UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

PHISON ELECTRONICS CORPORATION,

Petitioner

v.

VERVAIN, LLC,

Patent Owner

Inter Partes Review Case No. IPR2025-00212

Patent 8,891,298

PETITION FOR *INTER PARTES* REVIEW OF U.S. PATENT NO. 8,891,298

TABLE OF CONTENTS

TABLE OF AUTHORITIES

CASES

STATUTES

OTHER AUTHORITIES

LIST OF EXHIBITS

Exhibit	Description
1001	U.S. Patent No. 8,891,298, titled "LIFETIME MIXED LEVEL NON-VOLATILE MEMORY SYSTEM," to Rao
1002	Declaration of Carl Sechen, Ph.D.
1003	U.S. Provisional Application No. 61/509,257, titled "IMPROVED LIFETIME MIXED LEVEL NAND FLASH SYSTEM" by Rao
1004	U.S. Patent No. 7,855,916, titled "NONVOLATILE MEMORY SYSTEMS WITH EMBEDDED FAST READ AND WRITE MEMORIES," to Rao
1005	*Intentionally omitted*
1006	Excerpts from prosecution history of '298 patent
1007-24	*Intentionally omitted*
1025	Excerpts from Rino Micheloni, et al., INSIDE NAND FLASH MEMORIES (Springer 2010)
1026	Excerpts from NONVOLATILE MEMORY TECHNOLOGIES WITH EMPHASIS ON FLASH: A COMPREHENSIVE GUIDE TO UNDERSTANDING AND USING NVM DEVICES (Joe E. Brewer & Manzur Gill eds. Wiley-IEEE Press 2008)
1027	Excerpts from MICROSOFT COMPUTER DICTIONARY (5th ed. 2002)
1028-29	*Intentionally omitted*
1030	Yu Cai et al., *Program Interference in MLC NAND Flash Memory: Characterization, Modeling and Mitigation* (IEEE 2013)

1031	Kang-Deog Suh et al., *A 3.3 V 32 Mb NAND Flash Memory with Incremental Step Pulse Programming Scheme*, 30 IEEE J. Solid State Circuits 1149 (Nov. 1995)
1032	Eitan Yaakobi et al., *Error Characterization and Coding Schemes for Flash Memories* (2010 IEEE Globecom Workshops, Miami, FL at 1856-60)
1033	Evans, Chris, "Enterprise MLC: How flash vendors are boosting MLC write endurance" (June 3, 2011)
1034	Nelson Duann, SLC & MLC Hybrid (Silicon Motion Flash Memory Summit, Santa Clara, CA, Aug. 12, 2008)
1035	Ashish Jagmohan et. al., *Adaptive Endurance Coding for NAND Flash* (2010 IEEE Globecom Workshops, Miami, FL at 1841-45)
1036	Sungjin Lee et al., *FlexFS: A Flexible Flash File System for MLC NAND Flash Memory* (2009 USENIX Annual Technical Conference, San Diego, CA, June 14-19, 2009)
1037	Minyoung Sung & Kanghee Kim, *Asymmetric Flash Volume Management*, 58 *IEEE Transactions on Consumer Electronics* 455 (May 2012)
1038	Greg Atwood et al., *Intel StrataFlash*™ Memory Technology Overview at 4, Intel Technology J. (Q4 1997)
1039	Taehee Cho et al., "A 3.3 V 1 Gb multi-level NAND flash memory with non-uniform threshold voltage distribution," *2001 IEEE International Solid-State Circuits Conference. Digest of Technical Papers,* ISSCC (Cat. No. 01CH37177) (2001)
1040	U.S. Patent No. 6,456,528, titled for "SELECTIVE OPERATION OF A MULTI-STATE NON-VOLATILE MEMORY SYSTEM IN A BINARY MODE" to Chen

1041	U.S. Patent No. 8,078,794, titled "HYBRID SSD USING A COMBINATION OF SLC AND MLC FLASH MEMORY ARRAYS," to Lee et al.
1042	U.S. Patent Appl. Pub. No. 2008/0140918, titled "HYBRID NON-VOLATILE SOLD STATE MEMORY SYSTEM" by Sutardja
1043	U.S. Patent Appl. Pub. No. 2009/0327591, titled "SLC-MLC COMBINATION FLASH STORAGE DEVICE," by Moshayedi
1044	U.S. Patent Appl. Pub. No. 2010/0172179, titled "SPARE BLOCK MANAGEMENT OF NON-VOLATILE MEMORIES," by Gorobets et al.
1045	U.S. Patent No. 8,634,240, titled "NON-VOLATILE MEMORY AND METHOD WITH ACCELERATED POST-WRITE READ TO MANAGE ERRORS," to Gavens et al.
1046	U.S. Patent No. 8,806,301, titled "DATA WRITING METHOD FOR A FLASH MEMORY, AND CONTROLLER AND STORAGE SYSTEM USING THE SAME," to Yu et al.
1047	U.S. Patent Appl. Pub. No. 2010/0115192, titled "WEAR LEVELING METHOD FOR NON-VOLATILE MEMORY DEVICE HAVNG SINGLE AND MULTI LEVEL MEMORY CELL BLOCKS," by Lee
1048	U.S. Patent Appl. Pub. No. 2011/0271043, titled "SYSTEM AND METHOD FOR ALLOCATING AND USING SPARE BLOCKS IN A FLASH MEMORY," to Segal et al.
1049	U.S. Patent Appl. Pub. No. 2011/0153912, titled "MAINTAINING UPDATES OF MULTI-LEVEL NON-VOLATILE MEMORY IN BINARY NON-VOLATILE MEMORY," by Gorobets et al.
1050	U.S. Patent No. 5,930,167, titled "MULTI-STATE NON-VOLATILE FLASH MEMORY CAPABLE OF BEING ITS OWN TWO STATE WRITE CACHE," to Lee et al.

1051	U.S. Patent Appl. Pub. No. 2010/0172180, titled "NON-VOLATILE MEMORY AND METHOD WITH WRITE CACHE PARTITIONING," by Paley et al.
1052	*Intentionally omitted*
1053	U.S. Patent Appl. Pub. No. 2009/0268513, titled "MEMORY DEVICE WITH DIFFERENT TYPES OF PHASE CHANGE MEMORY," by De Ambroggi et al.
1054	U.S. Patent Appl. Pub. No. 2009/0307418, titled "MULTI-CHANNEL HYBRID DENSITY MEMORY STORAGE DEVICE AND CONTROL METHOD THEREFOR," by Chen et al.
1055	U.S. Patent Appl. Pub. No. 2010/0058018, "MEMORY SCHEDULER FOR MANAGING INTERNAL MEMORY OPERATIONS," by Kund et al.
1056-60	*Intentionally omitted*
1061	Duane H. Oto et al., *High-Voltage Regulation and Process Considerations for High-Density 5 V-Only E^2PROM's*, SC-18 IEEE J. Solid State Circuits 532 (Oct. 1983)
1062	Gheorghe Samachisa et al., *A 128K Flash EEPROM Using Double-Polysilicon Technology*, SC-22 IEEE J. Solid State Circuits 676 (Oct. 1987)
1063	Masaki Momodomi et al., *An Experimental 4-Mbit CMOS EEPROM with a NAND-Structured Cell*, 24 IEEE J. Solid State Circuits 1238 (Oct. 1989)
1064	Tae-Sun Chung et al., *A survey of Flash Translation Layer*, 55 J. Systems Architecture 332-43 (2009)
1065	Yoshiba Iwata et al., *A High-Density NAND EEPROM with Block-Page Programming for Microcomputer Applications*, 25 IEEE J. Solid-State Circuits 417 (Apr. 1990)

1066	Simona Boboila and Peter Desmoyers, *Write Endurance in Flash Drives: Measurements and Analysis*, Usenix Conference (Feb. 2010)
1067	SiliconSystems, *Increasing Flash SSD Reliability*, StorageSearch.com (April 2005)
1068	Taehee Cho et al., *A Dual-Mode NAND Flash Memory: 1-Gb Multilevel and High-Performance 512-Mb Single Level Modes*, 36 IEEE J. Solid State Circuits 1700 (Nov. 2001)
1069	Ken Takeuchi, *A Multipage Cell Architecture for High-Speed Programming Multilevel NAND Flash Memories*, 33 IEEE J. Solid-State Circuits 1228 (Aug. 1998)
1070-79	*Intentionally omitted*
1080	Docket list for *Vervain, LLC v. Phison Electronics Corp.*, No. 1:24-cv-00259 (W.D. Tex., filed March 8, 2024)
1081	Docket list for *Vervain, LLC v. Kingston Technology Company, Inc*., No. 1:24-cv-00254 (W.D. Tex. filed March 7, 2024)
1082	First Amended Complaint, Doc. No. 9 in *Vervain, LLC v. Phison Electronics Corp.*, No. 1:24-cv-00259 (W.D. Tex.)
1083	Signed Agreed Scheduling Order, Doc. No. 16 in *Vervain, LLC v. Phison Electronics Corp.*, No. 1:24-cv-00259 (W.D. Tex.)

I. INTRODUCTION:

Phison Electronics Corporation ("Phison" or "Petitioner") requests *inter partes* review ("IPR") of all claims 1-11 (the "Challenged Claims") of U.S. Patent No. 8,891,298, entitled "LIFETIME MIXED LEVEL NON-VOLATILE MEMORY SYSTEM," to Rao (the "'298 patent") (Ex. 1001), which is assigned to Vervain, LLC ("Vervain" or "Patent Owner"). The '298 patent, which issued on Nov. 18, 2014, from U.S. Appl. No. 13/455,267 filed Apr. 25, 2012, incorporated by reference Dr. Rao's U.S. Patent No. 7,855,916,[1] and claimed priority from U.S. Provisional App. No. 61/509,257,[2] filed on July 19, 2011 ("Provisional," Ex. 1003).

The '298 patent and seven other patents, U.S. Patent Nos. 9,196,385 ("the '385 patent"), 9,997,240 ("the '240 patent"), 10,950,300 ("the '300 patent"), 11,830,546 ("the '546 patent"), 11,854,612 ("the '612 patent"), 11,967,369 and

---

[1] ('298 1:12-16, "Rao '916," Ex. 1004.) Rao '916, entitled "NONVOLATILE MEMORY SYSTEMS WITH EMBEDDED FAST READ AND WRITE MEMORIES," incorporated along with a then-pending application by Dr. Rao of the same title and specification ('298 1:16-21).

[2] ('298 1:7-12 [also incorporated by reference].) Note that the '298 application's title omitted "improved" from the Provisional title.

11,967,370, all with the same title and based on the same Apr. 25, 2012, specification and claiming July 19, 2011, priority, are currently asserted (First Amended Complaint, Ex. 1082) by Patent Owner against Petitioner in *Vervain, LLC v. Phison Electronics Corp.*, No. 1:24-cv-00259 (W.D. Tex. filed March 8, 2024).[3] Petitioner has petitioned for post-grant reviews of the '546 patent in pending *Phison*

[3] All are also asserted against a customer of Phison in *Vervain, LLC v. Kingston Technology Company, Inc*., *Kingston Digital, Inc., and Kingston Technology Corp*., No. 1:24-cv-00254 (W.D. Tex. filed March 7, 2024). The first four patents were earlier asserted in *Vervain, LLC v. Micron Technology, Inc.,* No. 6:21-cv-00487 (W.D. Tex. filed May 10, 2021, settled April 21, 2023) and *Vervain, LLC v. Western Digital Corp.*, No. 6:21-cv-00488 (W.D. Tex. filed May 10, 2021, settled Aug. 8, 2023). Micron petitioned this Board for *inter partes* review in *Micron Technology, Inc. v. Vervain, LLC*, IPR2021-01547 ('298 patent), -01548 ('385 patent), -01549 ('240 patent), and -01550 ('300 patent), which were instituted on April 8 and 11, 2022, and terminated on April 4, 2023, on settlement after hearing but prior to final written decision.

Electronics Corp. v. Vervain, LLC, PGR2024-00047[4] and of the '612 patent in *Phison Electronics Corp. v. Vervain, LLC*, PGR2024-00048.[5]

This Petition demonstrates by a preponderance of the evidence that the Challenged Claims are unpatentable as claiming a long-known hybrid memory system of distinct Single-Level Cell ("SLC") and Multi-Level Cell ("MLC") NAND flash memory modules "adapted" to perform two optional functions of directing data preferentially to the SLC module under two circumstances: (1) failure of a "data integrity test" (7:26-13); and (2) "allocation" by "transfer" of the "contents" of "blocks" in the MLC and SLC modules "that receive the most frequent writes" (8:1-9). These functions were suggested in the specification as improving "lifetime (endurance)" but those functions in the detail recited had long been disclosed in the extensive NAND flash memory art which included many error management and wear leveling techniques for avoiding premature failure of hybrid systems. For Ground 1 of invalidity, Petitioner starts with a published hybrid memory system with multiple error management techniques and shows that a POSITA would apply known wear leveling techniques for the common objective of avoiding premature failure, thereby showing claim 1 obvious in view of the knowledge of the POSITA.

---

[4] (Accorded filing date Aug. 27, 2024 [Paper No. 5].)

[5] (Accorded filing date Sept. 25, 2024 [Paper No. 5].)

Conversely, for Grounds 2 and 3, Petitioner starts with two published hybrid memory systems with multiple wear leveling techniques and shows that a POSITA would apply known error management techniques for the common objective of avoiding premature failure, thereby showing claim 1 obvious in view of the knowledge of the POSITA. Under each Ground, the remaining Challenged Claims are also obvious over the asserted document in view of that document and other knowledge of the POSITA.

II. MANDATORY NOTICES UNDER 37 C.F.R. § 42.8:

A. Real Party-In-Interest Under 37 C.F.R. § 42.8(b)(1)):

Petitioner Phison Electronics Corporation is the real party-in-interest.

B. Related Matters Under 37 C.F.R. § 42.8(b)(2):

As set forth further at note 3 supra, the '298 patent (along with seven other patents based on the same specification) is currently asserted by Vervain against Phison and its customer (separately) in the Western District of Texas. In the Phison case a claim construction hearing is tentatively scheduled for December 5, 2024, and jury selection is tentatively scheduled for December 15, 2025 (*Vervain, LLC v. Phison Electronics Corp.*, No. 1:24-cv-00259, Doc. No. 16 (W.D. Tex. July 15, 2024) (Ex. 1083). Resolution of this Petition may simplify that litigation and promote settlement.

C. Lead and Back Up Counsel Under 37 C.F.R. § 42.8(b)(3):

Pursuant to 37 C.F.R. § 42.8(b)(3), lead counsel for this Petition is Hsuanyeh Chang, PhD (Reg. No. 73,431) and back-up counsel are Stephen Y. Chow (Reg. No. 31,338), Douglas E. Chin (Reg. No. 66,713), and Peter Yi (Reg. No. 61,790). Pursuant to 37 C.F.R. § 42.10(b), Petitioner has filed a power of attorney designating the above-identified counsel.

D. Service Information Under 37 C.F.R. § 42.8(b)(4):

Service information under 37 C.F.R. § 42.8(b)(4) for the Petition is as follows:

Lead Counsel	Back-Up Counsel
Hsuanyeh Chang, PhD Reg. No. 73,431 HSUANYEH LAW GROUP PC 11 Beacon Street, Suite 900 Boston MA 02108 (617) 886-9088 (Phone) hsuanyeh@hsuanyeh.com	Stephen Y. Chow Reg. No. 31,338 HSUANYEH LAW GROUP PC 11 Beacon Street, Suite 900 Boston MA 02108 (617) 886-9288 (Phone) Stephen.Y.Chow@hsuanyeh.com
	Douglas E. Chin Reg. No. 66,713 HSUANYEH LAW GROUP PC 11 Beacon Street, Suite 900 Boston MA 02108 (617) 886-9488 (Phone) Doug.Chin@hsuanyeh.com
	Peter Yi Reg. No. 61,790 HSUANYEH LAW GROUP PC 11 Beacon Street, Suite 900 Boston MA 02108 (617) 886-9188 (Phone) Peter.Yi@hsuanyeh.com

Petitioner consents to electronic service at the email addresses above.

III. ADDITIONAL REQUIREMENTS:

A. Payment of Fees Under 37 C.F.R. § 42.15:

The required fees are submitted from Deposit Account No. 50-6685 (Order No. 1280-0009). If any additional fees are due at any time during this proceeding, the Office is authorized to charge such fees to Deposit Account No. 50-6685 (Order No. 1280-0009).

B. Timing Under 37 C.F.R. § 42.102:

The present petition for *inter partes* review is filed more than nine months after the grant of the '298 patent.

C. Grounds for Standing Under 37 C.F.R. § 42.104(a):

Petitioner certifies that: (1) the '298 patent is eligible for *inter partes* review; and (2) Petitioner is not barred or estopped from requesting *inter partes* review of any claims of the '298 patent on the grounds identified herein.

IV. NAND FLASH TECHNOLOGY BACKGROUND

Petitioner's expert, Dr. Carl Sechen, has provided in his declaration ("Sechen," Ex. 1002) a deep analysis of the state of NAND flash technology implicitly relied on by the '298 patent (Sechen Section V), the embodiments admitted or advanced in the '298 specification (Sechen Section VII), the evolution of the '298 claims through prosecution and their remaining defects (Sechen Section VIII), and invalidity of those claims over the prior art in view of the POSITA's

knowledge of the art (Sechen Sections IX through XII). Following is a summary of the NAND flash memory technology underlying the '298 patent base claim to a "system for storing data," based on known MLC and SLC NAND flash modules, and a known controller "adapted" to perform essential NAND flash mapping and optional, also known, error management and wear leveling

A. Characteristic Nature of NAND Flash Memory

A **NAND flash cell** is type of transistor with an analog floating gate that can retain (store) a range of electrical charges associated with "threshold voltages" that, applied to a control gate, allows conduction of electrical current from drain to source. (Sechen ¶¶ 26 [Fig. 2.1(a)], 40 [Figures 6 and 7 from Atwood {Ex. 1038}] and 41.)

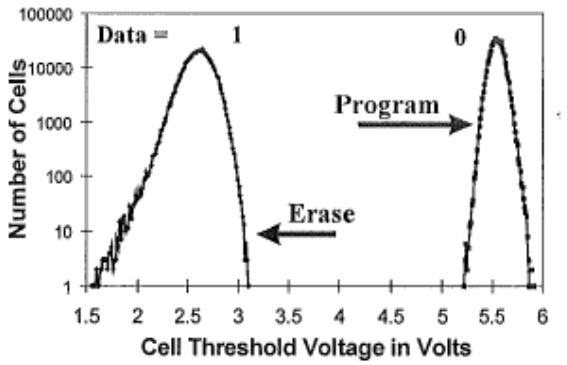

Figure 6: Single bit/cell array threshold voltage histogram

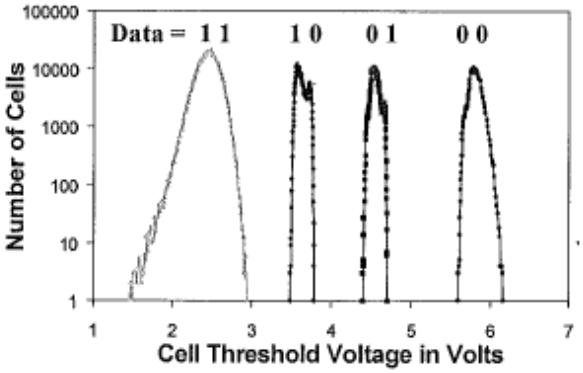

Figure 7: Two bit/cell array threshold voltage histogram

These threshold voltages/charges may be chosen to represent information to be stored, in earlier NAND memory operation, using two threshold voltage levels (Figure 6), the lower level to represent ("store") a logical "1" and the higher level to store a logical "0" of one bit of (binary) digital computer data. (*Id.*) The same NAND

storage cell can use four threshold voltage levels (Figure 7) to represent four values corresponding to two bits of binary computer data (which in NAND memory operation is sourced from two distinct groups, "pages," of logical data). (*Id.*) An *n* power of two (2^n: eight, sixteen, . . .) threshold levels may be used in NAND flash operation to represent corresponding bits of *n* NAND flash pages (*id.* ¶¶ 40 and 42-43). But a storage cell or group of storage cells do not constitute "memory," which requires circuitry for selecting, programming (writing) and reading the cell or group of cells (*id.* ¶¶ 26 and 29).

To form **NAND flash memory**, NAND storage cells are hard wired in strings (bit-line "columns") and word-line "rows" forming NAND "**pages**" each typically 2K storage cells **programmed page by page**. (*Id.* ¶ 26.) A defining characteristic of NAND flash is that it is hard wired such that "**blocks**" **of** 64 **pages** are **erased in unison** in required preparation for writing to pages in that block that are not empty or previously erased (*id.* ¶¶ 28-30). Thus, unlike magnetic hard disk drives (and the '298 patents' proposed alternative technologies) which allow "write in place" of data assigned to a particular physical location, updates to data (considered "logical") in NAND flash must be written to a new physical location that is empty or must be erased in preparation ("erase before write"). (*Id.*) To maintain logical visibility to the user (host computer) of a NAND flash memory device such as a Solid State Drive ("SSD") (shown at Sechen ¶ 30), there are maintained and updated tables in or

available to a **NAND flash device controller**, shown as a hardware module or chip on an SSD circuit board (*id.*) to "map" logical locations visible to the user to physical locations in the NAND flash memory array ("**L2P mapping**"), which update for each data write (*id.* ¶¶ 30 and 31). The physical instances of NAND flash pages and blocks determine the **Physical Block** Addresses ("PBA") mapped from the construct of **Logical Block** Addresses ("LBA"). (*Id.* ¶¶ 32-34) This mapping and remapping is typically done with a software **Flash Translation Layer** ("**FTL**"). (*Id.*) The physical reading, writing and erasing of NAND flash cells is performed by circuity in a NAND flash memory module, shown in multiples in Sechen paragraphs 29 and 30 ("flash," "flash memory" or just "NAND") on a board, along with a block diagram identifying reading, writing and erasing functions (*id.* ¶ 29). A POSITA would understand this to be a **NAND flash memory or module**. (*Id*. ¶ 30.)

B. Characteristic NAND Flash Memory Management

Sechen paragraphs 34-38 review well-known processes to manage the peculiarities of NAND flash memory. These include **bad block management** and **garbage collection**. Typically, "hot" data, *i.e.*, frequently updated (such as system) data, and "cold" (not-"hot") data are clustered in respective physical NAND flash blocks in order to facilitate efficient (fewer erase-and-re-writes) collecting of "valid" (un-updated) stored data and freeing up the blocks they otherwise occupy. In addition to avoiding unnecessary burdening of the NAND system operation with

long "erases," this and other "**wear leveling**" processes even out write and erase stresses on NAND memory blocks – which have limited expected "endurance" ("lifetime") – aiming to avoid premature failure or retirement of enough blocks that would compromise system use. (Sechen ¶ 39 [**system lifetime**].)

C. Introduction of MLC and SLC Flash

Multi-level cell (MLC) NAND flash, developed in 1992-97, as described in Atwood (Ex. 1038), was the use of the NAND storage cell with four threshold voltage levels to represent two bits of information. (Sechen ¶¶ 40-41.) The first known use of the term "SLC" was in a 2001 article that called the established technology **"single-level program cell NAND flash (SLC)"** and compared it with then new MLC. (Ex. 1039, Sechen ¶ 42.) Although the terms have been shortened over the years to their acronyms, a POSITA would understand that "MLC" refers to the use of NAND storage cells, each storing two bits written in two passes – a "lower" (least significant bit, "LSB") logical page and then a distinct "upper" (most significant bit, "MSB") logical page. (*Id.* ¶ 43.) That is, MLC flash memory allows writing **two logical data pages into one physical page of memory**. As reviewed at Sechen paragraph 43 (particularly Sechen note 17), the programming of MLC memory is more complex and stressful than for SLC because of MLC's narrower window for distinct threshold charges and its multiple passes of writing; also, a word-wide set of latches (word buffer) is required for reading *each* level of logical

word programming. Thus, a POSITA would understand that different circuitry is required for **MLC NAND flash memory** than for **SLC NAND flash memory,** and that **SLC NAND flash memory is incapable of storing more than one logical data page in one physical page of memory**.

Two-state/level operation of MLC (later called **pseudo-SLC)** was disclosed as early as 2002. (Sechen ¶ 49; U.S. Patent No. 6,456,528 to Chen for "Selective Operation of a Multi-State Non-Volatile Memory System in a Binary Mode" ["Chen '528," Ex. 1040] [one-bit *mode* of writing only one logical data page into one MLC physical page of memory]). This mode of MLC use, applying the MLC memory circuitry and pulsed/phased writing, has inferior performance and endurance relative to "real" SLC. A SanDisk suite of patents (incorporating by reference among others, Chen '528) disclosed lower-density (pSLC) *mode* use of MLC modules through partitioning and configuration for multiple purposes, including faster processing for caching writes and error management (*e.g.*, U.S. Patent No. 8,634,240 to Gavens et al., published Apr. 28, 2011 ["Gavens," Ex. 1045]; U.S. Patent No. 8,806,301 to Yu et al., published Apr. 14, 2011 [Ex. 1046], explained at Sechen ¶¶ 53-56).

D. Known Hybrid Use of SLC and MLC NAND Flash Modules with Strategies of Preferential Mapping, Wear Leveling and Error Management

Because of the compression of its threshold levels and its multiple passes of writing, MLC NAND flash, less costly per bit of information stored, was well-

known to wear more quickly (have less "endurance") but perform more slowly than SLC NAND flash. (Sechen ¶¶ 40-42.) Thus, it was proposed as early as 2008 to use hybrid systems of **SLC modules** and **MLC modules** in which more **frequently written**/updated (to logical locations) **data**, such as system data, user directories, and other user files accessed more frequently) would be written **preferentially to** a more robust **SLC module** while less frequently updated data such as user data would be written to the MLC module. (Sechen ¶ 46; Nelson Duann, SLC & MLC Hybrid (Silicon Motion Flash Memory Summit, Santa Clara, CA, Aug. 12, 2008) ["Duann," Ex. 1034]; U.S. Patent No. 8,078,794 to Lee et al., published Sept. 4, 2008 ["Lee '794," Ex. 1041].)

This **SLC module preference** was applied in 2008-09 to a variety of **wear leveling** techniques to avoid premature failure of a system by some components reaching their endurance limits before others. (Sechen ¶¶ 47-48 and 58-60; Sutardja, U.S. Patent Appl. Pub. No. 2008/0140918, "Hybrid Non-Volatile Sold State Memory System" ["Sutardja," Ex. 1042]; Moshayedi, U.S. Patent Appl. No. 2009/00327591, "SLC-MLC Combination Flash Storage Device" ["Moshayedi," Ex. 1043].)

Additional tools or techniques were known in the art before the priority date to avoid premature failure by applying **error management** near the end of life of the system with **preferential writing to SLC** upon determining failure of some

"data integrity test" at an MLC location. (Sechen ¶¶ 53-56 and 58-61, e.g., Gavens (Ex. 1046.)

The availability to the POSITA of these **preferential writing to SLC** techniques to avoid premature failure make obvious the base claim 1 of the '298 patent. (Sechen ¶ 62, Sections X(A), XI(A) and XII(A).)

V. THE '298 PATENT: Specification

The core disclosed system of known MLC NAND flash and SLC NAND flash modules is described as such, multiple times in the specification. (*E.g.*, 4:51-56, 6:36-46; Sechen ¶ 69.) Fig. 4 shows a physical memory device module including two banks each including 8+ **MLC flash memory modules** and 2+ **SLC NAND flash memory modules**). A controller "coupled" to the MLC and SLC modules performs L2P/FTL functions essential for NAND flash operation and also performs common wear leveling. (3:65-4:4; Sechen ¶¶ 70 and 72.) Only two "embodiments" are features that are optional for hybrid NAND flash memory systems.

The first, "data integrity test" feature is shown as the NAND flash method steps of Figs. 3a and 3b, largely, but not perfectly, explained in the text (5:55-6:23). Dr. Sechen identifies (¶ 71) ambiguities as to what is compared in the test as "stored data" and "retained data."

The second, feature is functionally explained in two sentences that describe what appears to be "wear leveling" by "**allocat[ing]** 'hot' blocks i.e., those blocks

that receive frequent writes, into the **SLC NAND flash memory module 28**, while allocating 'cold' blocks; i.e., those blocks that only receive infrequent writes, into the MLC NAND flash memory module 26." (6:24-35 [emphasis added].) Dr. Sechen notes (¶ 73) various ambiguities, including its unclear reference to logical or physical blocks.

VI. The Challenged Claims of the '298 Patent (37 C.F.R. § 42.104(b)(1):

Claim 1 (7:6-8:9). In this "system for storing data" (7:8) the first two elements of physical MLC and SLC non-volatile memory ("NVM") modules (7:9-10 and 7:11-12) are "coupled" to a third element of a controller (7:13-15) defined by four functions that reflect with widely varying faithfulness DETAILED DESCRIPTION "embodiments." Controller function a) (7:16-25) corresponds to the FTL/L2P functions admitted in the BACKGROUND (2:30-3:13) and reviewed at Sechen paragraphs 70 and 72. Controller function b) (7:26-33) is recited as a severe abbreviation or abstraction of the DETAILED DESCRIPTION of the Figs. 3a and 3b data integrity test (5:55-6:23) reviewed at Sechen paragraph 71. Controller functions c) and d) are recited in a rearrangement of terms and concepts of the two-sentence DETAILED DESCRIPTION of apparent wear leveling (6:24-35) reviewed at Sechen paragraph 73. The Examiner found controller function d), "allocation" of "those blocks that receive the most frequent writes" by "transfer" to the SLC module, to be the sole point of invention. (Sechen ¶¶ 76-81, *citing* June 17,

2014, Office Action at 4 and 7-8 [Ex. 1006 at 42 and 45-46] [limitations through controller function b) anticipated by Gorobets '179, controller function c) obvious in view of Segal]and Sept. 2, 2014, Notice of Allowability at 2 [Ex. 1006 at 80].)

Dependent Claims: Claims 2 and 3 limit the system of claim 1 to a minimum quanta of addresses to one block and one page respectively, without specifying logical or physical. Claims 4-10 limit the system of claim 1 to MLC and SLC NVM modules of NAND flash and alternative technologies. Claim 11 limits the system of claim 1 to transfer of content on a periodic basis.

In view of the literal claims and disclosure of the '298 patent, the accompanying prior art references and supporting declaration of Dr. Sachen (Ex. 1002), Petitioner respectfully requests cancellation of the Challenged Claims as summarized (pursuant to 37 C.F.R. § 42.104(b)(1), -(2), -(4) and –(5)) in the following table.

Grounds	**Exhibits**
Ground 1: Claims 1-11 are unpatentable under U.S.C. § 103 as obvious over Gavens and incorporated disclosures in view of a POSITA's knowledge of NAND flash technology.	1001, 1002, 1003, 1039, 1040, 1041, 1042, 1043, 1045, 1049, 1050, 1051, 1053, 1054, 1055

Ground 2: Claims 1-11 are unpatentable under U.S.C. § 103 as obvious over Moshayedi in view of a POSITA's knowledge of NAND flash technology.	1001, 1002, 1042, 1043, 1045, 1053, 1054, 1055
Ground 3: Claims 1-11 are unpatentable under U.S.C. § 103 as obvious over Sutardja in view of a POSITA's knowledge of NAND flash technology.	1001, 1002, 1042, 1043, 1045, 1053, 1054, 1055

VII. Person of Ordinary Skill in the Art

A person of ordinary skill in the art at the time of the claimed inventions would have a bachelor's degree in computer engineering, electrical engineering, computer science, or a closely related field, along with at least two years of experience in the design, development, implementation, or management of memory devices and systems. A person with an advanced degree in a relevant field, such as computer or electrical engineering, would require less experience in the development and use of memory devices and systems. As is common, one could obtain equivalent knowledge and perspective from other life experiences as well. (Sechen ¶¶ 60-61.)

VIII. 37 C.F.R. § 42.104(b)(3): CLAIM CONSTRUCTION

Under 37 C.F.R. § 42.100(b), claims in an *inter partes* review proceeding are construed using the same claim construction standard that is used to construe claims in a civil action. This includes "construing the claim in accordance with the ordinary and customary meaning of such claim as understood by one of ordinary skill in the art and the prosecution history pertaining to the patent," *id*. The standard for claim construction is set forth in *Phillips v. AWH Corp.,* 415 F.3d 1303 (Fed. Cir. 2005).

The terms "MLC non-volatile memory modules" and "SLC non-volatile memory module" (7:9 and 7:11) have distinct, established, customary meanings for a POSITA (Section IV *supra* and Sechen Section V) where "memory" and "module" also have distinct meanings:

- "MLC memory modules" means "modules comprising MLC non-volatile memory" *where*

  "**MLC non-volatile memory**" means "**non-volatile memory cells arranged with circuitry capable of storing multiple logical pages in a single physical page of cells**" and

- "SLC memory module" means "module comprising SLC non-volatile memory" *where*

  "**SLC non-volatile memory**" means "**non-volatile memory cells arranged with circuitry incapable of storing multiple logical pages in a single physical page of cells**"

Dr. Sechen explained (¶ 90) why a POSITA would only understand "SLC" and "MLC" to refer to distinct NAND flash memory circuity. Most tellingly, despite his knowledge of use of MLC in 2-state *mode*, corroborated by his Information Disclosure Statements submitted with the original application,[6] Dr. Rao chose to describe and claim his invention as a system with distinct, known MLC and SLC memory modules, "adapted" to two operational features – but not 2-state *mode* operation of only one MLC memory. In addition to explicit references to distinct modules (e.g., 4:51-56, 6:36-46, Fig. 4), Dr. Rao incorporated Rao '916 (Ex. 1004), his proposal to operate a known SLC memory module (Figs. 6a and 8 [100K p-e cycle endurance]) in a purportedly new way.

"**Blocks**": Petitioner maintains its position in the court litigation that the term is indefinite as to whether it refers to logical or physical blocks (Sechen ¶¶ 73 and 90) but will show here that the recited limitations are met under either construction.

---

[6] (Ex. 1006 at 26-36, notably Taehee Cho et al., *A Dual-Mode NAND Flash Memory: 1-Gb Multilevel and High-Performance 512-Mb Single Level Modes*, 36 IEEE J. Solid State Circuits 1700, 1700 (Nov. 2001) [Ex. 1068] and Ken Takeuchi, *A Multipage Cell Architecture for High-Speed Programming Multilevel NAND Flash Memories*, 33 IEEE J. Solid-State Circuits 1228, 1230 Figs. 5-7 (Aug. 1998) [Ex. 1069])

"**Controller**": Petitioner maintains its position in the court litigation that the term is indefinite for, among other things, failing to provide necessary structure or algorithms, but will show that the recited functions are met by the art.

"Data Integrity Test": Petitioner maintains its position in the court litigation that this term (across multiple patents) means "a test that compares stored data to retained data." As noted by Dr. Sechen (¶ 87), the claim recitation of that test is defective. To the extent understandable, the art discloses the claimed use of the test.

* * * * * * *

Grounds for Review Under 37 C.F.R. § 42.104(b)(4)-(5)):

IX. GROUND 1: The Challenged Claims Are Obvious Over Gavens, Including Incorporated References, in View of Knowledge of the POSITA

The Challenged Claims are obvious under *KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398 (2007) because there is no gap between the hybrid SLC-MLC NAND flash system "adapted" to perform ordinary, necessary L2P/FTL functions and optional functions of directing data to the SLC module upon failure of a data integrity test or for wear leveling. Starting with a Gavens NAND flash system with multiple alternative error management operations including directing data to SLC, a POSITA would be motivated to use that system consistently with known wear leveling operations together meeting the Challenged Claim limitations directed to avoiding premature system failure.

A. Claim 1 Is Obvious Over Gavens, Including Incorporated References, in View of Knowledge of the POSITA.

(Sechen ¶ 96): Gavens (Ex. 1045), among the SanDisk portfolio and incorporating by reference others of that portfolio,[7] is a published reference that discloses a complete multi-modal NAND flash system with the ordinary L2P/FTL function of its controller, a variety of post-write read error detection and management schemes, and the capability of and compatibility for employing other known tools (disclosed in incorporated documents and exemplified in [Sechen ¶¶ 58-61] used to avoid premature NAND system failure.

---

[7] *See, e.g.*, U.S. Patent No. 6,456,528 to Chen ("Chen '528"), Ex. 1040, incorporated by reference at Gavens 16:29-40; U.S. Patent Appl. Pub. No. 2011/0153912 ("Gorobets '912"), Ex. 1049, incorporated by reference at Gavens 20:53-59; U.S. Patent No. 5,930,167 to Lee et al. ("Lee '167"), Ex. 1050, incorporated by reference at Gavens 17:57-67; U.S. Patent Appl. Pub. No. 2010/0172180, to Paley et al. ("Paley"), Ex. 1051, incorporated by reference at Gavens 8:41-42. Hereinafter, it should be understood that reference to a specific portion of any of the above references is understood to be disclosed by Gavens due to incorporation by reference, as if the incorporated document was explicitly contained in Gavens.

(Sechen ¶ 97): Since the 2001 (*e.g.*, Cho, Ex. 1039) emergence of MLC and recharacterization of previously established NAND flash as SLC, hybrid MLC memory modules and SLC memory modules have availed of the endurance advantages of SLC. (*E.g.*, Lee '794 [Ex. 1041], Sutardja [Ex. 1042], and Moshayedi [Ex. 1043]). Many implementations involve partitioning and configuration of some MLC blocks for one-level (two states) or pseudo-SLC (pSLC) operation. As reviewed at [Sechen ¶ 49] and in Gavens itself (*e.g.*, 12:6-10) the 2-state mode of operation enjoys superior endurance compared to MLC in normal mode – though not as superior as "real" SLC memory which uses different circuitry. It would be obvious to a POSITA to use a "real" SLC module for the purposes contemplated in Gavens. Gavens and its incorporated references, in view of the knowledge of the POSITA, disclose the limitations recited in claim 1:

1. **[1a] at least one MLC non-volatile memory module comprising a plurality of individually erasable blocks;**

(Sechen ¶ 101): Gavens' memory chip 100 "includes a memory array 200 of memory cells with each cell capable of being configured as a multi-level cell ("MLC") (8:19-21). The memory array may be organized in erasable blocks erased together in a "flash" operation (Fig. 6, 11:4-20). Thus, Gavens discloses in memory chip 100 (Fig. 1 (100), 8:13-21) discloses claim element **[1a]**.

2. [1b] at least one SLC non-volatile memory module comprising a plurality of individually erasable blocks;

(Sechen ¶ 102): Gavens discloses, as a base system for multiple operational embodiments, "[a] first portion has each memory cell storing one bit of data and [a] second portion has each memory cell storing more than one bit of data" (4:32-34, 12:6-8 [memory operating in 1-bit mode "binary" or "SLC" memory]). It is explained that the first portion, being "configured as lower density storage[,] . . . operates with a wider margin of error than that of the second portion" (18:16-20, also 12:6-10, 21:6-9). As explained in [Sechen ¶ 96], it would be obvious to a POSITA to substitute a "real" SLC flash memory module for the Gavens first portion, with a similar, if not wider margin of error than binary mode operation of MLC (*see* [Sechen ¶ 49]) and achieve similar, if not better results with the same NAND flash controller Fig. 1(102) with great expectation of success. Gavens (4:32-34, 12:6-8), in view of the knowledge of the POSITA's knowledge of NAND flash, discloses the obvious use in the claimed system of claim element **[1b]**.

3. [1c] a controller coupled to the at least one MLC non-volatile memory module and the at least one SLC non-volatile memory module wherein the controller is adapted to;

(Sechen ¶ 103): Gavens discloses controller 102 which "cooperates with the memory chip and controls and manages high level operations" (Fig. 1(102), 8:32-33). It is noted that "one or more memory chip[s] 100 [are] managed by a controller

102" (8:17-18), multiple (NAND flash) chips shown in Fig. 1 by ellipsis with bi-directional coupling to the controller. It would be obvious to have one or more of those NAND flash chips be "native" SLC flash chips (modules), with the controller directing host commands and data to the SLC and MLC NAND chips with their particular circuitry "to perform memory operations on the memory array 200" (8:25-29.) Gavens (8:17-18), in view of the knowledge of the POSITA's knowledge of NAND flash, discloses the controller coupled to an MLC module and obviously substituted SLC module, meeting claim element **[1c]**.

4. **[1d] a) maintain an address map of at least one of the MLC and SLC non-volatile memory modules, the address map comprising a list of logical address ranges accessible by a computer system, the list of logical address ranges having a minimum quanta of addresses, wherein each entry in the list of logical address ranges maps to a similar range of physical addresses within either the at least one SLC non-volatile memory module or within the at least one MLC non-volatile memory module;**

(Sechen ¶¶ 104 and 105): Gavens discloses this controller function a) as the ordinary NAND controller function: "A memory block management system implemented in the controller stages the sectors and maps and stores them to the physical structure of the memory array" (8:36-39.) Gavens discloses updating of "a directory in a block management system embedded in the firmware of the controller (see Fig. 1)" to direct accesses to one or the other portions of the memory (17:17-19, 18:47-49, 18:64-66). Ordinary NAND flash organization and mapping operation

(*see* [Sechen ¶¶ 27-29, 33, and 34]) is explained, with "the page being a minimum unit of programming and reading" (11:9-29). [*Also* Paley (Ex. 2051) ¶¶ [0020] ("prevalent LBA interface") and [0155].)

(Sechen ¶ 106): Gavens (8:36-39. 11:9-29, 17:17-19, 18:47-49, 18:64-66) and its incorporated Paley reference (¶¶ [0020], [0155]) disclose the ordinary NAND flash controller function a) of claim element **[1d]**.

5. [1e] b) determine if a range of addresses listed by an entry and mapped to a similar range of physical addresses within the at least one MLC non-volatile memory module, fails a data integrity test, and, in the event of such a failure, the controller remaps the entry to the next available equivalent range of physical addresses within the at least one SLC non-volatile memory module;

(Sechen ¶ 107): As explained in [Sechen ¶ 88], the literal reading of controller function b), determining whether a range of (logical) addresses listed by an entry in an L2P map "fails a data integrity test," is absurd. There are multiple types of data integrity tests, including by ECC that compares the original data and written data only indirectly and with "failure" not typically remapping at all. [Sechen ¶¶ 51 and 52.] Gavens discloses multiple operations for error management, most relying on ECC, and deployment near the end-of-life of a memory system. [Sechen ¶¶ 53-55.] However, two embodiments describe the controller function b) comparing read data to original data and remapping to SLC if there is no match.

(Sechen ¶ 108): In one variant of the Gavens post-write-read error management embodiment of Figs. 14 and 15 (explaining text at 16:13-17:42), user data is written to the higher density (*e.g.*, MLC) portion as a "first copy." (16:41-45.) Then, in a "post write read," there may be a check for errors "by comparison with the **original copy** which may be cached" (16:46-49 [emphasis added]). If the number of error bits is less than a predetermined amount, the first copy is deemed valid and subsequent reads of the data page will be from that copy with any errors corrected by ECC. (16:50-56.) If the number is greater, then a copy is written to the lower density, lower error rate (more robust, *e.g.*, "SLC") portion; this is performed in one embodiment by copying from a cached copy ("of the original data") or in another embodiment from application of ECC to the first copy. (17:9-14.) In incorporated Paley (Ex. 2051), it was noted that "[u]sing RAM in a write cache operating with flash memory has been disclosed" (¶ [0027]); the '298 specification shows in Fig. 1 DRAM 20 which is where data to be written in the data integrity check apparently is first held (5:57-60). Thus, this variant discloses that written (stored) data that fails a data integrity test against retained (cached) original data to be written causes remapping of the original data to the SLC portion.

(Sechen ¶ 109): In another group of Gavens alternative post-write-read error management embodiments of Figs. 16 and 17 (explaining text at 17:43-19:24), the more robust portion is further divided into a first section for caching incoming data

(simultaneously with a "first copy" written into the higher density portion) and a second section for storing rewrites from the higher density portion. (18:21-29.) In a variant of that embodiment, the first copy (in higher density) is compared with the cached copy in the first section of the more robust portion in the on-chip data latches (18:30-44), and if the number of error bits is greater than a threshold, then a copy is rewritten from the cached copy to the rewrite section of the more robust portion. (18:53-66.).

(Sechen ¶ 110): It is understood that these actions are directed by the controller. The use of a separate SLC module instead of a binary-mode portion of one MLC module does not significantly increase latency, since any "copying" involves decoding, temporarily storing in a buffer, and coding (see notes 11 and 45 above) and comparison may be more efficiently done off the memory chip.

(Sechen ¶ 111): Thus, Gavens, in both the variants recited (16:13-17:42, 17:43-19:24) discloses the data integrity test controller function b) of element **[1e]**.

6. [1f] c) determine which of the blocks of the plurality of the blocks in the MLC and SLC non-volatile memory modules are accessed most frequently by maintaining a count of the number of times each one of the blocks is accessed; and

(Sechen ¶ 112): This recited controller function c) is not tied to controller function d) explicitly, temporally or even by use of the same terms. [Sechen ¶ 88]. All that is required is that the controller is "adapted" to "determine" which of the

blocks in the NAND flash system are "accessed most frequently by maintaining a count of the number of times each one of the blocks is accessed." This limitation is met on its face by Gavens's "[t]racking the age of each block by maintaining a hot count that records the number of erase/program cycling each block has undergone" (Fig. 19 (720), 20:18-21). Although (and because) the controller uses that information to determine according to a threshold count whether a particular post-write-read scheme should be implemented for the block (Fig. 19 (730), 20:22-27), the controller clearly is adapted to maintain the counts and determine from those counts the blocks accessed most frequently.

(Sechen ¶¶ 113 and 114): Incorporated by reference Lee '167 (Ex. 1050) (2:54-3:7) recites "[t]he usual desire to evenly wear the memory" that is met by its wear leveling approach that maintains "separate counts of the number of times that each sector has been erased and programmed." Incorporated by reference Gorobets '912 (Ex. 1049) ¶ [0126] also discloses wear leveling erase/rewrite counts among all blocks. (*Also* Chen '528 (Ex. 1040) 11:7-10; Paley (Ex. 1051) ¶ [0542].)

(Sechen ¶ 115): Gavens and incorporated references disclose the determining number of accesses controller function c) of claim element **[1f]**.

7. [1g] d) allocate those blocks that receive the most frequent writes by transferring the respective contents of those blocks to the at least one SLC non-volatile memory module.

(Sechen ¶ 116): The recited controller function d) is not tied to controller function c) explicitly, temporally or even by use of the same terms. [Sechen ¶ 88]. For at least this reason, it is unclear whether "those blocks" are logical or physical. Both are constructions are considered here.

(Sechen ¶ 117): The Gavens embodiments of post-write-read mapping to 2-state memory that meet controller function b) (element [1e], [Sechen ¶¶ 108 and 109]), as deployed after the physical block reaches a threshold erase/program count (Fig. 19 (730, 740), 20:21-29) literally meet limitation [1g], which has no temporal or completeness qualification: when those blocks exceeding the threshold (receiving the most frequent writes) are written with data that fails a data integrity test, they are rewritten (transferred) to SLC.

(Sechen ¶ 118): Also, various tools are disclosed by Gavens and its incorporated references disclose transfer of data to avoid premature failure of the system because of excessive wear on a component, tracked at a physical block level. For example, [Sechen ¶ 113] recited such disclosures in the incorporated by reference Lee '167 (Ex. 1050) (2:54-3:7) and Gorobets '912 (Ex. 1049) (¶ [0126]). Also, as admitted prior art: "[W]ear leveling algorithms within the FLASH devices . . . to attempt to ensure that hot blocks; i.e., those that are frequently written, are not

rendered unusable much faster than other blocks." (Provisional (Ex. 1003) ¶ [007], Sechen ¶ 120.)

(Sechen ¶ 121): As to logical blocks, Sutardja (Ex. 1042) discloses a controller directing a base algorithm similar to that of Lee '794 (Ex. 1041), mapping logical addresses with low write frequency (as reported by the host) to the lower endurance memory and logical write addresses with high write frequency to the higher endurance memory. (Fig. 7A (506) & (508), Sechen ¶¶ 47 and 58-60.)

(Sechen ¶ 122): Similarly, Moshayedi (Ex. 1043)

> keeps track of the number of times that data for each logical block address (LBA) has been written to the flash memory, and determines whether to store newly received data associated with a particular LBA in SLC flash or in MLC flash, depending on the number of writes that have occurred for that particular LBA

(¶ [0024].)

(Sechen ¶ 123): Thus, Gavens and incorporated references (with the POSITA's knowledge of obvious substitution of a "pure" SLC module for its partitioned and 2-state (pseudo-SLC) configured MLC module), in view of background knowledge of wear leveling exemplified by Sutardja, Moshayedi, and Lee '794 (with "pure SLC" modules, and as background processes compatible with Gavens' error management techniques), disclose the allocation by transfer controller function d) of claim element **[1g]**.

* * * * * * * * * * *

(Sechen ¶ 124): Gavens, with its incorporated references, meets each limitation of '298 claim 1 with known hybrid SLC-MLC NAND flash memory module systems and mixed SLC-mode and MLC-mode operation of MLC memory meeting elements [1a]-[1d], with the POSITA's understanding that Gavens's extensive technology-justified teachings include or render obvious, with the knowledge of the POSITA of many NAND flash operations to avoid premature failure, the features of elements [1e]-[1g].

B. Claim 2 Is Obvious Over Gavens, Including Incorporated References, in View of Knowledge of the POSITA.

1. [2Pre] The system of claim 1,

See Section VII(A).

2. [2] wherein the minimum quanta of address is equal to one block.

(Sechen ¶ 125): Gavens discloses or renders obvious limitation [2]. (Fig. 20C, 21:22 ["virtual block"].)

(Sechen ¶ 126): Incorporated Paley (Ex. 1051) explains block-level addressing (¶ [0016] ("logical blocks"), ¶ [0163] ("logical address LBA").

(Sechen ¶ 127): Thus, Gavens and incorporated references, in view of knowledge of the POSITA, discloses claim 2.

C. Claim 3 Is Obvious Over Gavens, Including Incorporated References, in View of Knowledge of the POSITA.

1. [3Pre] The system of claim 1,

See Section VII(A).

2. [3] wherein the minimum quanta of address is equal to one page.

(Sechen ¶ 130): Gavens discloses a minimum quanta of addresses equal to one page. (Fig. 4 (70), 10:19-20, 41-47, 20:2-4 [a page 70 is a group of memory cells enabled to be sensed or programmed in parallel], 3:21-22[a page of memory elements are read or programmed together],11:20-29 [A page is a minimum unit of programming and reading], 9:36-40 [the memory array 200 is arranged in rows and columns of memory cells, addressable by word lines and bit lines].)

(Sechen ¶ 131): Gavens and incorporated references, in view of the knowledge of the POSITA, discloses claim 3.

D. Claim 4 Is Obvious Over Gavens, Including Incorporated References, in View of Knowledge of the POSITA.

1. [4Pre] The system of claim 1,

See Section VII(A).

2. [4] wherein the MLC non-volatile memory module is NAND flash memory.

(Sechen ¶ 133): Gavens would be understood by a POSITA to describe a NAND flash system with multiple MLC NAND flash modules. (Fig. 4, 9:49-50 [page of memory cells in the memory array 200 is organized in the NAND configuration],1:18-19, 4:14-18 [flash memory having a first portion and a second

portion storing higher density compared to the first portion], 2:65-66, 8:65-67 [flash memory devices with NAND string structures; Fig. 1 (90) [flash memory device].)

(Sechen ¶ 134): Gavens and incorporated references, in view of the knowledge of the POSITA, discloses claim 4.

E. Claim 5 Is Obvious Over Gavens, Including Incorporated References, in View of Knowledge of the POSITA.

1. [5Pre] The system of claim 1,

See Section VII(A).

2. [5] wherein the SLC non-volatile memory module is NAND flash memory.

(Sechen ¶ 136): See [Sechen ¶ 133] (Gavens's memory is NAND flash).

(Sechen ¶ 137): A POSITA would readily and successfully substitute for Gavens' partitions of MLC memory elements configured for 2-state operation an SLC non-volatile memory module which is understood to include SLC NAND flash memory elements. [Sechen ¶¶ 97 and 102.]

(Sechen ¶ 138): Gavens and incorporated references, in view of the knowledge of the POSITA, discloses claim 5.

F. Claim 6 Is Obvious Over Gavens, Including Incorporated references, in View of Knowledge of the POSITA.

1. [6Pre] The system of claim 1,

See Section VII(A).

2. [6] wherein the MLC non-volatile memory module is resistive random-access memory (RRAM).

(Sechen ¶ 140): To the extent that this claim is enabled by the disclosures in the specification, it is obvious in view of the assumed background knowledge.

(Sechen ¶ 141): This background knowledge is corroborated by the Examiner's rejection ([Sechen ¶ 78], citing Gorobets '179 (Ex. 1044) in view of De Ambroggi (Ex. 1053) (RRAM, PCM) and Kund (Ex. 1055) (alternative technologies having advantages over flash memory).

(Sechen ¶ 142): Gavens and incorporated references, in view of the knowledge of the POSITA, discloses claim 6.

G. Claim 7 Is Obvious Over Gavens, Including Incorporated references, in View of Knowledge of the POSITA.

1. [7Pre] The system of claim 1,

See Section VII(A).

2. [7] wherein the SLC non-volatile memory module is resistive random-access memory (RRAM).

(Sechen ¶ 143): *See* [Sechen ¶¶ 140 and 141].

(Sechen ¶ 144): Gavens and incorporated references, in view of the knowledge of the POSITA, discloses claim 7.

H. Claim 8 Is Obvious Over Gavens, Including Incorporated references, in View of Knowledge of the POSITA.

1. [8Pre] The system of claim 1,

See Section VII(A).

2. [8] wherein the MLC non-volatile memory module is phase change memory (PCM).

(Sechen ¶ 147): To the extent that this claim is enabled by the disclosures in the specification, it is obvious in view of the assumed background knowledge.

(Sechen ¶ 148): This background knowledge is corroborated by the Examiner's rejection ([Sechen ¶ 78], citing Gorobets '179 (Ex. 1044) in view of De Ambroggi (Ex. 1053) (RRAM, PCM) and Kund (Ex. 1055) (alternative technologies having advantages over flash memory).

(Sechen ¶ 149): Gavens and incorporated references, in view of the knowledge of the POSITA, discloses claim 8.

I. Claim 9 Is Obvious Over Gavens, Including Incorporated references, in View of Knowledge of the POSITA.

1. [9Pre] The system of claim 1,

See Section VII(A).

2. [9] wherein the SLC non-volatile memory module is phase change memory (PCM).

(Sechen ¶ 151): *See* [Sechen ¶¶ 147 and 148].

(Sechen ¶ 152): Gavens and incorporated references, in view of the knowledge of the POSITA, discloses claim 9.

J. Claim 10 Is Obvious Over Gavens, Including Incorporated references, in View of Knowledge of the POSITA.

1. [10Pre] The system of claim 1,

See Section VII(A).

2. [10] wherein the SLC non-volatile memory module is magnetic random-access memory (MAGRAM).

(Sechen ¶ 154): To the extent that this claim is enabled by the disclosures in the specification, it is obvious in view of the assumed background knowledge.

(Sechen ¶ 155): This background knowledge is corroborated by the Examiner's rejection ([Sechen ¶ 78], citing Gorobets '179 (Ex. 1044) in view of Chen '418 (Ex. 1054) (MRAM), and Kund (Ex. 1055) (alternative technologies having advantages over flash memory).

(Sechen ¶ 156): Gavens and incorporated references, in view of the knowledge of the POSITA, discloses claim 10.

K. Claim 11 Is Obvious Over Gavens, Including Incorporated references, in View of Knowledge of the POSITA.

1. [11Pre] The system of claim 1,

See Section VII(A).

2. [11] wherein the controller causes the transfer of content on a periodic basis.

(Sechen ¶ 158): There is ambiguity in the term "periodic." [Sechen ¶ 73.] To the extent that "periodic" means "occasional," the actions disclosed in Gavens and incorporated references and background as reviewed at [Sechen ¶¶ 117-121] are "occasional" or responsive to conditions such as reaching a threshold.

(Sechen ¶ 159): Also, Gavens discloses that the controller performs operations in the background on a periodic basis. (15:39-41 [data retention errors can be

alleviated by periodically refreshing the threshold levels of the cells in a "read scrub" operation. (Also, Lee '167 (Ex. 1050) (2:7-10, 30-37 [preferably in the background without slowing down other operations, the data is read out of SLC and reprogrammed into MLC]; Sutardja [¶ [0148] [time to perform shift analysis].).

(Sechen ¶ 161): Gavens and incorporated references, in view of the knowledge of the POSITA, discloses claim 11.

X. Ground 2: The Challenged Claims Are Obvious Over Moshayedi in View of Knowledge of the POSITA

The Challenged Claims are obvious under *KSR Int'l Co. v. Teleflex Inc*., 550 U.S. 398 (2007) because starting with a Moshayedi (Ex. 1043) NAND flash system with multiple alternative wear leveling operations including directing data to SLC, a POSITA would be motivated to use that system consistently with error management operations together meeting the Challenged Claim limitations directed to avoiding premature system failure.

A. Claim 1 Is Obvious Over Moshayedi in View of Knowledge of the POSITA.

(Sechen ¶ 162): Moshayedi (Ex. 1043) is a published reference that discloses a complete hybrid NAND flash system both MLC and SLC modules, with the ordinary L2P/FTL function of its controller, a variety of wear leveling schemes, and the capability of and compatibility for employing other known tools (exemplified in

[Sechen 58-61]) used to avoid premature NAND system failure. Moshayedi, in view of the knowledge of the POSITA, discloses the limitations recited in claim:

1. **[1a] at least one MLC non-volatile memory module comprising a plurality of individually erasable blocks;**

(Sechen ¶ 166): Fig. 1 (114) shows four channels of "chips of flash (e.g., the first chip can be a K9K8G08 SLC flash, and the rest can be K9G8G08 MLC flash)" (¶ [0038]).

(Sechen ¶ 167): Paragraph [0003] explains that flash memory is non-volatile; paragraph [0005] explains that flash memory is organized into blocks and that it writes individual segments and pages but can only erase entire blocks.

(Sechen ¶ 168): Moshayedi discloses the first physical component of the claimed system, MLC module claim element **[1a]**.

2. **[1b] at least one SLC non-volatile memory module comprising a plurality of individually erasable blocks;**

(Sechen ¶ 169): See [Sechen ¶¶ 166 and 167], which recite the use of SLC flash chips shown as Fig. 1(112), thus disclosing the second physical component of the claimed system, SLC module claim element **[1b]**.

3. **[1c] a controller coupled to the at least one MLC non-volatile memory module and the at least one SLC non-volatile memory module wherein the controller is adapted to;**

(Sechen ¶ 170): Fig. 1 (102) shows a controller 102 coupled to the four channels of SLC and MLC chips.

(Sechen ¶ 171): Paragraph [0025] explains that controller 102 "controls operations of the flash storage device 10." Paragraph [0026] explains how the controller 102 is coupled to the SLC and MLC flash memory chips by an I/O bus. Paragraph [0036] explains that the controller controls the writing and erasing of data to the channels of flash memory in the memory array.

(Sechen ¶ 172): Moshayedi discloses the third physical component of the claimed system, controller claim element **[1c]**.

4. **[1d] a) maintain an address map of at least one of the MLC and SLC non-volatile memory modules, the address map comprising a list of logical address ranges accessible by a computer system, the list of logical address ranges having a minimum quanta of addresses, wherein each entry in the list of logical address ranges maps to a similar range of physical addresses within either the at least one SLC non-volatile memory module or within the at least one MLC non-volatile memory module;**

(Sechen ¶ 173): Paragraph [0006] explains that a logical block address (LBA) is mapped to a location within a physical block in the flash memory system. Paragraphs [0036] and [0038] explain "virtual-to-physical mapping (V2P) that is standard L2P mapping for NAND flash. Paragraph [0044] explains that the system firmware assigns LBAs to the physical address. Paragraph [0037] explains that V2P RAM 214 in controller 200 (architecture of controller 102) includes one or more tables that serve to record the physical block address of a specific virtual block. Fig. 3 and paragraph [0039] show V2P entries.

(Sechen ¶ 174): Moshayedi discloses the ordinary NAND flash controller function a) of claim element **[1d]**.

5. [1e] b) determine if a range of addresses listed by an entry and mapped to a similar range of physical addresses within the at least one MLC non-volatile memory module, fails a data integrity test, and, in the event of such a failure, the controller remaps the entry to the next available equivalent range of physical addresses within the at least one SLC non-volatile memory module;

(Sechen ¶ 175): Paragraph [0031] discloses that "data structures or linked lists may store information about the number of data errors that have occurred in read operations corresponding to each of the data blocks" and that the information "may allow controller to selecta data block from which to move dynamic data in favor of static data." Additional information about read error relative to their frequency or accumulation is explained. Paragraph [0033] explains that the information regarding the number data read errors associated with a given data block may be used to determine whether the data is dynamic or static, where dynamic data may be "relocated to data blocks with less wear." As SLC is indicated in paragraphs [0007] and [0022] to have 100 times the endurance of MLC, a POSITA would understand "blocks with less wear" to mean SLC. Moshayedi says as much at paragraph [0009] which explains that "static" and dynamic ("frequently written") data are to be kept in MLC and SLC, respectively. This discloses that failures of some data integrity test on a physical block (which is "read") resulting in "read errors" may be relocated

(moved by decoding and encoding as explained in [Sechen note 45] to a less worn block. It would be understood by a POSITA that an SLC block would have less wear per p-e cycle than an MLC block, or likely at the end-of-life of MLC blocks to have more life. This would align with the apparent rationale for the '298 patent's Figs. 3a and 3b direction to SLC of test-failed data stored in a location presumed to be worn, wear being the focus of the patent and of the Gavens late-in-life error management.

(Sechen ¶ 177): Alternatively, a POSITA seeking operations to help avoid premature failure of a NAND flash memory system would look not only to wear leveling but to known error management techniques such as those disclosed in Gavens. [Sechen ¶¶ 108 and 109.]

(Sechen ¶ 178): Moshayedi, by its read error management of paragraphs [0031] and [0033], discloses or through knowledge of the POSITA of Gavens, the data integrity test controller function b) of claim element **[1e]**.

6. [1f] c) determine which of the blocks of the plurality of the blocks in the MLC and SLC non-volatile memory modules are accessed most frequently by maintaining a count of the number of times each one of the blocks is accessed; and

(Sechen ¶ 179): Moshayedi discloses that the controller determines which physical blocks in the MLC flash and the SLC flash are accessed most frequently (¶ [0009] explains that write (access) frequency is reflected in erase frequency) by

maintaining an erase count of the number of times each physical block is erased. Additionally, Moshayedi discloses that the controller determines which logical blocks in the MLC flash and the SLC flash are written most frequently by maintaining a write count of the number of times each logical block is written.

(Sechen ¶ 180): Fig. 7A (710-712) and paragraph [0047] disclose checking the erase count of the block. If the erase count is over a specified number (e.g., 1000), then swap MLC flash data to SLC flash.

(Sechen ¶ 181): Fig. 8 (802) and paragraphs [0049] and [0050] disclose that if an MLC block has an erase count that reaches 500, then move the data in the MLC block to a free block of SLC.

(Sechen ¶ 182): Paragraphs [0030] and [0051] disclose tracking the number of times blocks have been erased. (*Also* ¶¶ [0051] and [0071] [keep the number of times that each logical block has been written to NAND] ¶ [0052] [write counts].)

(Sechen ¶ 183): The Abstract and paragraphs [0009] and [0024] state the general operation of the Moshayedi main embodiment: Keep track of the number of times that data for each logical block address (LBA) has been written to the flash memory. If the write count of an LBA is above the threshold, the logical block is written to SLC flash.

(Sechen ¶ 184): Thus, Moshayedi discloses the counting number of accesses controller function c) claim element **[1f]**.

7. [1g] d) allocate those blocks that receive the most frequent writes by transferring the respective contents of those blocks to the at least one SLC non-volatile memory module.

(Sechen ¶ 185): Moshayedi discloses the common transfer of the "hot data" of logical blocks that receive the most frequent writes to the SLC module (Abs., ¶¶ [0009] and [0024], claim 4, Sechen ¶¶ 122, 186, and 187) plus, in another embodiment, the physical blocks are counted:

> once a block in MLC flash reaches a threshold erase count (e.g., 500), the next write operation to that block triggers a swap where the data from the MLC flash block is written to a block in SLC flash. In this manner, data for an LBA that is frequently written and causes frequent erasures is moved to SLC flash which can perform more erase cycles than MLC flash.

(¶ [0032], *also* Figs. 7A (712) and 8 (802) and (812), ¶¶ [0030], [0047]-[0050], and [0060] Sechen ¶¶ 188-191.)

(Sechen ¶ 192): Thus, Moshayedi, under either logical or physical block construction, discloses the allocating by transfer controller function d) of claim element **[1g]**.

* * * * * * * * * * *

(Sechen ¶ 193): Moshayedi meets each limitation of '298 claim 1 with known hybrid SLC-MLC NAND flash memory module systems meeting elements [1a]-[1d], with its "read error" meeting the "data integrity test" claim element [1e] and

its wear leveling embodiments meeting, under either logical or physical block construction, controller functions c) and d) of claim elements [1f] and [1g].

B. Claim 2 Is Obvious Over Moshayedi in View of Knowledge of the POSITA.

1. [2Pre] The system of claim 1,

See Section VIII(A).

2. [2] wherein the minimum quanta of address is equal to one block.

(Sechen ¶ 195): Paragraph [0005] explains that flash memory is organized into blocks, which are each divided into pages. Flash memory writes individual segments. Paragraph [0006] explains the use of L2P addressing using blocks. The controller maps "physical block address [to] a specific virtual block." (¶ [0037].)

(Sechen ¶ 196): Although a physical flash block may not be a minimum quanta of address to be written in standard NAND flash operation, it is the minimum quanta to be erased (*see* [Sechen ¶ 28]). A POSITA would understand that physical and logical blocks may not correspond in size. (*See* [Sechen note 9].)

(Sechen ¶ 197): Moshayedi, in view of knowledge of the POSITA, discloses claim 2.

C. Claim 3 Is Obvious Over Moshayedi in View of Knowledge of the POSITA.

1. [3Pre] The system of claim 1,

See Section VIII(A).

2. [3] wherein the minimum quanta of address is equal to one page.

(Sechen ¶ 199): Paragraph [0005] explains that flash memory is organized into blocks, which are each divided into pages. Flash memory writes individual segments.

(Sechen ¶ 200): It would be obvious to set the page as a minimum quanta of address to be written, as it is a design choice that has been adopted by the industry. (*See* [Sechen ¶ 27].)

(Sechen ¶ 201): Thus, Moshayedi, in view of the knowledge of the POSITA, discloses claim 3.

D. Claim 4 Is Obvious Over Moshayedi in View of Knowledge of the POSITA.

1. [4Pre] The system of claim 1,

See Section VIII(A).

2. [4] wherein the MLC non-volatile memory module is NAND flash memory.

(Sechen ¶ 203): *See* [Sechen ¶¶ 166 and 167], which recite the use of MLC flash chips shown as Fig. 1(114). The chips identified at ¶ [0038] are NAND flash.

(Sechen ¶ 204): Thus, Moshayedi, in view of the knowledge of the POSITA, discloses claim 4.

E. Claim 5 Is Obvious Over Moshayedi in View of Knowledge of the POSITA.

1. [5Pre] The system of claim 1,

See Section VIII(A).

2. [5] wherein the SLC non-volatile memory module is NAND flash memory.

(Sechen ¶ 206): See [Sechen ¶¶ 166, 167 and 169], which recite the use of SLC flash chips shown as Fig. 1(112). The chips identified at ¶ [0038] are NAND flash.

(Sechen ¶ 207): Moshayedi, in view of the knowledge of the POSITA, discloses claim 5.

F. Claim 6 Is Obvious Over Moshayedi in View of Knowledge of the POSITA.

1. [6Pre] The system of claim 1,

See Section VIII(A).

2. [6] wherein the MLC non-volatile memory module is resistive random-access memory (RRAM).

(Sechen ¶ 209): To the extent that this claim is enabled by the disclosures in the specification, it is obvious in view of the assumed background knowledge.

(Sechen ¶ 210): This background knowledge is corroborated by the Examiner's rejection ([Sechen ¶ 78], citing Gorobets '179 (Ex. 1044) in view of De Ambroggi (Ex. 1053) (RRAM, PCM) and Kund (Ex. 1055) (alternative technologies having advantages over flash memory).

(Sechen ¶ 211): Moshayedi, in view of the knowledge of the POSITA, discloses claim 6.

G. Claim 7 Is Obvious Over Moshayedi in View of Knowledge of the POSITA.

1. [7Pre] The system of claim 1,

See Section VIII(A).

2. [7] wherein the SLC non-volatile memory module is resistive random-access memory (RRAM).

(Sechen ¶ 213): *See* [Sechen ¶¶ 209 and 210]

(Sechen ¶ 214): Moshayedi, in view of the knowledge of the POSITA, discloses claim 7.

H. Claim 8 Is Obvious Over Moshayedi in View of Knowledge of the POSITA.

1. [8Pre] The system of claim 1,

See Section VIII(A).

2. [8] wherein the MLC non-volatile memory module is phase change memory (PCM).

(Sechen ¶ 216): To the extent that this claim is enabled by the disclosures in the specification, it is obvious in view of the assumed background knowledge.

(Sechen ¶ 217): This background knowledge is corroborated by the Examiner's rejection ([Sechen ¶ 78], citing Gorobets '179 (Ex. 1044) in view of De Ambroggi (Ex. 1053) (RRAM, PCM) and Kund (Ex. 1055) (alternative technologies having advantages over flash memory).

(Sechen ¶ 218): Moshayedi, in view of the knowledge of the POSITA, discloses claim 8.

I. Claim 9 Is Obvious Over Moshayedi in View of Knowledge of the POSITA.

1. [9Pre] The system of claim 1,

See Section VIII(A).

2. [9] wherein the SLC non-volatile memory module is phase change memory (PCM).

(Sechen ¶ 220): *See* [Sechen ¶¶ 216 and 217].

(Sechen ¶ 221): Moshayedi, in view of the knowledge of the POSITA, discloses claim 9.

J. Claim 10 Is Obvious Over Moshayedi in View of Knowledge of the POSITA.

1. [10Pre] The system of claim 1,

See Section VIII(A).

2. [10] wherein the SLC non-volatile memory module is magnetic random-access memory (MAGRAM).

(Sechen ¶ 223): To the extent that this claim is enabled by the disclosures in the specification, it is obvious in view of the assumed background knowledge.

(Sechen ¶ 224): This background knowledge is corroborated by the Examiner's rejection ([Sechen ¶ 78], citing Gorobets '179 (Ex. 1044) in view of Chen '418 (Ex. 1054) (MRAM), and Kund (Ex. 1055) (alternative technologies having advantages over flash memory).

(Sechen ¶ 225): Moshayedi, in view of the knowledge of the POSITA, discloses claim 10.

K. Claim 11 Is Obvious Over Moshayedi in View of Knowledge of the POSITA.

1. [11Pre] The system of claim 1,

See Section VIII(A).

2. [11] wherein the controller causes the transfer of content on a periodic basis.

(Sechen ¶ 227): There is ambiguity in the term "periodic." [Sechen ¶ 73.] To the extent that "periodic" means "occasional," the actions disclosed in Moshayedi as reviewed at paragraphs 186-191 above are "occasional" or responsive to conditions such as reaching a threshold. (*E.g.* ¶ [0032] [reaching threshold erase count triggers a swap].)

(Sechen ¶ 228): It was common knowledge of the POSITA that wear leveling functions may be performed at the time of write (dynamic) or in the background when the device is not performing host-initiated read and write operations. ([Sechen ¶ 36], admitted art at 3:6-13; *also*, Sutardja (Ex. 1042) ¶ [0148].)

(Sechen ¶ 229): Moshayedi, in view of the knowledge of the POSITA, discloses claim 11.

XI. Ground 3: The Challenged Claims Are Obvious Over Sutardja in View of Knowledge of the POSITA

The Challenged Claims are obvious under *KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398 (2007) because starting with a Sutardja (Ex. 1042) NAND flash system with multiple alternative wear leveling operations including directing data to SLC, a POSITA would be motivated to use that system consistently with error management operations together meeting the Challenged Claim limitations directed to avoiding premature system failure.

A. Claim 1 Is Obvious Over Sutardja

(Sechen ¶ 230): Sutardja (Ex. 1042) is a published reference that discloses a complete hybrid NAND flash system using MLC and SLC modules, with the ordinary L2P/FTL function of its controller, a variety of wear leveling schemes, and the capability of and compatibility for employing other known tools (exemplified in [Sechen ¶¶ 58-61]) used to avoid premature NAND system failure. Sutardja, in view of the knowledge of the POSITA, discloses the limitations recited in claim 1:

1. [1a] at least one MLC non-volatile memory module comprising a plurality of individually erasable blocks;

(Sechen ¶ 234) Fig. 2 shows a First Solid-State Nonvolatile Memory ("NVM") 204 and a Second Solid-State NVM 206. Paragraph [0108] explains the memories 204 and 206 each "may include single-level cell (SLC) flash memory or multi-level cell (MLC) flash memory". Claim 37 recites the system wherein the

second memory includes single-level cell (SLC) flash memory, and the first memory includes multi-level cell (MLC) flash memory. As well-known to the POSITA, a flash memory includes many blocks, and the individually erasable block is the defining portion of NAND "flash" (*see* [Sechen ¶ 28], Sutardja ¶ [0157] [block is a group of memory cells erased together].)

(Sechen ¶ 235): Paragraph [106] would be understood by a POSITA that the first NVM may be MLC flash memory modules.

(Sechen ¶ 236): Sutardja discloses the first physical component of the claimed system, MLC module claim element **[1a]**.

2. [1b] at least one SLC non-volatile memory module comprising a plurality of individually erasable blocks;

(Sechen ¶ 237): *See* [Sechen ¶ 234], which recites the use of SLC flash memory, which may be included in the second memory. Paragraph [106] would be understood by a POSITA that the second NVM may be an SLC module

(Sechen ¶ 238): Sutardja discloses the second physical component of the claimed system, SLC module claim element **[1b]**.

3. **[1c] a controller coupled to the at least one MLC non-volatile memory module and the at least one SLC non-volatile memory module wherein the controller is adapted to;**

(Sechen ¶ 239): Fig. 3 shows a controller 252 that is coupled to both the first NVM 204 and the second NVM 206. Paragraph [118] explains this coupling as communication.

(Sechen ¶ 240): Paragraph [0009] explains, in the generic (prior art) SSD system that controller 102 "reads or writes data to the flash memory 104."

(Sechen ¶ 241): Sutardja discloses the third component of the claimed system, controller claim element **[1c]**.

4. **[1d] a) maintain an address map of at least one of the MLC and SLC non-volatile memory modules, the address map comprising a list of logical address ranges accessible by a computer system, the list of logical address ranges having a minimum quanta of addresses, wherein each entry in the list of logical address ranges maps to a similar range of physical addresses within either the at least one SLC non-volatile memory module or within the at least one MLC non-volatile memory module;**

(Sechen ¶ 242): Sutardja uses NAND flash standard L2P mapping (see [Sechen ¶¶ 32-34]) as corroborated by Sutardja's specific use of that mapping in its wear leveling and other features. (*E.g.*, Abs., ¶ [0107] [map logical addresses corresponding to data to physical addresses]). A POSITA would understand that Sutardja's various modules, which collect and respond to different defined conditions, use the basic NAND flash system of L2P mapping.

(Sechen ¶ 243): Sutardja in view of the knowledge of the POSITA discloses the ordinary NAND flash controller function a) of claim element **[1d]**.

5. [1e] b) determine if a range of addresses listed by an entry and mapped to a similar range of physical addresses within the at least one MLC non-volatile memory module, fails a data integrity test, and, in the event of such a failure, the controller remaps the entry to the next available equivalent range of physical addresses within the at least one SLC non-volatile memory module;

(Sechen ¶ 244): Figs. 5 (406) and 7D, explained at paragraphs [0134]-[0139] and [0151], disclose a "degradation" test of a physical location (range of physical addresses). Fig. 7D and paragraph [0151] explain the test as writing data to that location, reading back the data, writing the data to that location a second time (after a predetermined time), reading back the data a second time, and then comparing the data read back the first time and the data read back the second time, resulting in a "degradation value" for that physical address. One use for values that show degradation (in a sense, failing that data integrity test), is for estimating the maximum life cycle for a block based on its degradation and for the wear leveling process to "normalize" wear. A POSITA understands that such normalization is the wear leveling to avoid failure of one portion of the memory ahead of others. Paragraph [0139] discloses that using the number of write cycles remaining (from the estimated maximum), the wear leveling process may assign all new writes to another one of the memories rather than the one which is approaching the end of its

useful life. Paragraph [0135] discloses periodically testing the first, lesser write cycle lifetime (MLC) NVM, so that paragraph [0139] would call for remapping to the second, greater write cycle lifetime (SLC) NVM if the degradation test on the MLC location has "failed" by degradation values that have increased to the level suggesting near end-of-life for the location. A POSITA would understand that this test could result in remapping to SLC through the Sutardja processes directed to avoiding failure of some physical memory locations ahead of others.

(Sechen ¶ 245): Starting from that objective of a Sutardja system, it would also be obvious to a POSITA to apply the knowledge of other tools to avoid premature failure, for example, the two Gavens post-write-read error management embodiments described at [Sechen ¶¶ 108 and 109]. The Gavens post-write-read direction to SLC can be fit after Sutardja Fig. 7A (510).

(Sechen ¶ 246): Sutardja's degradation test of paragraphs [0134]-[0139] and [0151], applying knowledge of the POSITA, discloses the data integrity test controller function b) of element **[1e]**.

6. [1f] c) determine which of the blocks of the plurality of the blocks in the MLC and SLC non-volatile memory modules are accessed most frequently by maintaining a count of the number of times each one of the blocks is accessed; and

(Sechen ¶ 247): Sutardja discloses that the controller includes various modules that maintain a count of the number of times each physical block is written to and/or erased.

(Sechen ¶¶ 248-254): For example, Fig. 3 (260) and paragraphs [0110]-[0111] and [0121] disclose that wear leveling module 260 tracks the number times the write and/or erase operations are performed on each physical block in the first and second solid-state nonvolatile memories 204 and 206. *Also* Figs. 4A (write monitoring module 306), 4B (write mapping module 356), 7C (520) (Fig. 7E (548), paragraphs [0129], [0113], [0149], [0153], [0157] and [0159].

(Sechen ¶ 255): Fig. 7A (504) and (512) and paragraphs [0146] and [0147], disclose in the base algorithm, receiving write frequencies for logical addresses from the host and measuring actual write frequencies at which data is in fact written to the logical addresses.

(Sechen ¶ 257): Sutardja discloses the counting number of accesses controller function c) of claim element **[1f]**.

7. [1g] d) allocate those blocks that receive the most frequent writes by transferring the respective contents of those blocks to the at least one SLC non-volatile memory module.

(Sechen ¶ 258): Sutardja discloses operation of a NAND flash system in a "flow" diagram (Figs. 7A-7E) that has several branches (Figs. 7B-7E) that are not intended to follow each data write as they are qualified by whether it is "[t]ime to perform" certain analyses (Fig. 7B [514]["data shift analysis"], [516]["degradation analysis"] and [518]["wear analysis"]). The base algorithm is shown in Fig. 7A: step 504 ("Receive Write Frequencies For Logical Addresses From Host"), step 506 ("Map Logical Addresses With Low Write Frequencies To First Memory"], step 508 ("Map Logical Addresses With High Write Frequencies To Second Memory"), and step 510 ("Write Data According To Mapping"). This is the ordinary mapping and transfer to SLC memory of the data of the most frequently written LBAs ("hot data" of logical blocks) reviewed at [Sechen ¶¶ 46-48, 58 and 60]. Paragraph [0146] discloses the controller mapping the logical addresses having write frequencies greater than a predetermined threshold to the second NVM (SLC).

(Sechen ¶¶ 259 and 263): In the "data shift" branch, Fig. 7C (522) and paragraph [0149] discloses that if the number of write operations to the first [physical] block of the first NVM (MLC) during a predetermined time is greater than the predetermined threshold, then map the logical addresses that correspond to the first block of the first NVM (MLC) to a second block of the second NVM (SLC).

Paragraph [0126] discloses that static data may be moved, as in ordinary static wear leveling. Paragraphs [0131] and [0132] disclose that the write mapping module 356 can remap and move data between physical NVMs.

(Sechen ¶¶ 260-262 and 265): In the "wear leveling" branch, Fig. 7E (550) and paragraph [0153] disclose that if the wear level of the first [physical] NVM (MLC) is greater than a predetermined threshold, then map all logical blocks to physical blocks of the second NVM (SLC). *Also*, paragraph [0128] (biasing mapping upon threshold), [0164] (ordinary wear leveling of writing to less worn available block), [0168] (data stored in blocks "surpassing" lifetime written to other blocks),

(Sechen ¶ 266): Sutardja under either logical or physical block construction, discloses the allocate by transfer controller function d) claim element **[1g]**.

* * * * * * * * * * *

(Sechen ¶ 267): In summary, Sutardja meets each limitation of '298 claim 1 with known hybrid SLC-MLC NAND flash memory module systems meeting elements [1a]-[1d], with its "degradation test" meeting the "data integrity test" claim element [1e] and its wear leveling embodiments meeting controller functions c) and d) claim elements [1f] and [1g].

B. Claim 2 Is Obvious Over Sutardja in View of Knowledge of the POSITA.

1. [2Pre] The system of claim 1,

See Section IX(A).

2. [2] wherein the minimum quanta of address is equal to one block.

(Sechen ¶ 269): As well-known to the POSITA, a flash memory includes many blocks, and the individually erasable block is the defining portion of NAND "flash" (*see* [Sechen ¶ 28], Sutardja ¶ [0157] [block is a group of memory cells erased together].)

(Sechen ¶ 270): Although a physical flash block may not be a minimum quanta of address to be written in standard NAND flash operation, it is the minimum quanta to be erased (*see* [Sechen ¶ 28]). A POSITA would understand that physical and logical blocks may not correspond in size. (*See* [Sechen note 9].)

(Sechen ¶ 271): Sutardja, in view of knowledge of the POSITA, discloses claim 2.

C. Claim 3 Is Obvious Over Sutardja in View of Knowledge of the POSITA.

1. [3Pre] The system of claim 1,

See Section IX(A).

2. [3] wherein the minimum quanta of address is equal to one page.

(Sechen ¶ 273): It would be obvious to set the page as a minimum quanta of address to be written, as it is a design choice that has been adopted by the industry. (*See* [Sechen ¶ 27].)

(Sechen ¶ 274): Sutardja, in view of the knowledge of the POSITA, discloses claim 3.

D. Claim 4 Is Obvious Over Sutardja in View of Knowledge of the POSITA.

1. [4Pre] The system of claim 1,

See Section IX(A).

2. [4] wherein the MLC non-volatile memory module is NAND flash memory.

(Sechen ¶ 276): *See* [Sechen ¶¶ 234 and 235], which recite the use of MLC flash memory (¶ [0108].) This would be understood to be NAND flash as at least disclosed at paragraphs [0103] and [0172].

(Sechen ¶ 277): Thus, Sutardja, in view of the knowledge of the POSITA, discloses claim 4.

E. Claim 5 Is Obvious Over Sutardja in View of Knowledge of the POSITA.

1. [5Pre] The system of claim 1,

See Section IX(A).

2. [5] wherein the SLC non-volatile memory module is NAND flash memory.

(Sechen ¶ 279): *See* [Sechen ¶¶ 234 and 235], which recite the use of SLC flash memory (¶ [0108].) This would be understood to be NAND flash as at least disclosed at paragraphs [0103] and [0172].

(Sechen ¶ 280): Sutardja, in view of the knowledge of the POSITA, discloses claim 5.

F. Claim 6 Is Obvious Over Sutardja in View of Knowledge of the POSITA.

1. [6Pre] The system of claim 1,

See Section IX(A).

2. [6] wherein the MLC non-volatile memory module is resistive random-access memory (RRAM).

(Sechen ¶ 282): To the extent that this claim is enabled by the disclosures in the specification, it is obvious in view of the assumed background knowledge.

(Sechen ¶ 283): This background knowledge is corroborated by the Examiner's rejection ([Sechen ¶ 78], citing Gorobets '179 (Ex. 1044) in view of De Ambroggi (Ex. 1053) (RRAM, PCM) and Kund (Ex. 1055) (alternative technologies having advantages over flash memory).

(Sechen ¶ 284): Sutardja, in view of the knowledge of the POSITA, discloses claim 6.

G. Claim 7 Is Obvious Over Sutardja in View of Knowledge of the POSITA.

1. [7Pre] The system of claim 1,

See Section IX(A).

2. [7] wherein the SLC non-volatile memory module is resistive random-access memory (RRAM).

(Sechen ¶ 286): *See* [Sechen ¶¶ 282 and 283].

(Sechen ¶ 287): Sutardja, in view of the knowledge of the POSITA, discloses claim 7.

H. Claim 8 Is Obvious Over Sutardja in View of Knowledge of the POSITA.

1. [8Pre] The system of claim 1,

See Section IX(A).

2. [8] wherein the MLC non-volatile memory module is phase change memory (PCM).

(Sechen ¶ 289): To the extent that this claim is enabled by the disclosures in the specification, it is obvious in view of the assumed background knowledge.

(Sechen ¶ 290): This background knowledge is corroborated by the Examiner's rejection ([Sechen ¶ 78], citing Gorobets '179 (Ex. 1044) in view of De Ambroggi (Ex. 1053) (RRAM, PCM) and Kund (Ex. 1055) (alternative technologies having advantages over flash memory). Sutardja paragraphs [0103], [0104], and [0172] also mention PCM chips.

(Sechen ¶ 291): Sutardja, in view of the knowledge of the POSITA, discloses claim 8.

I. Claim 9 Is Obvious Over Sutardja in View of Knowledge of the POSITA.

1. [9Pre] The system of claim 1,

See Section IX(A).

2. [9] wherein the SLC non-volatile memory module is phase change memory (PCM).

(Sechen ¶ 293): *See* [Sechen ¶¶ 289 and 290].

(Sechen ¶ 294): Sutardja, in view of the knowledge of the POSITA, discloses claim 9.

J. Claim 10 Is Obvious Over Sutardja in View of Knowledge of the POSITA.

1. [10Pre] The system of claim 1,

See Section IX(A).

2. [10] wherein the SLC non-volatile memory module is magnetic random-access memory (MAGRAM).

(Sechen ¶ 296): To the extent that this claim is enabled by the disclosures in the specification, it is obvious in view of the assumed background knowledge.

(Sechen ¶ 297): This background knowledge is corroborated by the Examiner's rejection ([Sechen ¶ 78], citing Gorobets '179 (Ex. 1044) in view of Chen '418 (Ex. 1054) (MRAM), and Kund (Ex. 1055) (alternative technologies

having advantages over flash memory). Sutardja paragraph [0172] also mentions "magnetic RAM."

(Sechen ¶ 298): Sutardja, in view of the knowledge of the POSITA, discloses claim 10.

K. Claim 11 Is Obvious Over Sutardja in View of Knowledge of the POSITA.

3. [11Pre] The system of claim 1,

See Section IX(A).

4. [11] wherein the controller causes the transfer of content on a periodic basis.

(Sechen ¶ 300): There is ambiguity in the term "periodic." [Sechen ¶ 73.] To the extent that "periodic" means "occasional," the actions disclosed in Sutardja as reviewed at [Sechen ¶¶ 259-265] are "occasional" or responsive to conditions such as reaching a threshold.

(Sechen ¶ 301): As reviewed at [Sechen ¶¶ 36, 227 and 228], it was common knowledge of the POSITA that wear leveling functions may be performed at the time of write (dynamic) or in the background when the device is not performing host-initiated read and write operations. (Admitted art at 3:6-13.)

(Sechen ¶ 302): Sutardja specifically discloses that its system "determines whether [the] time to perform data shift analysis has arrived in step 514." (¶ [0148].) (*Also* ¶¶ [0135][periodic degradation testing] and [0167]["At various times, such as

periodically, the wear leveling module analyzes the wear levels of the blocks, and remaps relatively frequently rewritten logical addresses to blocks with low wear levels"].)

(Sechen ¶ 303): Sutardja, in view of the knowledge of the POSITA, discloses claim 11.

XII. THE PENDING LITIGATIONS IN THE WESTERN DISTRICT OF TEXAS DO NOT WARRANT DENYING INSTITUTION.

With eight current patents asserted by Vervain against Phison based on varying word claims to the same two "data integrity test" and "wear leveling" functions described in the same specification for a known hybrid SLC-MLC NAND flash memory, review by this Board is more appropriate than by the District Court where those word variations can escape scrutiny. The Board may take notice that construction of the multiple terms of the eight patents are restricted in number (12) in District Court (per "side" in related cases as present here), even counting contentions of indefiniteness. As demonstrated in the above-stated grounds (Sections VII, VIII, and IX), the case against the '298 patent is compelling. Balancing of the six factors set forth by the Board in *Apple Inc. v. Fintiv, Inc*., IPR2020-00019, Paper 11 (P.T.A.B. Mar. 20, 2020) (precedential) does not warrant denying institution here. Rather, institution would best serve the interest in "system efficiency, fairness, and patent quality" (*id.* at 5), allowing due resolution of the substantial patentability questions (*id.* at 15).

As to *Fintiv* Factor 5, the petitioner and defendant in the parallel proceeding are the same party. Phison has acted to file this petition expeditiously, where it received first notice of Vervain's assertion in the Complaint filed March 8, 2024, and the First Amended Complaint filed May 13, 2024 (Ex. 1082). The District Court has not issued any substantive opinions regarding the scope or validity of the Challenged Claims. The parallel proceeding is in an early stage: the Court has not yet decided Phison's motion to dismiss (*see* Docket List [Ex. 1080] entries 12 and 24), the Petitioners have not otherwise answered, fact discovery has not yet proceeded beyond automatic disclosure and Phison's invalidity and subject-matter eligibility contentions were filed on August 16, 2024, and claim construction materials exchanged in accordance with the Scheduling Order (Ex. 1083). A *Markman* hearing is initially scheduled for December 5, 2024. (*Id.*) Thus, as to Factor 3, there has not been overriding investment in the parallel proceeding; to the contrary, the requirements of the parallel proceeding limited Markman construction to twelve terms between the parties for eight patents of verbal variations of the same specification. Thus, as to Factor 4, there is overlap, in the greater inclusiveness of the requested review as the District Court will not consider all the issues properly raised in this Petition, the pending PGR2024-00047 of the '546 patent and PGR2024-00048 of the '612 patent, and others for review of the five other variant patents that Petitioner may petition.

Factors 1 (stay) and 2 (proximity of trial dates) do not significantly weigh for or against instituting the requested IPR. Petitioner does not know if the District Court will stay the case if trial by this Board is instituted and the Court has *tentatively* set jury selection for December 15, 2025 (Scheduling Order [Ex. 1083] at 4).

Regarding the sixth factor (merits, other circumstances), the merits strongly weigh in favor of instituting trial as shown through the strength of the grounds in this Petition. "Other circumstances" favor institution. As described at note 3 *supra*, there are separate suits by Patent Owner Vervain (following the same template) against Phison and a customer (Kingston Docket List [Ex. 1081]). This presents the complexity of eight patents with method and apparatus claims that Vervain asserts against multiple, sometimes competing products (that thus must be held confidential between the defendants) that implicate third-party actors who may be implicated in divided or indirect infringement as users who select the configuration or performance of the claimed system. Moreover, unlike Phison, its customer defendant in the other case has moved to dismiss for improper venue and to transfer (Kingston Docket List [Ex. 1081] entries 13 and 20) with attendant discovery and confidentiality proceedings (*id*. entries 15 and 16).

The interest of the public in cleaning away such verbal variants of questionable innovation in the eight Vervain patents are best served by the requested review. As with Factor 4, the Board may take notice that such litigation, if it reaches

trial, will have been drastically reduced in claims and theories asserted – this begins with the limitation of terms for construction for eight patents. An IPR trial, in contrast, allows a focus on resolving all Challenged Claims in a single patent, and may substantially simplify district court trial without the extensive district court discovery for which the America Invents Act offered this alternative before the Board. And, of course, Petitioner has shown its case to be compelling.

XIII. CONCLUSION

Petitioner submits that for the reasons set forth above, supported by the declaration of Dr. Sechen and the Exhibits, it has been shown that more likely than not, Challenged Claims 1-11 are invalid under all of 35 U.S.C. § 103 and should reviewed by the Board and canceled.

Date: November 25, 2024

Respectfully submitted,
Attorneys for Petitioner

/Hsuanyeh Chang, Ph.D./

Hsuanyeh Chang, Ph.D. (Reg. No. 73,431)
hsuanyeh@hsuanyeh.com
Stephen Y. Chow (Reg. No. 31,338)
Stephen.Y.Chow@hsuanyeh.com
Douglas E. Chin (Reg. No. 66,713)
Doug.Chin@hsuanyeh.com
Peter Yi (Reg. No. 61,790)
Peter.Yi@hsuanyeh.com
Hsuanyeh Law Group, PC
11 Beacon Street, Suite 900
Boston, MA 02108
(617) 886-0988

CERTIFICATION OF SERVICE (37 C.F.R. § 42.6(e))

The undersigned hereby certifies that the above-captioned Petition for *Inter Partes* Review of U.S. Patent No. 8,891,298 (and accompanying exhibits), was served in its entirety on November 25, 2024, upon Patent Owner's correspondent of record via overnight courier:

Patent owner's correspondent
address of record for U.S. Patent No. 11,854,298

BILL R. NAIFEH
P.O. BOX 803423
DALLAS, TX 75380

/Hsuanyeh Chang, Ph.D./
Hsuanyeh Chang, Ph.D. (Reg. No. 73,431)
hsuanyeh@hsuanyeh.com
Stephen Y. Chow (Reg. No. 31,338)
Stephen.Y.Chow@hsuanyeh.com
Douglas E. Chin (Reg. No. 66,713)
Doug.Chin@hsuanyeh.com
Peter Yi (Reg. No. 61,790)
Peter.Yi@hsuanyeh.com
Hsuanyeh Law Group, PC
11 Beacon Street, Suite 900
Boston, MA 02108
(617) 886-9088

CERTIFICATION OF WORD COUNT (37 C.F.R. § 42.6(e))

The undersigned hereby certifies that the attached Petition, including footnotes, but not the cover page, exhibit list, table of contents, mandatory notices, and certifications, contains 13,875 words, as measured by the Word Count function of Microsoft Word. This is less than the limit of 14,000 words as specified by 37 C.F.R. § 42.24(a)(i).

/Hsuanyeh Chang, Ph.D./
Hsuanyeh Chang, Ph.D. (Reg. No. 73,431)