronous DRAM is “*operat[ed] ... at a first clock frequency*” because it is operated synchronously by a clock signal. This corresponds to the usual operation of DRAM memory, as would have been well-known by one of ordinary skill in the art. Ex. 1003, ¶282. For example, according to the DDR2 specification at the time of Best’s priority date, a conforming DDR2 DRAM device receives two input clock signals. DDR2 SDRAM Specification, JESD79-2B (Jan. 2005) (Ex. 1017), 6, 29. A skilled artisan would thus understand Best’s DRAM to be operated by a clock signal at a particular frequency. Ex. 1003, ¶283.

Best also discloses “*a first mode of operation in which data is communicated between the volatile memory subsystem and the memory controller.*” As explained above with regarding claim 5, Best discloses a “Shadow Operation” embodiment where a portion of the DRAM is used as a write buffer for the non-volatile memory. Ex. 1006, ¶24; Ex. 1003, ¶¶207-214. “[D]ata is stored in a fast-access volatile storage die” and later “a write-back trigger is detected within the shared interface circuitry,” such as a “power-loss” event. Ex. 1006, ¶¶25-26. Once the

power-loss event is detected, “internal data transfer operations are performed to transfer data ... from the DRAM to ... the NV memory.” *Id.*; Ex. 1003, ¶284.

Best therefore describes at least two modes of operation: (1) the operation of the system when no write-back trigger is detected (“*a first mode of operation*”), and (2) the operation of the system when a power-loss event write-back trigger is detected (“*a second mode of operation*”). Ex. 1006, ¶¶25-26. Both modes would normally be operated at the DRAM’s frequency of operation. Ex. 1003, ¶285.

Best’s “*first mode of operation*” (*i.e.*, the operation of the system when no write-back trigger is detected) is one “*in which data is communicated between the volatile memory subsystem and the memory controller*” because initially “data is stored in a fast-access volatile storage die” and only later after a write-back trigger is written to non-volatile memory. Ex. 1006, ¶25. Ex. 1003, ¶286.

To the extent one might argue that the Board’s interpretation of “*clock frequency*” requires a particular numeric value (*i.e.*, as in number of cycles per second), it would have been obvious to operate Best’s DRAM at any of the frequencies described by the DDR DRAM and DDR2 DRAM standards, *e.g.*, 67-400MHz. Ex. 1003, ¶287 (discussing Ex. 1018, 58; Ex. 1017, 69). Best teaches that “any [] type of volatile storage technology may be used” Ex. 1006, ¶13. It would have been obvious to use one of the known DDR or DDR2 DRAM modules in Best’s system, and to subsequently clock the module at its standards-specified

rate. Ex. 1003, ¶288. To do so would have been merely the use of a known structure (a standard DRAM module and clock frequency) for its known use (operation as a DRAM module) to achieve a predictable result (operating a DRAM module at a standard clock frequency in Best's system). *Id.* A skilled artisan would have been motivated to operate DDR2 DRAM at those frequencies in order to ensure proper operations of the memory. *Id.* Therefore, operating Best's volatile memory “*at a first clock frequency*” would have been obvious. Therefore, Best discloses this claim element.

2. “Operating the Non-Volatile Memory Subsystem at a Second Clock Frequency”

Claim 15 recites “*operating the non-volatile memory subsystem at a second clock frequency when the memory module is in a second mode of operation in which data is communicated between the volatile memory subsystem and the non-volatile memory subsystem.*”

Best describes at least two modes of operation: (1) the operation of the system when no write-back trigger is detected (“*a first mode of operation*”), and (2) the operation of the system when a power-loss event write-back trigger is detected such as a power-loss (“*a second mode of operation*”). Ex. 1006, ¶¶25-26. Ex. 1003, ¶¶282-286, 291. Best's “*second mode of operation*” (*i.e.*, the operation of the system when a power-loss write-back trigger) involves “*operating the non-*

volatile memory subsystem” (Best’s NV memory) such that “*data is communicated between the volatile memory subsystem and the non-volatile memory subsystem.*”

Ex. 1006, ¶25. Ex. 1003, ¶292.

Best does not explicitly disclose “*operating the non-volatile memory subsystem*” (the Flash memory) “*at a second clock frequency*” during this second mode, but it would have been obvious to include that functionality in the system of Best. Ex. 1003, ¶293.

Synchronous Flash memory interfaces that were operated by clock signals were well-known in the art by the 831 Patent’s priority date. Ex. 1003, ¶294. One prior art Flash implementation is described by US Patent No. 6,026,465 to Mills (“Mills”) (Ex. 1010), which describes a synchronous flash interface: “FIG. 6 illustrates a block diagram of a synchronous flash interface (SFI) flash memory integrated circuit 600 that incorporates a complete synchronous flash interface in a single flash memory chip.” Ex. 1010, 16:60-63. The synchronous flash interface causes the Flash memory system to “*operat[e] ... at a second clock frequency*”: “A clock input is a part of the interface. ... All the external operations of the device are synchronized to the rising edge of the clock. ... The user can cycle the device at frequencies as high as 33 MHz.” Ex. 1010, 17:10-25; Ex. 1003, ¶294.

Mills teaches that this synchronous operation is used for both read operations and write operations: “When SFI is enabled, interlace control 670 and

[bank] select logic 674 operate to interlace read (*and write*) operations between flash bank A 610 and a flash bank B 620.” Ex. 1010, 17:28-39 (emphasis added). Because “the device is interleaved internally,” “this interface creates an average access time for sequential read accesses that is significantly less than the access time of an asynchronous flash device.” Ex. 1010, 17:1-9. Mills therefore teaches one of ordinary skill in the art a synchronous Flash interface where read and write operations are synchronized to the rising edge of a clock signal provided to the device and operating at a particular frequency (*i.e.*, “operating [a] non-volatile memory subsystem at a second clock frequency”). Ex. 1003, ¶295. To the extent one might argue that the Board’s interpretation of “clock frequency” requires a particular numeric value (*i.e.*, as in number of cycles per second), Mills discloses the use of his Flash memory at, for example, a frequency “as high as 33 MHz”. Ex. 1010, 17:10-25; Ex. 1003, ¶301.

Best and Mills are analogous art to the 831 Patent because each is from the field of memory systems. Ex. 1001, 1:20-23; Ex. 1006, ¶1; Ex. 1010, 6:57-59; Ex. 1003, ¶296. As of the priority date of the 831 Patent, it would have been obvious to one of ordinary skill in the art to employ a synchronous flash memory, such as disclosed in Mills, in the system of Best because to do so would have been merely an arrangement of old elements with each performing the same function it had

been known to perform and yielding no more than what one would expect from such an arrangement, *i.e.*, the non-volatile storage of data. Ex. 1003, ¶297.

As combined, Best's Flash interface would conform to Mills' synchronous Flash protocol and include a separate clock signal that controls read and write operations. Ex. 1010, 17:10-25, 17:28-39. This means that Best's Flash memory would operate at "*a second clock frequency,*" *i.e.*, the frequency of the synchronous Flash interface clock up to 33 MHz (Ex. 1010, 17:10-25), during its operation, including when a write-back trigger is detected and data is written from DRAM to Flash ("*when the memory system is in a second mode of operation in which data is communicated between the volatile memory subsystem and the non-volatile memory subsystem*"). Ex. 1003, ¶298.

A skilled artisan would have been motivated to make such a combination because, as Mills explains, a synchronous flash interface "creates an average access time for sequential read accesses that is significantly less than the access time of an asynchronous flash device." Ex. 1010, 17:6-9. In the context of Best, restoring data from the non-volatile flash memory would therefore have been faster by use of a synchronous flash memory, and reduced sequential read access times during other operations or uses of Best's Flash memory, motivating one of ordinary skill in the art to use a synchronous interface generally. Ex. 1003, ¶299.

Moreover, Best discloses that his “non-volatile storage IC 101 ... is implemented by a Flash memory die ... of either the NAND-Flash or NOR-Flash varieties, though any other electrically-erasable or electrically-alterable storage technology may alternatively be used.” Ex. 1006, ¶13. One of ordinary skill in the art would have therefore understood Best to suggest modification to work with any known Flash interface, including Mills’ synchronous Flash interface. Ex. 1003, ¶300. Therefore, Best renders obvious this claim element.

3. “Operating the Volatile Memory Subsystem at a Third Clock Frequency”

Claim 15 recites “*operating the volatile memory subsystem at a third clock frequency when the memory module is in the second mode of operation, the third clock frequency being less than the first clock frequency.*”

As explained above, Best discloses or renders obvious “*operating the volatile memory subsystem*” at a “*first clock frequency,*” but does not explicitly disclose “*operating the volatile memory subsystem at a third clock frequency when the memory module is in the second mode of operation, the third clock frequency being less than the first clock frequency.*” Ex. 1003, ¶¶282-288. Doing so would have been obvious at the time of the 831 Patent’s priority date for two reasons: (1) reducing power during volatile to non-volatile flush operations prompted by a power loss was a well-known technique, and (2) one known way to reduce the

power consumption of DRAM devices was to reduce their frequency of operation.

It would have therefore been obvious to operate Best's DRAM "*at a third clock frequency*" that is "*less than the first clock frequency*" when performing Best's flush operation (*i.e.*, the "*second mode*"). Ex. 1003, ¶304.

First, reducing power consumption during memory backup operations initiated in response to an interrupted power supply was known by the priority date of the 831 Patent, as were its advantages. For example, U.S. Patent Application Publication 2006/0212651 A1 ("Ashmore") (Ex. 1011) "reduc[es] battery power consumption during a main power loss to reduce the likelihood of loss of user write-cached data in a write-caching mass storage controller." Ex. 1011, ¶9; Ex. 1003, ¶305. As another example, U.S. Patent No. 7,421,552 to Long ("Long") (Ex. 1012), discloses that "if there is a loss of primary power 34, ... provid[ing] a significantly slower clock signal to ... the controller 40 while the controller 40 moves data from the volatile-memory storage cache 42 to the flash-based memory vault 44." Ex. 1012, 4:54-64. "As a result, less power is consumed thus enabling the use of a smaller-sized backup power source 28 (e.g., a relatively small battery)." *Id.* Reducing power consumption during memory backup operations initiated in response to an interrupted power supply was therefore well-known and would reduce the likelihood of data loss or reduce the necessary size for the backup power source. Ex. 1003, ¶306.

One of ordinary skill in the art would have been motivated to reduce the power consumption during Best's write flushing in response to a power loss. A skilled artisan would have been motivated to perform this power reduction technique for all the reasons that were known in the art: *e.g.*, decreasing the risk of data loss due to insufficient backup power (Ex. 1011, ¶7) and enabling the use of a smaller-sized backup power source (Ex. 1012, 4:54-64). Reducing power consumption during write flushing in response to a power loss would also have been the arrangement of old elements, each performing the same function it had been known to perform, in a way that yields no more than one of ordinary skill in the art would expect from such an arrangement (reducing power consumption during a power loss event, as suggested by Long and Ashmore). Ex. 1003, ¶307.

Second, it was known that one way to reduce the power consumption of a DRAM was to reduce the operating frequency of the DRAM, *i.e.*, by lowering its clock frequency. For example, U.S. Patent Application Publication No. 2007/0136523 A1 to Bonella ("Bonella") (filed Dec. 8, 2006) (Ex. 1013) explains that one way to reduce the power consumption of the memory module is to slow or reduce the operating frequency of the DRAM. Ex. 1013, ¶¶49-50; Ex. 1003, ¶308. Best and Bonella are both analogous art to the 831 Patent because each is from the field of memory systems. Ex. 1001, 1:20-23; Ex. 1006, ¶1; Ex. 1013, ¶2; Ex. 1003, ¶309.

One of ordinary skill in the art would also have found it obvious to reduce power consumption during Best's write flushing in response to a power loss using any known or conventional means, and would have also considered power consumption reduction techniques other than those of Ashmore and Long to obtain the same benefits, including those described in Bonella. Bonella describes reducing the operational frequency of the DRAM in order to reduce the power consumption generally. Ex. 1013, ¶50. An ordinary artisan would have found it obvious to perform Bonella's power reduction technique (DRAM operational frequency reduction) during the performance of Best's write flushing in response to a power loss, in view of the specific motivations evidenced by Ashmore and Long. Ex. 1003, ¶310.

Bonella also motivates the use of his technique because it results in "major power savings." Ex. 1013, ¶50. Using Bonella's power reduction technique during Best's write flushing would also have been the arrangement of old elements, each performing the same function it had been known to perform, in a way that yields no more than one of ordinary skill would expect from such an arrangement (reducing power consumption during a power loss event, as suggested by Long and Ashmore). Ex. 1003, ¶311.

Reducing the DRAM's operational frequency in such a situation would not have caused any performance issues because the transfer of data between DRAM

and FLASH could not occur faster than the speed of the FLASH, which in the case of Best would be substantially slower than the DRAM. Bonella teaches that various levels of a computer's memory hierarchy have different performance levels, and lists DRAM as higher performance than Flash memory. Ex. 1013, ¶65, Fig. 3; *see also* Ex. 1009, ¶97; Ex. 1003, ¶312. One of ordinary skill in the art would have understood that reducing the DRAM operating frequency in this way would not cause performance issues. Ex. 1003, ¶312. Such a modification would have therefore been obvious.

F. Claim 15 Is Obvious Over Best in view of Bonella and Ashmore, With or Without Roy

To the extent Best and Bonella would not render obvious claim 15 in view of the knowledge of one of ordinary skill as explained above (§V.E), it would have been obvious to use Bonella's power saving mode during performance of Best's power loss write flushing (“*second mode of operation*”) in view of the specific teachings of Ashmore. Best, Bonella, and Ashmore are analogous art to the 831 Patent because each is from the field of memory systems. Ex. 1001, 1:20-23; Ex. 1006, ¶1; Ex. 1013, ¶2; Ex. 1011, ¶1; Ex. 1003, ¶312.

Ashmore would have motivated one of ordinary skill to consider other power consumption reduction techniques to obtain the same benefits, *e.g.*, “to reduce the likelihood of user data loss” (Ex. 1011, ¶7), including those described in

Bonella. Ashmore thus motivates one of ordinary skill to reduce the power consumption in Best's system during the power loss event write flushing using any known or conventional means, including Bonella's DRAM frequency reduction technique. Ex. 1003, ¶314. Bonella's motivates the use of his own technique because it "can result in major power savings," Ex. 1013, ¶50, and using it would have been the arrangement of old elements, each performing the same function it had been known to perform, in a way that yields no more than one of ordinary skill in the art would expect from such an arrangement (reducing power consumption during a power loss event, as suggested by Ashmore). Ex. 1003, ¶314. Such a modification would have therefore been obvious.

VI. CONCLUSION

For the foregoing reasons, the challenged claims are unpatentable.

Dated: January 17, 2017

Respectfully Submitted,

/Joseph Micallef/
Joseph A. Micallef
Registration No. 39,772
Sidley Austin LLP
1501 K Street NW
Washington, DC 20005

CERTIFICATE OF COMPLIANCE

I hereby certify that this petition complies with the type-volume limitations of 37 C.F.R. § 42.24, because it contains 13,981 words (as determined by the Microsoft Word word-processing system used to prepare the petition), excluding the parts of the petition exempted by 37 C.F.R. § 42.24.

Dated: January 17, 2017

Respectfully Submitted,

/Joseph Micallef/
Joseph A. Micallef
Registration No. 39,772
Sidley Austin LLP
1501 K Street NW
Washington, DC 20005

Petition for *Inter Partes* Review of U.S. Patent No. 8,874,831

**PETITION FOR INTER PARTES REVIEW
OF U.S. PATENT NO. 8,874,831**

Attachment A:

Proof of Service of the Petition

Petition for *Inter Partes* Review of U.S. Patent No. 8,874,831

CERTIFICATE OF SERVICE

I hereby certify that on this 17th day of January, 2017, a copy of this Petition, including all attachments, appendices and exhibits, has been served in its entirety by overnight mail on the following counsel of record for patent owner:

Nixon Peabody LLP
P.O. Box 26769
San Francisco, CA 94126

Dated: January 17, 2017

Respectfully Submitted,

/Joseph Micallef/
Joseph A. Micallef
Registration No. 39,772
Sidley Austin LLP
1501 K Street NW
Washington, DC 20005

Petition for *Inter Partes* Review of U.S. Patent No. 8,874,831

**PETITION FOR INTER PARTES REVIEW
OF U.S. PATENT NO. 8,874,831**

Attachment B:

List of Evidence and Exhibits Relied Upon in Petition

Petition for *Inter Partes* Review of U.S. Patent No. 8,874,831

Exhibit #	Reference Name
1001	U.S. Patent No. 8,874,831 (“831 Patent”)
1002	File History of U.S. Patent No. 8,874,831 (“831 FH”)
1003	Declaration of Ron Maltiel
1004	Curriculum Vitae of Ron Maltiel
1005	U.S. Provisional Pat. App. No. 60/941,586 (“586 Application”)
1006	U.S. Patent App. Pub. No. 2010/0110748 (“ <u>Best</u> ”)
1007	U.S. Provisional Patent App. No. 60/912,321 (“’321 Application”)
1008	U.S. Patent No. 6,065,092 (“ <u>Roy</u> ”)
1009	U.S. Patent App. Pub. No. 2003/0028733 (“ <u>Tsunoda</u> ”)
1010	U.S. Patent No. 6,026,465 to <u>Mills</u>
1011	U.S. Patent App. Pub. No. 2006/0212651 (“ <u>Ashmore</u> ”)
1012	U.S. Patent No. 7,421,552 (“ <u>Long</u> ”)
1013	U.S. Patent App. Pub. No. 2007/0136523 (“ <u>Bonella</u> ”)
1014	U.S. Patent No. 8,301,833 (“833 Patent”)
1015	Microsoft Computer Dictionary (5th Ed.) (2002)
1016	Intel, <i>1.8 Volt Intel StrataFlash(R) Wireless Memory</i> (L18) (Apr. 2003)
1017	JEDEC Standard, DDR2 SDRAM Specification, JESD79-2B (Jan. 2005)
1018	JEDEC Standard, DDR SDRAM Specification, JESD79 (Jun. 2000)
1019	U.S. Patent App. Pub. No. 2005/0273548 (“ <u>Roohparvar</u> ”)