IN THE UNITED STATES DISTRICT COURT
FOR THE WESERN DISTRICT OF TEXAS
AUSTIN DIVISION

POLARIS POWERLED TECHNOLOGIES, LLC, Plaintiff, v. WESTERN DIGITAL CORPORATION and WESTERN DIGITAL TECHNOLOGIES, INC., Defendants.	§ § § § § § § § § § § § §	Civil Action No. 1:24-cv-00101 **JURY TRIAL DEMANDED**

COMPLAINT FOR PATENT INFRINGEMENT AGAINST DEFENDANTS WESTERN DIGITAL CORPORATION AND WESTERN DIGITAL TECHNOLOGIES, INC.

Plaintiff Polaris PowerLED Technologies, LLC ("Plaintiff" or "Polaris") brings this patent infringement action against Defendants Western Digital Corporation and Western Digital Technologies, Inc. (collectively, "Defendants" or "Western Digital") as follows:

NATURE OF THE ACTION

1. This is a civil action for infringement of U.S. Patent Nos. 8,554,968 ("the '968 Patent"), 9,183,085 ("the '085 Patent"), and 8,601,346 ("the '346 Patent"), and under the patent laws of the United States, 35 U.S.C. § 1 *et seq*.

2. The technology in this case involves flash memory. Flash memory is used in, among other things, computing, gaming, data storage, and e-commerce. The ability of solid-state flash memory to hold electric charges without moving parts has revolutionized how information is stored and has resulted in great improvements over older memory technologies. These electrical charges, held in miniscule transistors, are used to read, store, and write enormous amounts of

information in small, lightweight memory products that have transformed the daily lives of consumers.

3. Defendants' infringing devices are its solid-state drive ("SSD") products ("the Accused Products" or "the infringing devices").

4. Western Digital is among the largest manufacturers of memory products in the United States.

5. Plaintiff brings this patent infringement action to protect its valuable patented technology specifically relating to (1) nonvolatile memory controllers (NVMCs) and SSDs; (2) interrupt techniques used in NVMCs and SSDs; (3) how NVMCs and SSDs adaptively select among error correction coding (ECC) schemes; and (4) how NVMCs and SSDs generate parity data using a distributed processing technique.

THE PARTIES

6. Polaris is a California limited liability company having its address at 5150 E. Pacific Coast Highway, Suite 200, Long Beach, California 90804.

7. On information and belief, Defendant Western Digital Corporation is a Delaware corporation having offices and principal places of business at 7501 N. Capital of Texas Highway, Suite A-100, Austin, Texas 78731 and 5601 Great Oaks Parkway, San Jose, California 95119.

8. On information and belief, Defendant Western Digital Technologies, Inc. is a subsidiary of Western Digital Corporation. Western Digital Technologies, Inc. is a Delaware corporation having offices and principal places of business at 7501 N. Capital of Texas Highway, Suite A-100, Austin, Texas 78731 and 5601 Great Oaks Parkway, San Jose, California 95119. Western Digital Technologies, Inc. is registered with the Texas Secretary of State to do business in Texas and can be served through its registered agent, The Corporation Service Company d/b/a

CSC-Lawyers Incorporating Service Company, 211 E. 7th Street, Suite 620, Austin, Texas 78701-3218.

9. The Defendants control, participate in the commission of, and have a direct financial interest in the infringing acts set forth herein.

JURISDICTION AND VENUE

10. Plaintiff incorporates by reference and re-alleges the foregoing paragraphs as fully set forth herein.

11. This action arises under the patent laws of the United States, Title 35 of the United States Code. Accordingly, this Court has subject matter jurisdiction under 28 U.S.C. §§ 1331 and 1338(a).

12. Venue is proper in this judicial District pursuant to 28 U.S.C. §§ 1391(b), (c) and 1400(b). On information and belief, Defendants have transacted business in this District and have committed acts of direct and indirect infringement in this District by, among other things, making, using, offering to sell, selling, and/or importing products that infringe Polaris's patents as set forth below. Defendants have at least one regular and established place of business in this District. Defendants maintain operations out of their leased property at 7501 N. Capital of Texas Highway, Suite A-100, Austin, Texas 78731. Western Digital Corporation has admitted in another case that its subsidiary Western Digital Technologies, Inc. leases office space in Austin, Texas.[1]

13. Defendants have not maintained corporate separateness. On information and belief, Western Digital Corporation and Western Digital Technologies, Inc. have employees in common.

[1] Defendants' Answers and Defenses at ¶ 5, *Vervain, LLC v. Western Digital Corporation*, No. 6:21-cv-00488-ADA (W.D. Tex. July 16, 2021), ECF No. 20; *id*., Notice of Correction, (May 2, 2022), ECF No. 51.

As an example, Michael Charles Ray has previously signed Registration Statements on behalf of Western Digital Corporation and Western Digital Technologies, Inc.[2]

14. Moreover, Western Digital Technologies, Inc. has guaranteed some of Western Digital Corporation's securities without receiving consideration in return. Further, the Standard Terms and Conditions for Defendants' Performance Stock Units and Restricted Stock Units follow the Dispute Resolution Agreement made by Western Digital Technologies, Inc.[3] Western Digital Technologies, Inc. follows Western Digital Corporation's corporate policy, as evidenced by the published Western Digital Corporation Code of Business Ethics that applies to employees of both entities.[4] For purposes of infringement, there is no discernable difference between Western Digital Corporation and Western Digital Technologies, Inc.

15. Western Digital Technologies, Inc.'s presence and acts of infringement committed in this District are attributable and imputed to Western Digital Corporation for venue purposes. On information and belief, the Defendants function as an integrated organization in the operation of Defendants' business operations with respect to the infringing actions complained of herein.

16. The Court has personal jurisdiction over Defendants consistent with the Texas Long Arm Statute and the Due Process Clause of the Fourteenth Amendment. Defendants have purposely availed themselves of the benefits and protections of Texas. Defendants have maintained

---

[2] U.S. Securities and Exchange Commission, Form S-3 Registration Statement, Western Digital Corporation (Jan. 29, 2018), https://investor.wdc.com/node/19351/html, last accessed Jan 30, 2024.

[3] 2017 Notice of Grant of Performance Stock Units and Performance Stock Unit Award Agreement – TSR Measure, https://www.sec.gov/Archives/edgar/data/106040/000010604018000034/wdc-2019q1ex102.htm, last accessed Jan. 30, 2024.

[4] Western Digital, *Code of Business Ethics* (Feb. 4, 2015), https://documents.westerndigital.com/content/dam/doc-library/en_us/assets/public/western-digital/collateral/company/western-digital-code-of-business-ethics.pdf, last accessed Jan. 30, 2024.

a presence in this District for years.

17. Tax records from 2019 and 2023 show that Western Digital Corporation is Western Digital Technologies, Inc.'s parent company, and maintains a presence in Texas.[5]

18. On information and belief, Western Digital Corporation owns and operates the website www.westerndigital.com, on which it offers, distributes, and markets its memory products across the country, including within this District. The Western Digital website makes no distinction between Western Digital Corporation and Western Digital Technologies, Inc. regarding its products, services, employees, or leadership.

19. Western Digital Technologies, Inc. has admitted that it is registered with the Texas Secretary of State to conduct business in Texas, and can be served through the Corporation Service Company d/b/a CSC-Lawyers Incorporating Service Company, located at 211 E. 7th Street, Suite 620 in Austin, Texas 78701-3218.[6]

20. On information and belief, Defendants employ people in this District that further the usage and sale of the infringing products. This includes, upon information and belief, Chief Financial Officer Wissam Jabre, Channel Sales Representative Avery Bell, and Senior Director Bud Koch.[7] Defendants also employ Scott Glenn, the Director of Global Channel and OEMs, who also leads the marketing efforts for Western Digital's SSD enterprise storage products.[8]

---

[5] Texas Franchise Tax Public Information Report, Western Digital Technologies, Inc. (May 14, 2019); Texas Franchise Tax Public Information Report, Western Digital Technologies, Inc. (May 12, 2023).

[6] Defendants' Answers and Defenses at ¶ 16, *Viasat, Inc. v. Western Digital Corporation, Western Digital Technologies, Inc.*, No. 6:21-cv-01230-ADA (Jan. 6, 2023), ECF No. 117.

[7] See Wissam Jabre, LinkedIn, https://www.linkedin.com/in/wissam-jabre-cfa-b721a8/; Avery Bell, LinkedIn, https://www.linkedin.com/in/avery-bell7/; Bud Koch, LinkedIn, https://www.linkedin.com/in/budkoch/ (last visited Jan. 2, 2024).

[8] Scott Glenn, LinkedIn, https://www.linkedin.com/in/sglenn1/, (last visited Jan. 2, 2024).

21. On information and belief, Defendants also engage in the recruitment of individuals through LinkedIn and other websites. University Recruiter Austin Painchaud works out of nearby Georgetown, Texas[9] in furtherance of these efforts.

22. Defendants conduct continuous and systematic business in this District, including, among other acts, offering infringing products and services to those residing in this District and soliciting business from people residing in this District. Defendants make infringing sales of the accused products in this District. Defendants have committed infringing acts within the Western District of Texas giving rise to this action and has established minimum contacts within the forum state of Texas.

23. This Court has general jurisdiction over Defendants due to its continuous and systematic contacts with the State of Texas, including its ownership and/or lease of land in the State of Texas, and other business activities throughout the State of Texas.

PATENTS-IN-SUIT

24. Plaintiff owns the entire right, title, and interest in U.S. Patent No. 8,554,968 titled "Interrupt Technique for a Nonvolatile Memory Controller," including the right to assert all causes of action arising under said patent and to seek damages and all other remedies for the infringement thereof. The '968 Patent issued on October 8, 2013 to inventors Peter Z. Onufryk, Jayesh Patel and Ihab Jaser from the U.S. Patent Application No. 13/052,388 filed on March 21, 2011. A true and correct copy of the '968 Patent is attached as Exhibit A to this Complaint.

25. Plaintiff owns the entire right, title, and interest in U.S. Patent No. 9,183,085 titled "Systems and Methods for Adaptively Selecting from among a Plurality of Error Correction

[9] Austin Painchaud, LinkedIn, https://www.linkedin.com/in/austinpainchaud/, (last visited Jan. 2, 2024).

Coding Schemes in a Flash Drive for Robustness and Low Latency," including the right to assert all causes of action arising under said patent and to seek damages and all other remedies for the infringement thereof. The '085 Patent issued on November 10, 2015, to inventor Philip L. Northcott from the U.S. Patent Application No. 13/477,600, filed on May 22, 2012. A true and correct copy of the '085 Patent is attached as Exhibit B to this Complaint.

26. Plaintiff owns the entire right, title, and interest in U.S. Patent No. 8,601,346 titled "System and Method for Generating Parity Data in a Nonvolatile Memory Controller by Using a Distributed Processing Technique," including the right to assert all causes of action arising under said patent and to seek damages and all other remedies for the infringement thereof. The '346 Patent issued on December 3, 2013, to inventors Peter Z. Onufryk and Inna Levit from the U.S. Patent Application 13/052,835, filed on March 21, 2011. A true and correct copy of the '346 Patent is attached as Exhibit C to this Complaint.

'968 PATENT BACKGROUND

27. The '968 Patent relates generally to improvements to nonvolatile memory controllers, including the processing of memory commands and the generation of a completion status for such commands. The inventive nonvolatile memory controller transmits the completion status to a host processing unit for storage in a completion queue of the host processing unit. An interrupt manager in the nonvolatile memory controller determines whether the completion queue contains an unprocessed completion status and generates an interrupt message packet. The nonvolatile memory controller transmits the interrupt message packet to the host processing unit for triggering an interrupt in the host processing unit and alerting the host processing unit to the unprocessed completion status.

’085 PATENT BACKGROUND

28. The ’085 Patent relates generally to disclosed techniques providing relatively low uncorrectable bit error rates (BER) for flash memory; low write amplification; long life, fast and efficient retrieval; and efficient storage density such that a solid-state drive (SSD) or flash drive can be implemented using relatively inexpensive MLC flash for enterprise storage application.

’346 PATENT BACKGROUND

29. The ’346 Patent relates generally to a nonvolatile memory controller performing a data-stripe operation by processing a collection of commands. The collection of commands includes data update commands and a parity write command. The nonvolatile memory controller includes a number of command processing units, each of which receives a command in the collection of commands. Each of the command processing units receiving a data update command requests a data block from a controller memory; receives the data block from the controller memory through a data path in response to the request; and writes the data block to a nonvolatile memory device.

COUNT I

(INFRINGEMENT OF U.S. PATENT NO. 8,554,968)

30. Plaintiff incorporates by reference and re-alleges the foregoing paragraphs as fully set forth herein.

31. Defendants have directly infringed, and continue to directly infringe, one or more claims of the ’968 Patent, including at least claim 1 of the ’968 Patent, literally and/or under the doctrine of equivalents, by or through making, using, offering for sale, selling within the United States and/or importing into the United States its SSD products that support NVMe.

32. Claim 1 of the ’968 Patent, for example, reads as follows:

> 1. A nonvolatile memory controller for alerting a host processing unit to an unprocessed completion status contained in a completion queue of the host processing unit, the nonvolatile memory controller comprising:
>
> an interrupt manager configured to generate a completion queue state for indicating the occurrence of a completion queue event associated with the completion queue, generate an interrupt vector state based on the completion queue state, determine the completion queue of the host processing unit contains an unprocessed completion status based on the interrupt vector state, and generate an interrupt message packet for triggering an interrupt in the host processing unit to alert the host processing unit of the unprocessed completion status in the completion queue, and wherein the completion queue state includes a doorbell update status indicating whether the host processing unit has performed a doorbell update event in which the host processing unit updates a head pointer stored in the nonvolatile memory controller for the completion queue.

33. Defendants make, use, sell, offer for sale and import SSD products that support the NVMe standard and that include a nonvolatile memory controller configured for alerting a host processing unit to an unprocessed completion status contained in a completion queue of the host processing unit. For example, as shown below, WD Black-branded SSDs, which are used for gaming consoles and gaming PCs, Blue- and Green-branded SSDs, which are used for everyday PCs, Red-branded SSDs, which are used for Network Attached Storage (NAS), and Gold- and Ultrastar SSDs , which are used for data centers, all support NVMe:

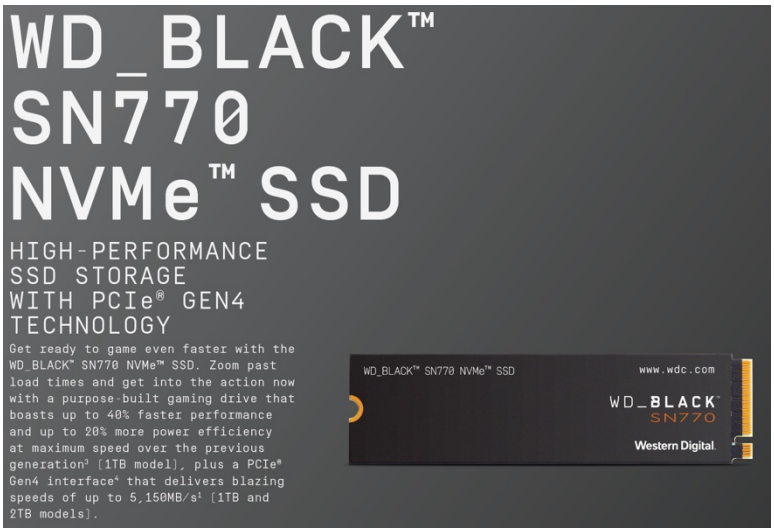

WD_BLACK SN770 NVMe SSD, Western Digital (Jan. 2022), https://documents.westerndigital.com/content/dam/doc-library/en_us/assets/public/western-

digital/product/internal-drives/wd-black-ssd/product-brief-wd-black-sn770-nvme-ssd.pdf [hereinafter WD_BLACK SN770 NVMe SSD].

34. The image below shows the benefits and improvements of NVMe over SATA:

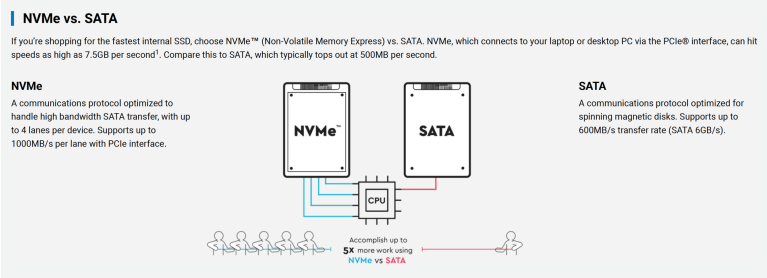

Internal SSD, Western Digital, https://www.westerndigital.com/solutions/internal-ssd, (last visited Jan. 2, 2024) [hereinafter Internal SSD].

35. The following describes how NVMe functions:

1.4 Theory of Operation

The NVM Express scalable interface is designed to address the needs of Enterprise and Client systems that utilize PCI Express based solid state drives or fabric connected devices. The interface provides optimized command submission and completion paths. It includes support for parallel operation by supporting up to 65,535 I/O Queues with up to 64 Ki - 1 outstanding commands per I/O Queue. Additionally, support has been added for many Enterprise capabilities like end-to-end data protection (compatible with SCSI Protection Information, commonly known as T10 DIF, and SNIA DIX standards), enhanced error reporting, and virtualization.

The interface has the following key attributes:

- Does not require uncacheable / MMIO register reads in the command submission or completion path;
- A maximum of one MMIO register write is necessary in the command submission path;
- Support for up to 65,535 I/O Queues, with each I/O Queue supporting up to 65,535 outstanding commands;
- Priority associated with each I/O Queue with well-defined arbitration mechanism;
- All information to complete a 4 KiB read request is included in the 64B command itself, ensuring efficient small I/O operation;
- Efficient and streamlined command set;
- Support for MSI/MSI-X and interrupt aggregation;
- Support for multiple namespaces;
- Efficient support for I/O virtualization architectures like SR-IOV;
- Robust error reporting and management capabilities; and
- Support for multi-path I/O and namespace sharing.

This specification defines a streamlined set of registers whose functionality includes:

- Indication of controller capabilities;
- Status for controller failures (command status is processed via CQ directly);
- Admin Queue configuration (I/O Queue configuration processed via Admin commands); and
- Doorbell registers for scalable number of Submission and Completion Queues.

The NVM Express interface is based on a paired Submission and Completion Queue mechanism. Commands are placed by host software into a Submission Queue. Completions are placed into the associated Completion Queue by the controller. Multiple Submission Queues may utilize the same Completion Queue. Submission and Completion Queues are allocated in memory.

An Admin Submission and associated Completion Queue exist for the purpose of controller management and control (e.g., creation and deletion of I/O Submission and Completion Queues, aborting commands, etc.). Only commands that are part of the Admin Command Set may be submitted to the Admin Submission Queue.

NVM Express, Inc., *NVM Express Base Specification, Revision 1.4*, (June 10, 2019), https://nvmexpress.org/wp-content/uploads/NVM-Express-1_4-2019.06.10-Ratified.pdf

[hereinafter NVM Express Specification].

Host software submits commands to a controller through pre-allocated Submission Queues. A controller is alerted to newly submitted commands through SQ Tail Doorbell register writes. The difference between the previous doorbell register value and the current register write indicates the number of commands that were submitted.

Upon completion of the command execution by the NVM subsystem, the controller presents completion queue entries to the host through the appropriate Completion Queues. If MSI-X or multiple message MSI is in use, then the interrupt vector indicates the Completion Queue(s) with possible new command completions for the host to process. If pin-based interrupts or single message MSI interrupts are used, host software interrogates the Completion Queue(s) for new completion queue entries. The host updates the CQ Head doorbell register to release Completion Queue entries to the controller and to clear the associated interrupt.

NVM Express Specification at 274.

36. Defendants' SSD products with NVMe support include an interrupt manager configured to generate a complete queue state for indicating occurrence of a completion queue event associated with the completion queue, as shown below:

1.6.9 controller

A controller is the interface between a host and an NVM subsystem. There are three types of controllers:

a) I/O controllers;
b) discovery controllers; and
c) administrative controllers.

A controller executes commands submitted by a host on a Submission Queue and posts a completion on a Completion Queue. All controllers implement one Admin Submission Queue and one Admin Completion

NVM Express Specification at 14.

5. After a command has completed execution, the controller places a completion queue entry in the next free slot in the associated Completion Queue. As part of the completion queue entry, the controller indicates the most recent Submission Queue entry that has been consumed by advancing the Submission Queue Head pointer in the completion entry. Each new completion queue entry has a Phase Tag inverted from the previous entry to indicate to the host that this completion queue entry is a new entry;
6. The controller optionally generates an interrupt to the host to indicate that there is a new completion queue entry to consume and process. In the figure, this is shown as an MSI-X interrupt, however, it could also be a pin-based or MSI interrupt. Note that based on interrupt coalescing settings, an interrupt may or may not be generated for each new completion queue entry;
7. The host consumes and then processes the new completion queue entries in the Completion Queue. This includes taking any actions based on error conditions indicated. The host continues consuming and processing completion queue entries until a previously consumed entry with a Phase Tag inverted from the value of the current completion queue entries is encountered; and
8. The host writes the Completion Queue Head Doorbell register to indicate that the completion queue entry has been consumed. The host may consume many entries before updating the associated Completion Queue Head Doorbell register.

NVM Express Specification at 283.

Figure 432: Command Processing

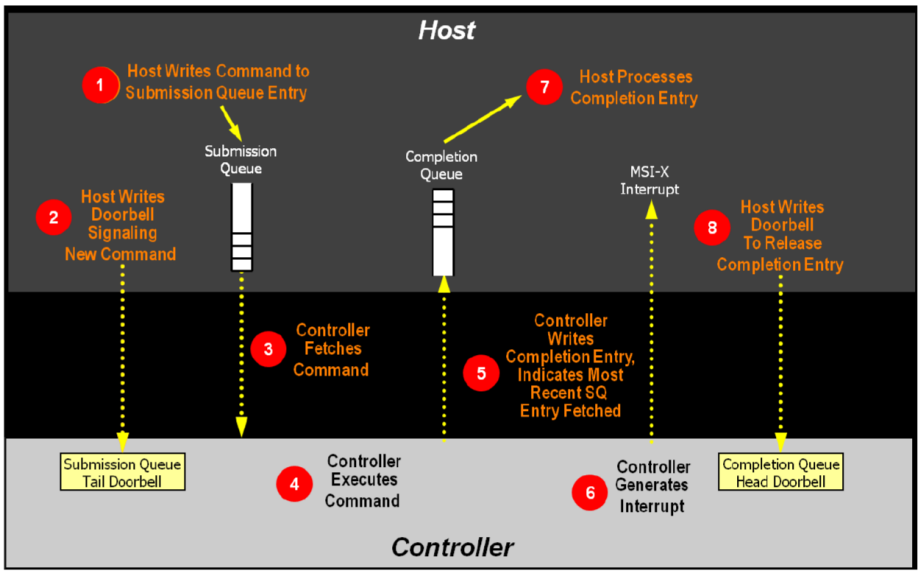

NVM Express Specification at 284.

37. Defendants' SSD products with NVMe support include an interrupt manager configured to generate an interrupt vector state based on the complete queue state, which is shown below:

Figure 151: Create I/O Completion Queue – Command Dword 11

Bits	Description
31:16	**Interrupt Vector (IV):** This field indicates interrupt vector to use for this Completion Queue. This corresponds to the MSI-X or multiple message MSI vector to use. If using single message MSI or pin-based interrupts, then this field shall be cleared to 0h. In MSI-X, a maximum of 2,048 vectors are used. This value shall not be set to a value greater than the number of messages the controller supports (refer to MSICAP.MC.MME or MSIXCAP.MXC.TS). If the value is greater than the number of messages the controller supports, the controller should return an error of Invalid Interrupt Vector.

NVM Express Specification at 102.

5. After a command has completed execution, the controller places a completion queue entry in the next free slot in the associated Completion Queue. As part of the completion queue entry, the controller indicates the most recent Submission Queue entry that has been consumed by advancing the Submission Queue Head pointer in the completion entry. Each new completion queue entry has a Phase Tag inverted from the previous entry to indicate to the host that this completion queue entry is a new entry;
6. The controller optionally generates an interrupt to the host to indicate that there is a new completion queue entry to consume and process. In the figure, this is shown as an MSI-X interrupt, however, it could also be a pin-based or MSI interrupt. Note that based on interrupt coalescing settings, an interrupt may or may not be generated for each new completion queue entry;
7. The host consumes and then processes the new completion queue entries in the Completion Queue. This includes taking any actions based on error conditions indicated. The host continues consuming and processing completion queue entries until a previously consumed entry with a Phase Tag inverted from the value of the current completion queue entries is encountered; and
8. The host writes the Completion Queue Head Doorbell register to indicate that the completion queue entry has been consumed. The host may consume many entries before updating the associated Completion Queue Head Doorbell register.

NVM Express Specification at 283.

Figure 432: Command Processing

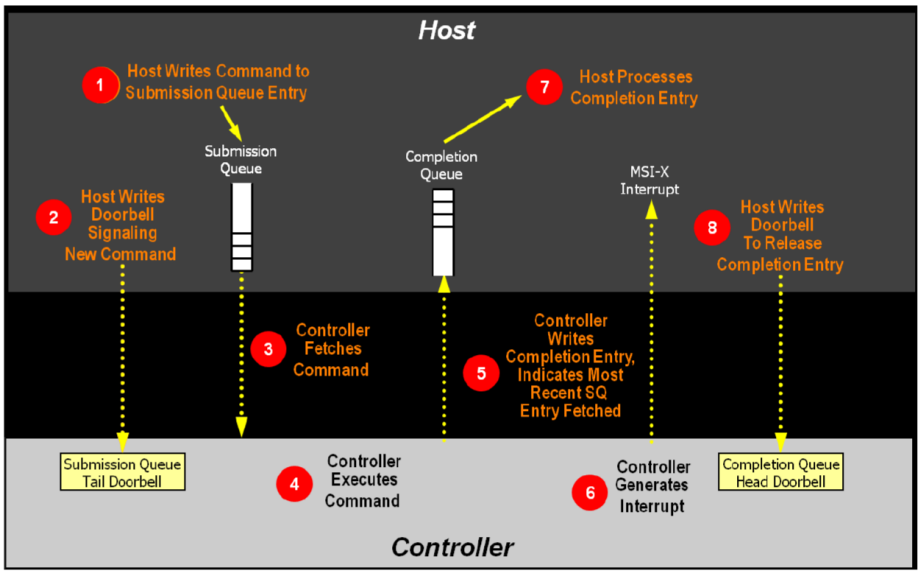

NVM Express Specification at 284.

38. The Accused Products with NVMe support include an interrupt manager configured to determine whether the completion queue of the host processing unit contains an unprocessed completion state based on the interrupt vector state. The interrupt manager is further configured to generate an interrupt message packet for triggering an interrupt in the host processing unit to alert the host processing unit of the unprocessed completion status in the completion queue. This is shown in the images below:

6. The controller optionally generates an interrupt to the host to indicate that there is a new completion queue entry to consume and process. In the figure, this is shown as an MSI-X interrupt, however, it could also be a pin-based or MSI interrupt. Note that based on interrupt coalescing settings, an interrupt may or may not be generated for each new completion queue entry;
7. The host consumes and then processes the new completion queue entries in the Completion Queue. This includes taking any actions based on error conditions indicated. The host continues consuming and processing completion queue entries until a previously consumed entry with a Phase Tag inverted from the value of the current completion queue entries is encountered; and
8. The host writes the Completion Queue Head Doorbell register to indicate that the completion queue entry has been consumed. The host may consume many entries before updating the associated Completion Queue Head Doorbell register.

Figure 432: Command Processing

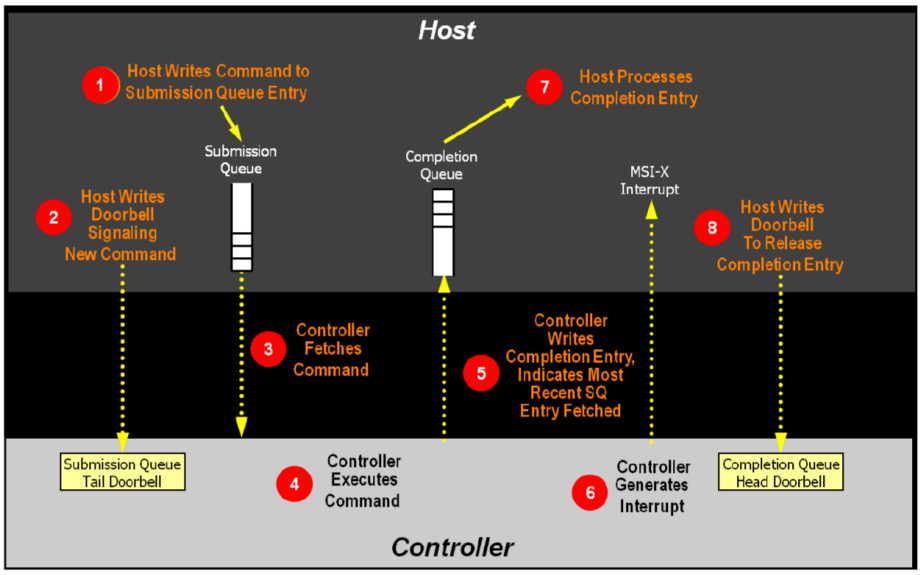

NVM Express Specification at 283-284.

39. The nonvolatile memory controller, which generates an MSI-X interrupt, is transmitted as a PCIe message packet. This is demonstrated in the quote below:

> NVMe, which connects to your laptop or desktop PC via the PCIe® interface, can hit speeds as high as 7.5GB per second. Compare this to SATA, which typically tops out at 500MB per second.

Internal SSD

PCI Express Base Specification Revision 2.1 (March 4, 2009), [hereinafter PCIe v2.1].

Message Signaled Interrupt (MSI/MSI-X)	Two similar but separate mechanisms that enable a Function to request service by writing a system-specified DWORD of data to a system-specified address using a Memory Write Request. Compared to MSI, MSI-X supports a larger maximum number of vectors and independent message address and data for each vector.
Packet	A fundamental unit of information transfer consisting of a header that, in some cases, is followed by a data payload.

PCIe v2.1 at 31.

Transaction Layer	The Layer that operates at the level of transactions (for example, read, write).
Transaction Layer Packet, TLP	A Packet generated in the Transaction Layer to convey a Request or Completion.

PCIe v2.1 at 33.

40. The image below shows a completion queue state includes a doorbell update status indicating whether the host processing unit has performed a doorbell update event in which the

host processing unit updates a head pointer stored in the nonvolatile memory controller for the completion queue.

6. The controller optionally generates an interrupt to the host to indicate that there is a new completion queue entry to consume and process. In the figure, this is shown as an MSI-X interrupt, however, it could also be a pin-based or MSI interrupt. Note that based on interrupt coalescing settings, an interrupt may or may not be generated for each new completion queue entry;
7. The host consumes and then processes the new completion queue entries in the Completion Queue. This includes taking any actions based on error conditions indicated. The host continues consuming and processing completion queue entries until a previously consumed entry with a Phase Tag inverted from the value of the current completion queue entries is encountered; and
8. The host writes the Completion Queue Head Doorbell register to indicate that the completion queue entry has been consumed. The host may consume many entries before updating the associated Completion Queue Head Doorbell register.

Figure 432: Command Processing

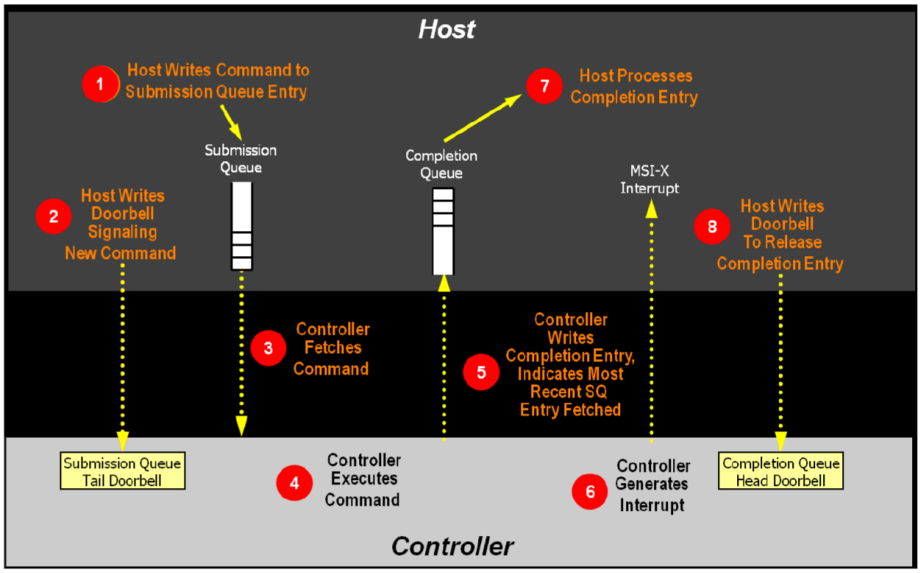

NVM Express Specification at 283-284.

41. Defendants have directly infringed, and continue to infringe, the claims of the '968 Patent in the United States, by making, using, offering for sale, selling, and/or importing the Accused Products in violation of 35 U.S.C. § 271(a). For example, Defendants provide SSD products configured with the hardware and software that satisfy the limitations of at least claim 1. Defendants further directly infringe the '968 Patent when its SSD products with NVMe are installed and operated by its employees in a computer system, such as for gaming, personal computing, or data centers. Direct infringement further occurs when Defendants' employees use and test the hardware and software.

42. Defendants also have infringed, and continue to infringe the claims of the '968 Patent by actively inducing others to use the Accused Products. Defendants' users, customers, agents or other third parties who use the Accused Products in accordance with Defendants' instructions infringe the claims of the '968 Patent, in violation of 35 U.S.C. § 271(a). Defendants induce its customers to use its SSD products with NVMe for the benefits of higher performance, reduced size and power, and increased reliability compared to SSD products with other interfaces, such as SATA. Defendants are thereby liable for infringement of the '968 Patent under 35 U.S.C. § 271(b).

43. Defendants' users, customers, agents or other third parties who use the Accused Products in accordance with Defendants' instructions infringe the claims of the '968 Patent, in violation of 35 U.S.C. § 271(b). Defendants intentionally instruct its customers to infringe through support information, demonstrations, brochures, videos, and user guides, such as those located at:
https://www.westerndigital.com/support;
https://support.wdc.com/contact.aspx?lang=en;
https://www.westerndigital.com/support/category-selection;
https://www.youtube.com/channel/UC_M0BlntDVSblWg6ggUa7UQ;
https://www.youtube.com/watch?v=L0zZtwRFj0E;
https://www.youtube.com/westerndigital; and
https://www.youtube.com/c/westerndigitalcorporation.

44. Defendants are on notice of its infringement by no later than the filing and service of this Complaint. By the time of trial, Defendants will have known and intended (since receiving such notice) that its continued actions would actively induce to the infringement of the '968 Patent. Defendants were further aware of the '968 Patent because it was cited as prior art by Western Digital in U.S. Patent Nos. 10,725,835, 10,296,249, 10,452,278, 10,466,903, and 10,509,569, which are all assigned to Western Digital.

45. On January 23, 2024, Plaintiff sent a notice letter to Defendants, notifying them

that they were infringing the Asserted Patents. Defendants' infringement of the '968 Patent has been willful and intentional under the standard announced in *Halo Elecs., Inc. v. Pulse Elecs., Inc.*, 579 U.S. 93, 136 S. Ct. 1923, 195 L. Ed 2d 278 (2016). Defendants have willfully infringed the '968 Patent by refusing to take a license and continuing to make, use, test, sell, license, and/or offer for sale/license the '968 Patent Accused Products. Instead of taking a license, Defendants have opted to make the business decision to "efficiently infringe" the '968 Patent. In doing so, Defendants willfully infringe the '968 Patent.

46. Defendants are liable as contributory infringers of the '968 Patent under 35 U.S.C. § 271(c) by having offered to sell, sold and imported and continuing to offer to sell, selling, and importing into the United States its SSD products that support NVMe, to be especially made or adapted for use in an infringement of the '968 Patent. Defendants' SSD products are key components in gaming consoles and gaming PCs, computers, everyday PCs, Network Attached Storage (NAS) devices, servers and data centers. These SSD products are material components for use in practicing the '968 Patent and are specifically made and are not a staple article of commerce suitable for substantial noninfringing use. Defendants supplied these components with knowledge of the '968 Patent and with knowledge that these components constitute material parts of the claimed inventions of the '968 Patent.

47. As a result of Defendants' infringement of the '968 Patent, Plaintiff has suffered monetary damages and is entitled to no less than a reasonable royalty for Defendants' use of the claimed inventions of the '968 Patent, together with interest and costs as determined by the Court. Plaintiff will continue to suffer damages in the future.

48. Defendants' acts of direct and indirect infringement have caused and continue to cause damage to Plaintiff. Plaintiff is entitled to damages in accordance with 35 U.S.C. §§ 271 and

281 sustained as a result of Defendants' wrongful acts in an amount to be proven at trial.

COUNT II

(INFRINGEMENT OF U.S. PATENT NO. 9,183,085)

49. Plaintiff incorporates by reference and re-alleges the foregoing paragraphs as fully set forth herein.

50. Defendants have directly infringed and continue to directly infringe the '085 Patent, literally and/or under the doctrine of equivalents, by making, using, offering for sale, selling, and/or importing into the United States its SSD products. Defendants directly infringe claim 1 of the '085 Patent when its SSD products are installed and operated by its employees in a computer system, such as for gaming, personal computing, or data centers. Defendants directly infringe claim 1 when Defendants' employees use and test its SSD products.

51. Claim 1 of the '085 Patent reads as follows:

> 1. A method of selecting an error correction coding (ECC) scheme, the method comprising:
>
> determining a bit error rate associated with a region comprising a fixed number of two or more flash memory pages or integer fractions thereof, wherein the two or more flash memory pages of a region can be read simultaneously, wherein the region stores at least data payload and primary and secondary ECC parity symbols corresponding to the data payload of the region;
>
> comparing the determined bit error rate to one or more predetermined thresholds corresponding to a set of predefined gears comprising at least a first gear and a second gear, wherein the predefined gears correspond to different predefined ECC schemes, wherein the first gear has a different data payload size and correction capability than the second gear, wherein the amount of memory space allocated for the storage of data payload within the region varies between the first gear and the second gear to accommodate a varying number of parity symbols between the first gear and the second gear; and
>
> selecting a gear from the set for the region based at least partly on the comparisons to the one or more predetermined thresholds;
>
> wherein determining, comparing, and selecting are performed by an integrated circuit.

52. Defendants' SSD products perform a method of selecting an error correction coding (ECC) scheme. For example, WD Black-branded SSDs, which are used for gaming consoles and gaming PCs, Blue- and Green-branded SSDs, which are used for everyday PCs, Red-branded SSDs, which are used for Network Attached Storage (NAS), and Gold- and Ultrastar SSDs, which are used for data centers, all perform a method of selecting an error correction coding (ECC) scheme, as shown below:

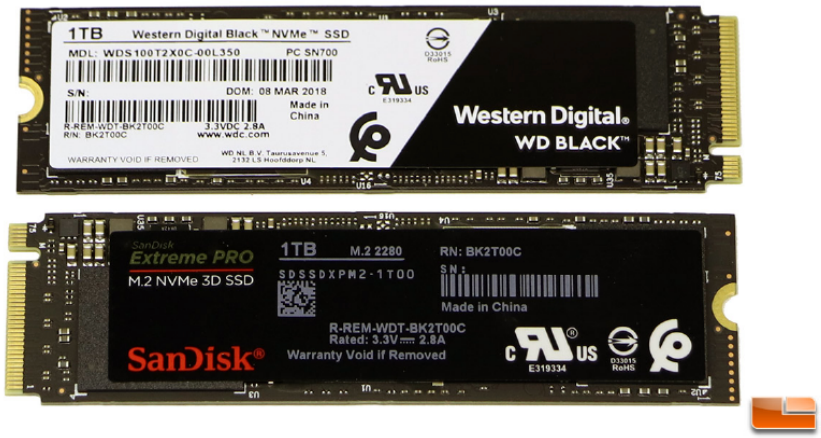

> Western Digital's new controller architecture means that the company has in-house development from the controller to the NAND. This gives WD tremendous design flexibility and they believe this new PCIe Gen 3 x4 controller is scalable and future-ready (think QLC NAND). Since WD designed everything they were able to design for low power and low latency as well as general performance. WD's new controller has three cores, is built on the 28nm process, features nCache 3.0 (marketing term for SLC Cache) for enhanced burst speeds, multi-gear LDPC (low-density parity-check) and hardware ECC to provide full SRAM and DRAM bit correction.

WD Black and SanDisk Review

> The Sentinel ECC&DSP LDPC engine is generic and can support different memory types (SLC, MLC, TLC, QLC) in terms of the page size and ECC redundancy. The Sentinel ECC&DSP LDPC codes provide near Shannon limit correction capability and therefore maximize the supported operational specs (endurance, retention, disturb resilience, quality of service, etc.) for a given ECC overprovisioning in the NAND. The Sentinel ECC&DSP LDPC decoder is based on a proprietary multi-gear architecture, which optimizes the power consumption under variable BER observed across memory pages, operational conditions and memory lifetime. It supports multiple read

ECC/DSP White Paper at 5.

53. The Accused Products determine a bit error rate (BER) associated with a region comprising a fixed number of two or more flash memory pages or integer fractions thereof, as demonstrated in the images below:

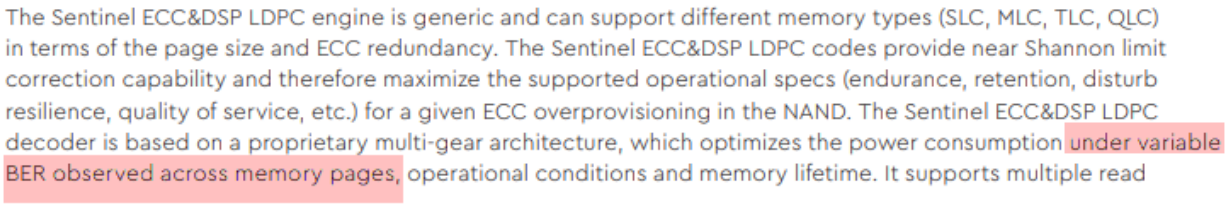
The Sentinel ECC&DSP LDPC engine is generic and can support different memory types (SLC, MLC, TLC, QLC) in terms of the page size and ECC redundancy. The Sentinel ECC&DSP LDPC codes provide near Shannon limit correction capability and therefore maximize the supported operational specs (endurance, retention, disturb resilience, quality of service, etc.) for a given ECC overprovisioning in the NAND. The Sentinel ECC&DSP LDPC decoder is based on a proprietary multi-gear architecture, which optimizes the power consumption under variable BER observed across memory pages, operational conditions and memory lifetime. It supports multiple read

Idan Alrod et al., *The Application of ECC/DSP to Flash Memory* (Mar. 2021), Western Digital, https://documents.westerndigital.com/content/dam/doc-library/en_us/assets/public/western-digital/collateral/white-paper/white-paper-the-application-of-ecc-dsp-to-flash-memory.pdf [hereinafter ECC/DSP White Paper].

> To power the SSD, Western Digital uses a proprietary Arm-based multi-core eight-channel PCIe 4.0 x4 NVMe SSD controller that leverages a Micron DDR4 DRAM chip to deliver responsive performance. Western Digital references the controller as its WD_BLACK G2. Outfitting the WD Black SN850 with a faster Gen4 PHY is great for performance, but with such fast bandwidth, power draw and heat output were a concern at 28 nm. Thus, WD_BLACK G2 on a newer process node to better control those variables with TSMC's 16nm FinFET technology.

WD Black SN850 Review

Solution Block Diagram

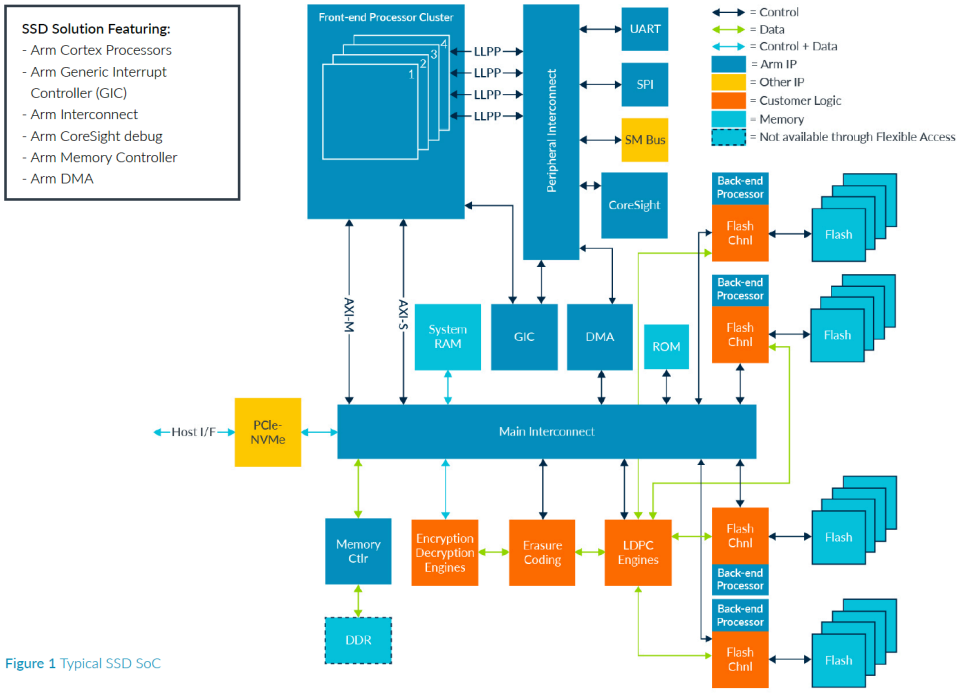

Figure 1 Typical SSD SoC

Front-end processors are usually Cortex-R processors but can sometimes be Cortex-A processors to support more intensive wear-levelling algorithms.

Flash Translation Layer (FTL)

The **front-end processors** are responsible for maintaining the FTL tables that map logical addresses to physical addresses. As part of the FTL, the front-end processors also handle wear-levelling algorithms to determine where the data should be written, which data should move, and which locations require garbage collection.

Wear-levelling is the process of systematically erasing and writing all NAND flash locations to ensure all locations wear out at about the same time. It may require moving unchanging data (cold data) already stored so a block can be erased and rewritten with rapidly changing data (hot data).

Read Requests and Write Requests

Read and write requests arrive from the host over the **PCIe/NVMe** interface. The requests arrive in the form of descriptors that describe the parameters of a transfer. Write requests also carry the data to be written. These descriptors are written to either the **DDR** memory or a system RAM using a **DMA** operation from the host. For write requests, the user data is transferred to the controller DDR using a separate DMA operation.

Arm Storage Solution for SSD Controllers at 4.

> A front-end processor is notified about the arrival of this descriptor. It then parses the descriptor to modify some of the parameters. In particular, the Logical Block Address (LBA) provided by the host must be modified to the physical address in the NAND where the read data is stored. With the correct physical address, the descriptor can be dispatched to the appropriate **flash channel** using a DMA.
>
> For read requests, the flash channel issues a read operation to the appropriate NAND die. When the data is available, it is moved from the NAND die buffer to the flash channel and then to the LDPC engines, as space is available. The **LDPC engines** perform ECC correction and data recovery from the media. If necessary, erasure coding operations are performed to recover data lost due to a block or page failure. Encrypted data is decrypted and stored in controller DDR. Finally, the user data is moved to host memory using a DMA operation to write into the DDR of the host, and then the host is notified that the data transfer completed.
>
> For write requests, the user data in DDR is sent to the SSD data path of encryption, erasure coding, and LDPC encoding. When the encoded user data reaches the flash channel, the flash channel issues an operation to write the data to the NAND buffers. When all the data is in the NAND buffer, a program operation is issued to program the NAND. A completion response is generated by the flash channel that is sent back to the host.

Arm Storage Solution for SSD Controllers at 5.

54. The highlighted text below shows The Accused Products support reading two or more memory pages of a region simultaneously.

> Another important aspect, especially for enterprise grade solutions, is to maintain quality of service (QoS). This means that the delay of providing the host with response for any command falls under tight profile of probabilities it is allowed to exceed certain latency numbers. In order to meet such harsh requirements the ECC engine in gears 1,2,3 is segregated in a manner that it can serve separate requests in separate gears simultaneously. In case a single command requires higher latency gear-3 decoding, a different command can be attended simultaneously with Gear-1 engine. Together with out of order processing and host PCIe™ interface using NVMe protocol the Sentinel ECC&DSP technology is tailored to serve QoS in near optimal manner.

ECC/DSP White Paper at 7.

Now, another restriction on multi-plane operations is being relaxed: the timing of read operations in different planes doesn't need to line up. This makes it possible for one plane to perform multiple reads from SLC pages while another plane is performing a single slower read from TLC or QLC pages. This capability is called Asynchronous Independent (Multi-)Plane Read. The practical effect is that for read operations, a large 4-plane die can now match the performance of four smaller 1-plane dies. This mitigates many of the performance downsides that higher per-die capacity brings to SSDs that only have one or two dies per channel.

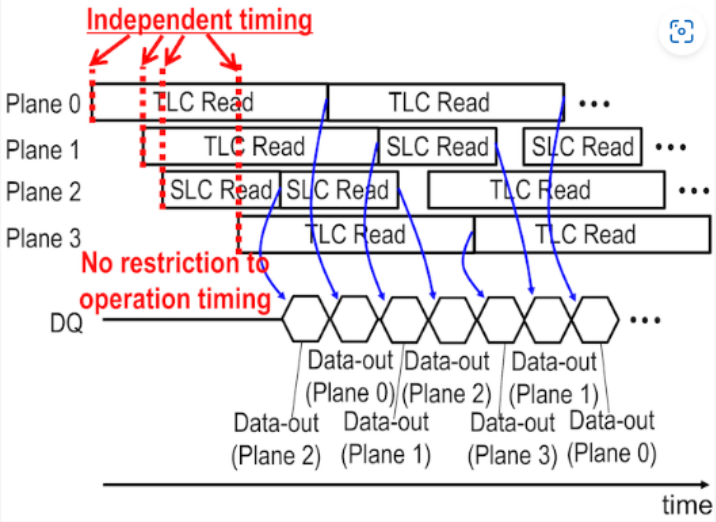

Billy Tallis, *2021 NAND Flash Updates from ISSCC: The Leaning Towers of TLC and QLC*, AnandTech (Feb. 19, 2021), https://www.anandtech.com/show/16491/flash-memory-at-isscc-2021.

55. The Accused Products can store in a region at least data payload and primary and secondary ECC parity symbols corresponding to the data payload of the region, as shown in the highlighted text below:

Error Detection and Correction

One of the key factors to increase flash reliability as well as write endurance is the implementation of an Error Detection and Correction mechanism. The three most popular Error Correction algorithms used with NAND flash technology today are:

- BCH (Bose, Ray-Chaudhuri, Hocquenghem): BCH codes were invented in 1959 by Hocquenghem, and independently in 1960 by Bose and Ray Chaudhuri.
- LDPC (Low Density Parity Codes) invented by Gallagher in 1961.

The maximum number of errors that can be detected and corrected is determined by the algorithm. No matter what algorithm is used within the ECC engine of a NAND flash controller, the method of error detection and correction remains the same:

1. Every time a page of data is written to the flash, data is passed through the controller's ECC Engine to create a unique ECC signature.
2. The data and ECC signature are stored together on the flash; the data in the page area, the ECC signature in the spare area.
3. When reading the data back, both data and stored ECC signature are read into the controller. A new ECC signature is generated, based upon the read-back data.
4. The newly created ECC signature is then compared to the original stored ECC signature. If both signatures are the same, no errors have occurred, and the data will be provided to the host. If the two signatures differ, the data will be corrected by the controller before being provided to the host.

Some flash controllers will write the corrected data back to the flash to optimize reliability, while others will not, since there is no guarantee that the data will not show errors again in the future. For 4KB page flash, typically 8 ECC signatures are created when writing data to the flash; one for each 512 bytes of data, as shown in Figure 10.

Western Digital, *Flash 101 and Flash Management* (Sep. 2023), https://documents.westerndigital.com/content/dam/doc-library/en_us/assets/public/western-digital/collateral/white-paper/white-paper-sandisk-flash101-management.pdf [hereinafter Flash 101 White Paper].

The figure shows a 4KB page with data payload of 4096 data payload and an associated spare area of 128 bytes of ECC bytes containing both primary and secondary ECC bits.

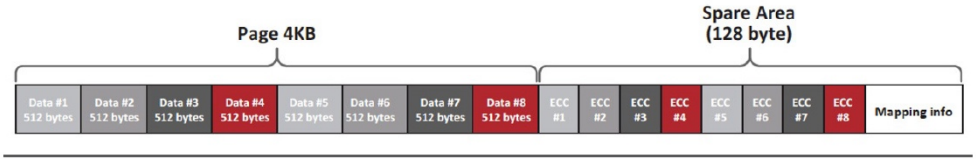

Figure 10: Four ECC signatures for 4KB page

Flash 101 White Paper at 7.

56. The image below shows how The Accused Products compare the determined bit error rate (BER) to one or more predetermined thresholds corresponding to a set of predefined gears comprising at least a first gear and a second gear, wherein the predefined gears correspond to different predefined ECC schemes.

In this example, three decoding gears are used. The color-coding is used to show the BER regions where each gear is likely to succeed the decoding.

Gear 1 (the green region in Figure 5 below) is based on the bit-flipping decoder mentioned above. Said Gear-1 is characterized by high-energy efficiency (low J/GB/sec). However, as the BER goes up, it fails at some point to decode with high probability. At that point, the higher resolution decoding gears kicks in.

The blue curve in Figure 5 indicates the throughput of the system. Naturally, once NAND approaches the BER point where the 2nd gear kicks in, the throughput declines rapidly assuming constant power consumption is maintained.

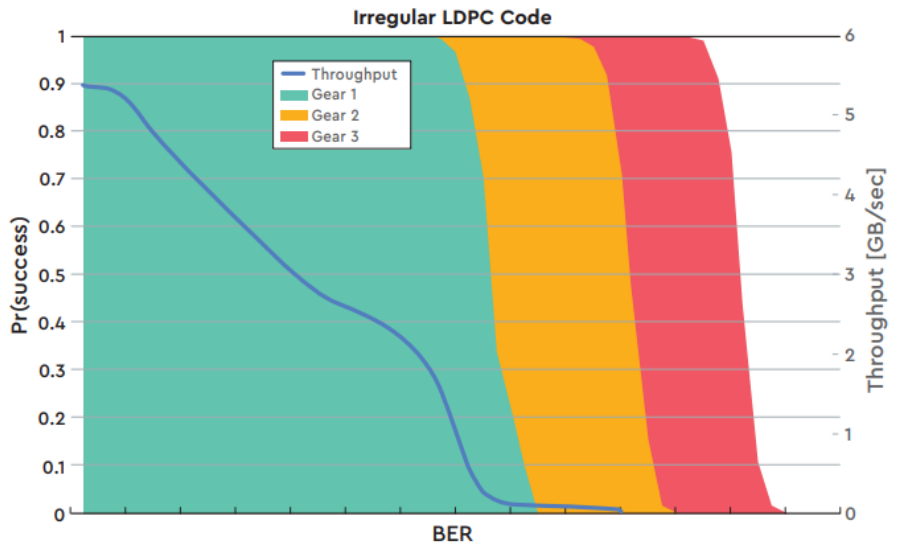

Figure 5: Example of various decoding gears correction capability and overall performance

When the BER is high, there is no point in spending decoder time trying to decode with lower gears, as this would only degrade performance by adding unnecessary latency to the decoding sequence. Consequently, the LDPC engine estimates the BER of the noisy page as part of its initialization process by counting the number of unsatisfied parity checks. This forms the basis on which it automatically chooses the appropriate decoding gear.

ECC/DSP White Paper at 6.

57. The Accused Products include a first gear with a different data payload size and

correction capability than a second gear. For example, Gear 1 corresponds to a high throughput and a ECC level while Gear 2 corresponds to a lower throughput and an increased ECC level (i.e., different payload size and correction capability), as shown in the image below:

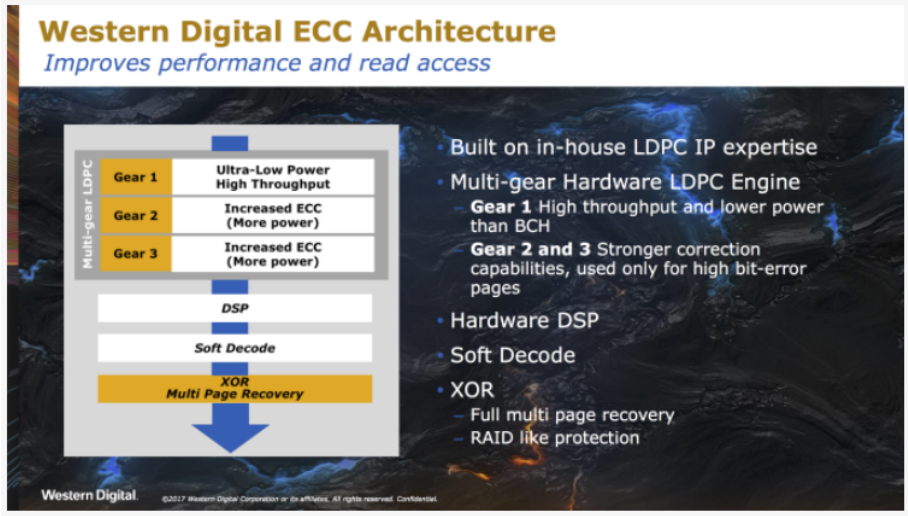

Ganesh T S & Billy Tallis, The Western Digital NVMe Architecture - NAND & Controller, *The Western Digital WD Black 3D NAND SSD Review: EVO Meets Its Match*, AnandTech (Apr. 5, 2018), https://www.anandtech.com/show/12543/the-western-digital-wd-black-3d-nand-ssd-review/2 [hereinafter WD Black 3D Review].

See also, U.S. Patent No. 11,527,300, *Level Dependent Error Correction Code Protection in Multi-Level Non-Volatile Memory*, by Yang et al., assigned to Western Digital [hereinafter Yang], describing the adjustment of payload size and ECC parity size (correction capability) based on BER, corresponding to the criteria for selecting the "gear".

> FIG. 12 illustrates an ECC codeword 1200 and a dynamic or variable size of Components of the ECC codeword in accordance with one embodiment. The ECC codeword 1200 may comprise a payload and a parity section as introduced in FIG. 6A through FIG. 6C. The payload size 1202 and parity size 1204 of the ECC codeword 1200 may vary as indicated by the dotted lines. "Payload size" refers to a size measured in data storage units for a payload of an ECC codeword. In one embodiment, a data storage unit is equal in size to a data sector. "Parity size" refers to a size measured in bits or bytes, or some other data storage measurement unit, for a parity section. The unit of measure for a parity size may depend on the ECC method being used to encode and decode the ECC codewords.

Yang at 33:12-39.

> Determining the coding rate may comprise increasing the payload size 1202 in response to an attribute indicating a greater data integrity for data stored on the multi-level page that is assigned to the ECC codeword 1200. In particular, the payload size 1202 may be increased, and the parity size 1204 decreased, based on the data integrity of the ECC codeword 1200 being higher in relation to other multi-level pages that will be stored on the same multi-level storage cells. Alternatively, or in addition, the payload size 1202 may be decreased, and the parity size 1204 increased, based on the data integrity of the ECC codeword 1200 being lower in relation to other multi-level pages that will be stored on the same multi-level storage cells.

Yang at 34:4-16.

58. The image below shows how The Accused Products allocate memory space for the storage of data payload within the region varies between the first gear and the second gear to accommodate a varying number of parity symbols between the first gear and the second gear. The memory space allocated for the storage of data payload is reduced as the flash controller transitions into a higher gear. The reduction of storage for data payload is caused by the increased of parity bits required by a more complex ECC used in higher gears.

The blue curve in Figure 5 indicates the throughput of the system. Naturally, once NAND approaches the BER point where the 2nd gear kicks in, the throughput declines rapidly assuming constant power consumption is maintained.

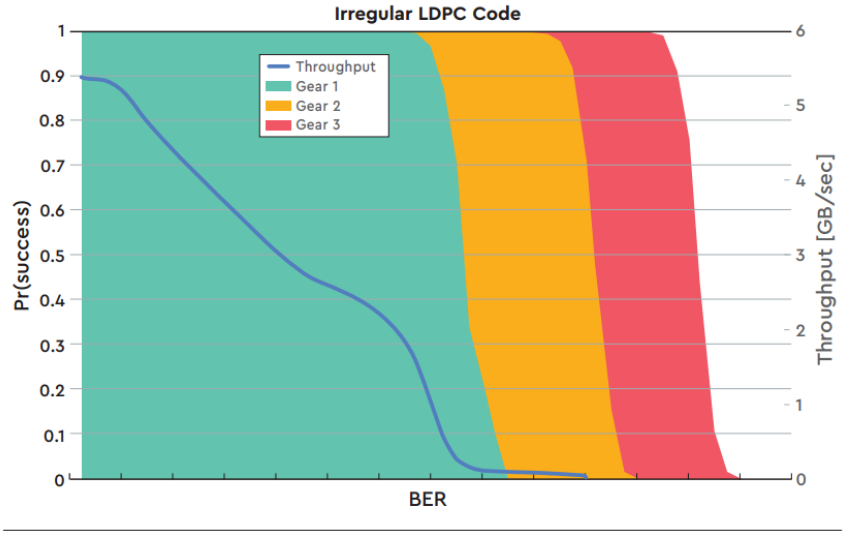

Figure 5: Example of various decoding gears correction capability and overall performance

When the BER is high, there is no point in spending decoder time trying to decode with lower gears, as this would only degrade performance by adding unnecessary latency to the decoding sequence. Consequently, the LDPC engine estimates the BER of the noisy page as part of its initialization process by counting the number of unsatisfied parity checks. This forms the basis on which it automatically chooses the appropriate decoding gear.

ECC/DSP White Paper at 6.

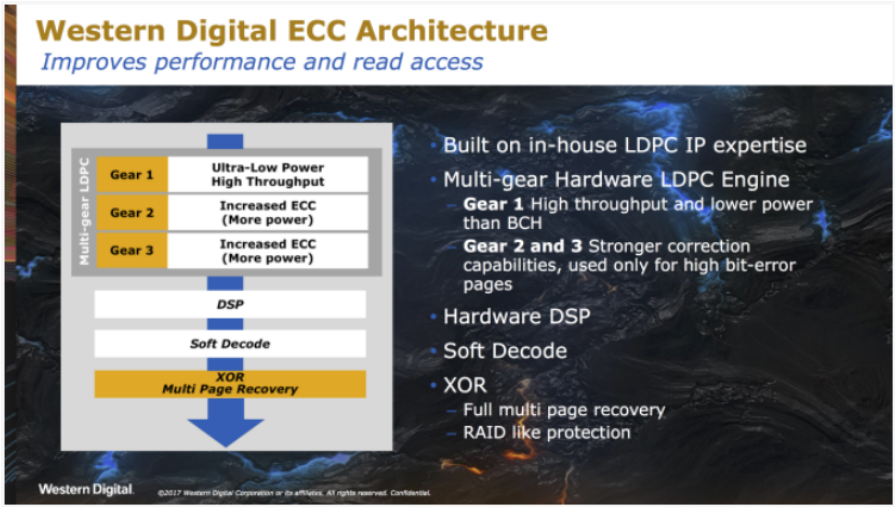

WD Black 3D Review

See also, ECC/DSP White Paper at 11 ("This feature is set to accommodate full usage of the available NAND cells in the physical page while enabling flexibility in code rate for numerous applications / scenarios.")

See also Yang, assigned to Western Digital.

> In certain embodiments, the payload size 1202 and parity size 1204 may be adjusted based on a data integrity attribute for a particular non-volatile storage media, storage cell, set of storage cells or the like. "Data integrity" refers to an attribute or measure of data, or a data sample, indicating whether the data is accurate, not erroneous, and unchanged from a prior transmission or recording of the data. In certain embodiments, data integrity is an objective characteristic. In other embodiments, data integrity may be expressed in relation to a spectrum in which one end represents no, or very low data integrity and the opposite end represents perfect, or very high data integrity.

Yang at 33:45-56.

> The receiver 1304 may coordinate with the address allocator 1306 to determine where the write data 1312 will be stored. The address allocator 1306 may determine a multilevel page to store a set of data blocks associated with a set of write commands 1310. "Address allocator" refers to any circuit, sub-circuit, electronic component, hardware, software, firmware, module, logic, device, or apparatus configured, programmed, designed, arranged, or engineered to determine, assign and/or allocate a physical block address for a particular logical block address.
>
> The packetizer 1308 may coordinate with the receiver 1304 and address allocator 1306 to prepare the write data 1312 for storage on the non-volatile storage media.

> "Packetizer" refers to any hardware, software, firmware, circuit, component, module, logic, device, or apparatus configured, programmed, designed, arranged, or engineered to organize a set of source data into one or more data packets. In one embodiment, the source data may comprise user data for one or more storage operations. The packetizer may be configured to include a header, footer and/or redundancy data in each data packet. The packetizer may be configured to include padding data or filler data to combine with a remainder of the source data that does not completely fill a data packet.
>
> The packetizer 1308 may combine the write data 1312 for the set of data blocks into a payload for an ECC codeword. The packetizer 1308 may coordinate with the address allocator 1306 to determine which multi-level page a particular ECC codeword is assigned to be stored on. Based on a determined multi-level page for the particular ECC codeword, the packetizer 1308 may change a payload size for the payload of the particular ECC codeword in response to a reliability attribute of the determined multi-level page satisfying a threshold.

Yang at 34:62-35:28.

59. The Accused Products select a gear from the set for the region based at least partly on the comparisons to the one or more predetermined thresholds, as shown in the image and highlighted text below:

In this example, three decoding gears are used. The color-coding is used to show the BER regions where each gear is likely to succeed the decoding.

Gear 1 (the green region in Figure 5 below) is based on the bit-flipping decoder mentioned above. Said Gear-1 is characterized by high-energy efficiency (low J/GB/sec). However, as the BER goes up, it fails at some point to decode with high probability. At that point, the higher resolution decoding gears kicks in.

The blue curve in Figure 5 indicates the throughput of the system. Naturally, once NAND approaches the BER point where the 2nd gear kicks in, the throughput declines rapidly assuming constant power consumption is maintained.

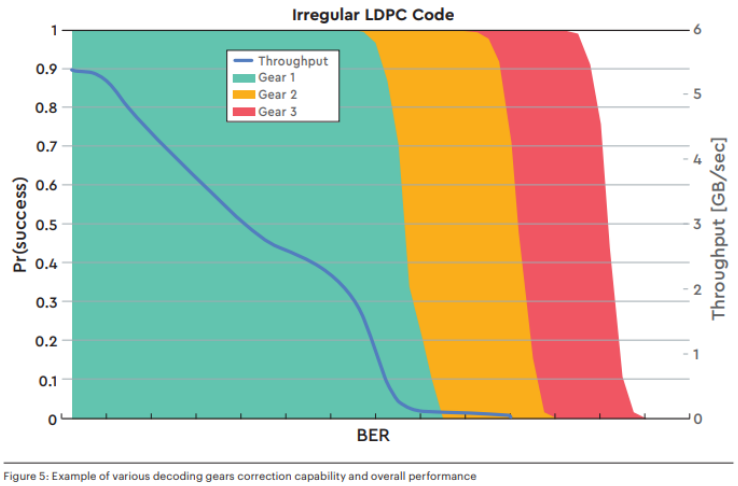

Figure 5: Example of various decoding gears correction capability and overall performance

When the BER is high, there is no point in spending decoder time trying to decode with lower gears, as this would only degrade performance by adding unnecessary latency to the decoding sequence. Consequently, the LDPC engine estimates the BER of the noisy page as part of its initialization process by counting the number of unsatisfied parity checks. This forms the basis on which it automatically chooses the appropriate decoding gear.

ECC/DSP White Paper at 6.

60. The Accused Products determine, compare and select using an integrated circuit.

> Furthermore, the parallelism of each decoding gear is dimensioned according to its usage probability (given the memory BER distribution) and the overall required decoding throughput. Thus, the number of costly high-resolution processing units instantiated for the full resolution BP decoder, which is rarely used ("safety net"), could be much lower than the number of simple BF processing units. This approach significantly reduces the ASIC footprint with a negligible impact on overall sustained decoding throughput.

ECC/DSP White Paper at 7.

61. Defendants have infringed, and continue to infringe, the claims of the ’085 Patent in the United States, by making, using, offering for sale, selling, and/or importing the Accused Products in violation of 35 U.S.C. § 271(a). For example, Defendants directly infringe at least claim 1 of the ’085 Patent when its SSD products are installed and operated by its employees in a computer system, such as for gaming, personal computing, or data centers, and when Defendants’ employees use and test its SSD products.

62. On January 23, 2024, Plaintiff sent a notice letter to Defendants, notifying them that they were infringing the Asserted Patents. Defendants’ infringement of the ’085 Patent has been willful and intentional under the standard announced in *Halo Elecs., Inc. v. Pulse Elecs., Inc.*, 579 U.S. 93, 136 S. Ct. 1923, 195 L. Ed 2d 278 (2016). Defendants have willfully infringed the ’085 Patent by refusing to take a license and continuing to make, use, test, sell, license, and/or offer for sale/license the ’085 Patent Accused Products. Instead of taking a license, Defendants have opted to make the business decision to “efficiently infringe” the ’085 Patent. In doing so, Defendants willfully infringe the ’085 Patent.

63. Defendants also have infringed, and continue to infringe the claims of the ’085 Patent by actively inducing others to use the Accused Products. Defendants’ users, customers, agents or other third parties who use the Accused Products in accordance with Defendants’ instructions infringe the claims of the ’085 Patent, in violation of 35 U.S.C. § 271(a). Defendants induce its customers to use its SSD products for the benefits of higher performance, reduced size

and power, and increased reliability compared to other SSD products. Defendants are thereby liable for infringement of the '085 Patent under 35 U.S.C. § 271(b).

64. Defendants' users, customers, agents or other third parties who use the Accused Products in accordance with Defendants' instructions infringe the claims of the '085 Patent, in violation of 35 U.S.C. § 271(b). Defendants intentionally instruct its customers to infringe through support information, demonstrations, brochures, videos, and user guides, such as those located at:

https://www.westerndigital.com/support;
https://support.wdc.com/contact.aspx?lang=en;
https://www.westerndigital.com/support/category-selection;
https://www.youtube.com/channel/UC_M0BlntDVSblWg6ggUa7UQ;
https://www.youtube.com/watch?v=L0zZtwRFj0E;
https://www.youtube.com/westerndigital; and
https://www.youtube.com/c/westerndigitalcorporation.

65. Defendants are liable as contributory infringers of the '085 Patent under 35 U.S.C. 271(c) by having offered to sell, sold and imported and continuing to offer to sell, selling, and importing into the United States its SSD products, to be especially made or adapted for use in an infringement of the '085 Patent. Defendants' SSD products are key components in gaming consoles and gaming PCs, computers, everyday PCs, Network Attached Storage (NAS) devices, servers and data centers. These SSD products are material components for use in practicing the '085 Patent and are specifically made and are not a staple article of commerce suitable for substantial noninfringing use. Defendants supplied these components with knowledge of the '085 Patent and with knowledge that these components constitute material parts of the claimed inventions of the '085 Patent.

66. As a result of Defendants' infringement of the '085 Patent, Plaintiff has suffered monetary damages and is entitled to no less than a reasonable royalty for Defendants' use of the

claimed inventions of the ’085 Patent, together with interest and costs as determined by the Court. Plaintiff will continue to suffer damages in the future.

67. Defendants are on notice of its infringement by no later than the filing and service of this Complaint. By the time of trial, Defendants will have known and intended (since receiving such notice) that their continued actions would actively induce the infringement of the ’085 Patent.

68. Defendants’ acts of direct and indirect infringement have caused and continue to cause damage to Plaintiff. Plaintiff is entitled to damages in accordance with 35 U.S.C. §§ 271 and 281 sustained as a result of Defendants’ wrongful acts in an amount to be proven at trial.

COUNT III

(INFRINGEMENT OF U.S. PATENT NO. 8,601,346)

69. Plaintiff incorporates by reference and re-alleges the foregoing paragraphs as fully set forth herein.

70. Defendants have directly infringed, and continue to directly infringe, one or more claims of the ’346 Patent, including at least claim 1 of the ’346 Patent, literally and/or under the doctrine of equivalents, by or through making, using, offering for sale, selling within the United States and/or importing into the United States its SSD products.

71. Claim 1 of the ’346 Patent reads as follows:

> 1. A nonvolatile memory controller for performing a data stripe operation on a plurality of data blocks, the nonvolatile memory controller comprising:
>
> a plurality of command processing units, each command processing unit of the plurality of command processing units configured to receive a command of a plurality of commands for performing the data stripe operation, the plurality of commands including a plurality of data update commands and a parity write command, each command processing unit of the plurality of command processing units receiving a data update command of the plurality of data update commands configured to request a data block of the plurality of data blocks based on the data update command, receive the data block in response to the request, and write the data block to a nonvolatile memory device; and

a parity calculator coupled to the plurality of command processing units, the parity calculator further comprising a context memory including a page frame, the parity calculator configured to receive the plurality of data blocks as a sequence of data blocks, to generate a parity block by storing a first data block of the sequence of data blocks into the page frame and updating the data block stored in the page frame with each data block following the first data block in the sequence of data blocks, without storing each data block in a data buffer, the command processing unit receiving the parity write command configured to write the parity block to a nonvolatile memory based on the parity write command.

72. Defendants make, use, sell, offer for sale and import SSD products that include a nonvolatile memory controller for performing a data stripe operation on a plurality of data blocks. For example, WD Black-branded SSDs, which are used for gaming consoles and gaming PCs, Blue- and Green-branded SSDs, which are used for everyday PCs, Red-branded SSDs, which are used for Network Attached Storage (NAS), and Gold- and Ultrastar SSDs, which are used for data centers, all include a nonvolatile memory controller for performing a data stripe operation on a plurality of data blocks, as shown below:

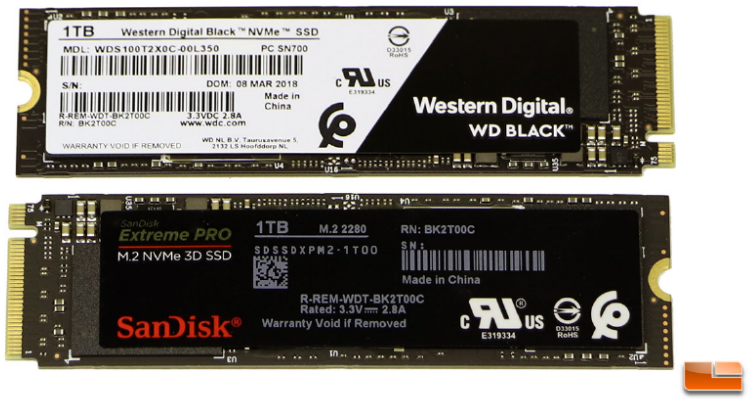

Western Digital's new controller architecture means that the company has in-house development from the controller to the NAND. This gives WD tremendous design flexibility and they believe this new PCIe Gen 3 x4 controller is scalable and future-ready (think QLC NAND). Since WD designed everything they were able to design for low power and low latency as well as general performance. WD's new controller has three cores, is built on the 28nm process, features nCache 3.0 (marketing term for SLC Cache) for enhanced burst speeds, multi-gear LDPC (low-density parity-check) and hardware ECC to provide full SRAM and DRAM bit correction.

Nathan Kirsch, *WD Black NVMe 3D and SanDisk Extreme PRO NVMe 3D 1TB SSD Review*, Legit Reviews (Apr. 5, 2018), https://www.legitreviews.com/wd-black-nvme-3d-sandisk-extreme-pro-nvme-3d-1tb-ssd-review_204268 [hereinafter WD Black and SanDisk Review].

73. The images and quotes below show how The Accused Products include command processing units, such as ARM processor cores. The processing units receive commands, including a command to perform a data stripe operation. For example:

> The new controller has a tri-core architecture (probably using Arm Cortex-cores Fabricated in a 28nm process. It is designed to be scalable – the current controller can interface with the host using a PCIe 3.0 x4 link, or an x2 link as in the Western Digital SN520. The architecture of the controller also allows future products using variants to come to market faster and with newer features. It also allows Western Digital to segment their NVMe product stack. The controller in the Western Digital Black 3D NAND SSD is optimized for client workloads including PC gaming and high-performance commercial applications. Western Digital expects this new controller architecture to last at least until NVMe SSDs move beyond PCIe 3 x4 interfaces.

WD Black 3D Review.

> WD's new controller has three cores, is built on the 28nm process, …

WD Black and SanDisk Review

> To power the SSD, WD uses a proprietary Arm-based multi-core eight-channel PCIe 4.0 x4 NVMe SSD controller that leverages a Micron DDR4 DRAM chip to deliver responsive performance. Western Digital references the controller as its WD_BLACK G2. Outfitting the Western Digital Black SN850 with a faster Gen4 PHY is great for performance, but with such fast bandwidth, power draw and heat output were a concern at 28nm. Thus, like the controllers from competing manufacturers, Western Digital opted to build the WD_BLACK G2 on a newer process node to better control those variables with TSMC's 16nm FinFET technology.

Sean Webster, *WD Black SN850 M.2 NVMe SSD Review: Top-Tier Storage for Gamers and Pros*

(Updated), Tom's Hardware (Apr. 1, 2021), https://www.tomshardware.com/reviews/wd-black-sn850-m-2-nvme-ssd-review [hereinafter WD Black SN850 Review].

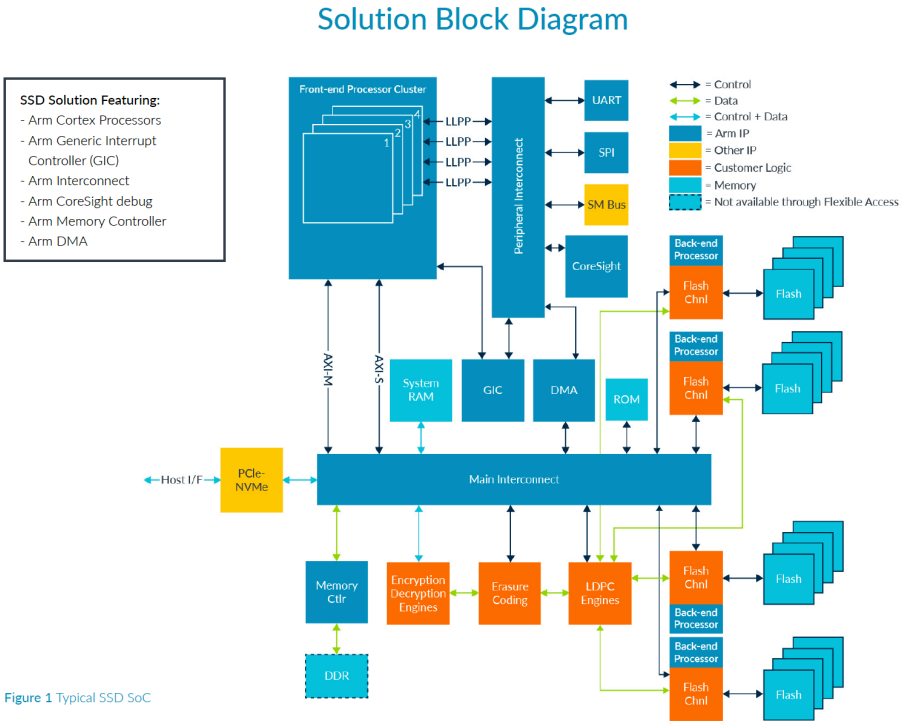

Figure 1 Typical SSD SoC

Front-end processors are usually Cortex-R processors but can sometimes be Cortex-A processors to support more intensive wear-levelling algorithms.

Flash Translation Layer (FTL)

The **front-end processors** are responsible for maintaining the FTL tables that map logical addresses to physical addresses. As part of the FTL, the front-end processors also handle wear-levelling algorithms to determine where the data should be written, which data should move, and which locations require garbage collection.

Wear-levelling is the process of systematically erasing and writing all NAND flash locations to ensure all locations wear out at about the same time. It may require moving unchanging data (cold data) already stored so a block can be erased and rewritten with rapidly changing data (hot data).

Read Requests and Write Requests

Read and write requests arrive from the host over the **PCIe/NVMe** interface. The requests arrive in the form of descriptors that describe the parameters of a transfer. Write requests also carry the data to be written. These descriptors are written to either the **DDR** memory or a system RAM using a **DMA** operation from the host. For write requests, the user data is transferred to the controller DDR using a separate DMA operation.

Arm Ltd., *Arm Storage Solution for SSD Controllers* at 4 (Sept. 2020), https://www.arm.com/-/media/global/solutions/storage/arm-storage-solution-for-ssd-solutions-brief.pdf?rev=3530e7536aae437aa2d7acf1704fe25b&revision=3530e753-6aae-437a-a2d7-acf1704fe25b [hereinafter Arm Storage Solution for SSD Controllers].

A front-end processor is notified about the arrival of this descriptor. It then parses the descriptor to modify some of the parameters. In particular, the Logical Block Address (LBA) provided by the host must be modified to the physical address in the NAND where the read data is stored. With the correct physical address, the descriptor can be dispatched to the appropriate **flash channel** using a DMA.

For read requests, the flash channel issues a read operation to the appropriate NAND die. When the data is available, it is moved from the NAND die buffer to the flash channel and then to the LDPC engines, as space is available. The **LDPC engines** perform ECC correction and data recovery from the media. If necessary, erasure coding operations are performed to recover data lost due to a block or page failure. Encrypted data is decrypted and stored in controller DDR. Finally, the user data is moved to host memory using a DMA operation to write into the DDR of the host, and then the host is notified that the data transfer completed.

For write requests, the user data in DDR is sent to the SSD data path of encryption, erasure coding, and LDPC encoding. When the encoded user data reaches the flash channel, the flash channel issues an operation to write the data to the NAND buffers. When all the data is in the NAND buffer, a program operation is issued to program the NAND. A completion response is generated by the flash channel that is sent back to the host.

Arm Storage Solution for SSD Controllers at 5.

Defendants' nonvolatile memory controllers include RAID scheme support, which inherently includes stripe operations.

The Sentinel ECC&DSP error correction is based on state-of-the-art Low Density Parity Check (LDPC) coding and provides a full suite of NAND DSP (Digital Signal Processing) services, including data randomization or shaping, NAND health metering via Bit Error Rate (BER) estimation, ECC-based read thresholds calibration, and NAND defect protection and recovery via XOR based RAID scheme support.

Idan Alrod et al., *The Application of ECC/DSP to Flash Memory* (Mar. 2021), Western Digital, https://documents.westerndigital.com/content/dam/doc-library/en_us/assets/public/western-digital/collateral/white-paper/white-paper-the-application-of-ecc-dsp-to-flash-memory.pdf [hereinafter ECC/DSP White Paper] at 5.

Joint LDPC and RAID Decoding for Enterprise SSD

Certain flash applications, such as enterprise SSD, require an exceptionally high level of data reliability. In these applications, the probability of losing data – no matter what caused it – must be exceptionally low. More specifically, the system must be able to cope with flash failure modes that do not only include random errors, but failures of entire pages. These might include word line failures of various sorts, block failures, or even die-level faults.

One existing strategy is to use RAID parity. This method has many variants depending on reliability requirements (how many failed sectors can the system recover from), locality of recovery (how many read and transfer operations required for a single recovery), and other system tradeoffs. The simplest of such schemes is the plain-vanilla variant, where recovery from a single sector failure is possible. This requires one parity sector, which holds the RAID parity of all other sectors. Upon failure of one of the sectors, reconstruction is possible by performing RAID recovery on all the non-failing sectors. Of course, this simple scheme is equally effective if the failing sector is due to a large number of random errors, which is not correctable by the ECC. The following two are simple observations on this scheme:

ECC/DSP White Paper at 11.

Appendix – Case Study [FMS 2016, Patent Application US20170255518A1]

Let's assume a standard RAID stripe of length 32. This means that 31 sectors are protected by an additional parity sector. The over-provisioning used by this RAID scheme is around 3%. Next, let's assume an LDPC-based ECC solution that has approximately 10% over-provisioning.

If the LDPC and RAID layers were to be combined, one could observe that the result is a code with extra 3% overprovisioning and a code word size that is x32 larger than the original. A Tanner graph representing a unified LDPC-RAID code is shown in Figure 13 below. It shows a Tanner graph of a combined LDPC-RAID scheme. The smaller graphs G_0,G_1,...,G_30 represent the information sectors and G_31 is the parity sector.

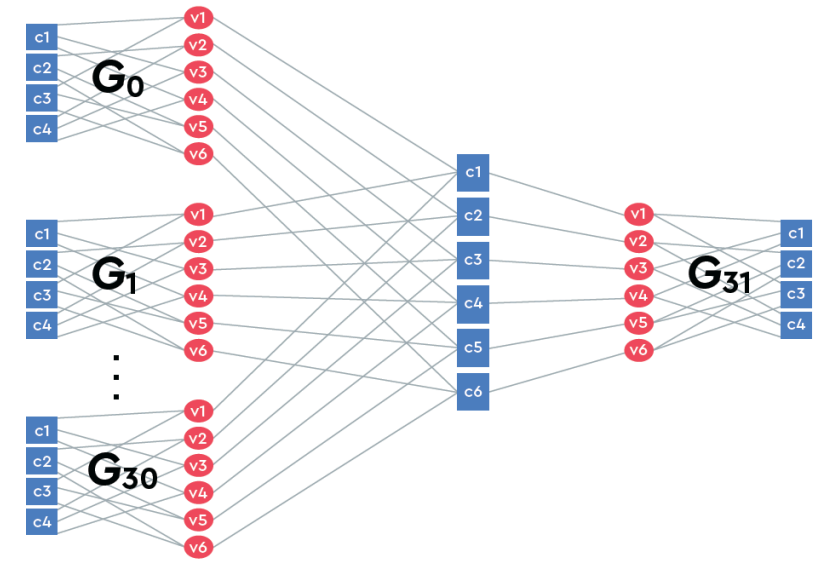

Figure 13: Tanner graph of a combined LDPC-RAID scheme

ECC/DSP White Paper at 12.

What are RAID 50 and 60? I see you have a product that supports them. –

RAID 50 and 60 are very similar to RAID 10, as they are all "nested" RAIDs. The first number in the type — 5 and 6, respectively — represents what subset they are built with, either RAID 5 arrays or RAID 6 arrays. The second number represents the RAID type the first one is "nested" into; in this case, it's RAID 0.

Let's dive a bit deeper.

So, with RAID 50:

- Two RAID 5 subsets stripe across each other in RAID 0.
- With eight drives, four drives act as RAID 5 with their own parity. Another four do the same with data split between two sets.
- Any one of the four drives in each set could fail. That's one to two drives — two if they're not within the same subset.

So, with RAID 60:

- Two RAID 6 subsets stripe across each other in RAID 0.
- With eight drives, any one of the four drives in each set could fail. That's two to four drives — four if two are from one subset and two are from the other subset.

RAID 50 vs. RAID 60:

- RAID 50 will have faster performance than RAID 60.
- RAID 60, like RAID 6, has parity for all data written twice. That's why you "lose" two disks of capacity per RAID 6 subset.
- With RAID 50, it's one disk per subset.

Western Digital, *What is RAID Storage?*, https://www.westerndigital.com/solutions/raid [hereinafter RAID Storage Guide].

74. The Accused Products include a plurality of data update commands, which write data to a nonvolatile controller to update the contents of nonvolatile memory, and a parity write command, as shown below:

> **6.15 Write command**
>
> The Write command writes data and metadata, if applicable, to the I/O controller for the logical blocks indicated.

NVMe Express Specification v1.4 at 268.

> The Sentinel ECC&DSP error correction is based on state-of-the-art Low Density Parity Check (LDPC) coding and provides a full suite of NAND DSP (Digital Signal Processing) services, including data randomization or shaping, NAND health metering via Bit Error Rate (BER) estimation, ECC-based read thresholds calibration, and NAND defect protection and recovery via XOR based RAID scheme support.

ECC/DSP White Paper at 5.

> **Joint LDPC and RAID Decoding for Enterprise SSD**
>
> Certain flash applications, such as enterprise SSD, require an exceptionally high level of data reliability. In these applications, the probability of losing data – no matter what caused it – must be exceptionally low. More specifically, the system must be able to cope with flash failure modes that do not only include random errors, but failures of entire pages. These might include word line failures of various sorts, block failures, or even die-level faults.
>
> One existing strategy is to use RAID parity. This method has many variants depending on reliability requirements (how many failed sectors can the system recover from), locality of recovery (how many read and transfer operations required for a single recovery), and other system tradeoffs. The simplest of such schemes is the plain-vanilla variant, where recovery from a single sector failure is possible. This requires one parity sector, which holds the RAID parity of all other sectors. Upon failure of one of the sectors, reconstruction is possible by performing RAID recovery on all the non-failing sectors. Of course, this simple scheme is equally effective if the failing sector is due to a large number of random errors, which is not correctable by the ECC. The following two are simple observations on this scheme:

ECC/DSP White Paper at 11.

Appendix – Case Study [FMS 2016, Patent Application US20170255518A1]

Let's assume a standard RAID stripe of length 32. This means that 31 sectors are protected by an additional parity sector. The over-provisioning used by this RAID scheme is around 3%. Next, let's assume an LDPC-based ECC solution that has approximately 10% over-provisioning.

If the LDPC and RAID layers were to be combined, one could observe that the result is a code with extra 3% overprovisioning and a code word size that is x32 larger than the original. A Tanner graph representing a unified LDPC-RAID code is shown in Figure 13 below. It shows a Tanner graph of a combined LDPC-RAID scheme. The smaller graphs G_0,G_1,...,G_30 represent the information sectors and G_31 is the parity sector.

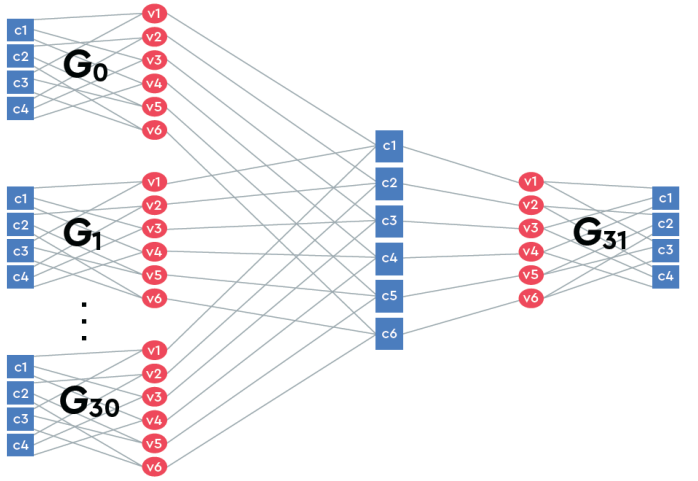

Figure 13: Tanner graph of a combined LDPC-RAID scheme

ECC/DSP White Paper at 12.

75. For The Accused Products, each command processing unit of the plurality of command processing units receiving a data update command of the plurality of data update commands configured to request a data block of the plurality of data blocks based on the data update command, receive the data block in response to the request, and write the data block to a nonvolatile memory device, as shown in the quotes and pictures below:

> **6.15 Write command**
>
> The Write command writes data and metadata, if applicable, to the I/O controller for the logical blocks indicated.

NVMe Express Specification at 268.

Solution Block Diagram

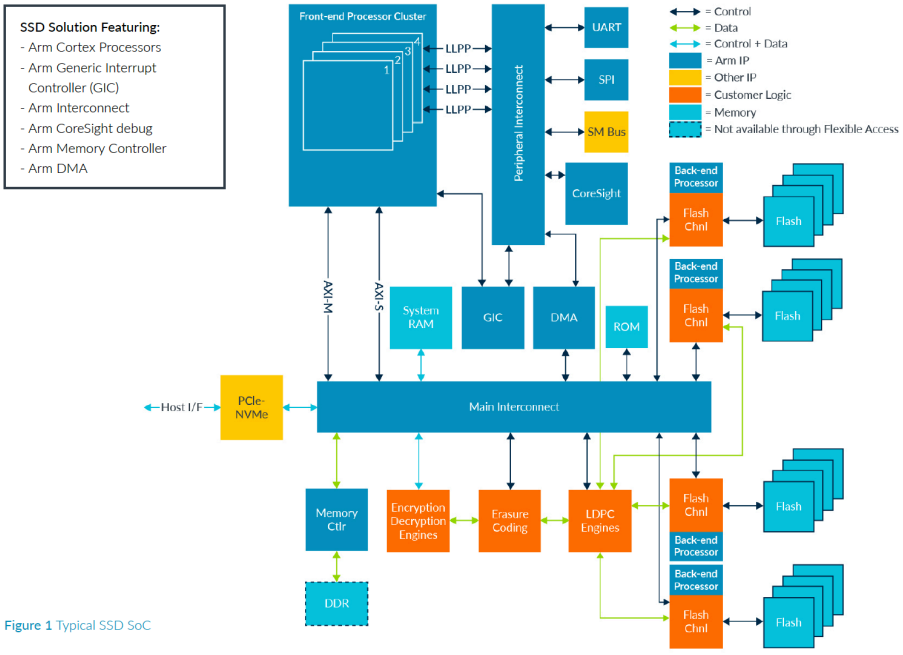

Figure 1 Typical SSD SoC

Front-end processors are usually Cortex-R processors but can sometimes be Cortex-A processors to support more intensive wear-levelling algorithms.

Flash Translation Layer (FTL)

The **front-end processors** are responsible for maintaining the FTL tables that map logical addresses to physical addresses. As part of the FTL, the front-end processors also handle wear-levelling algorithms to determine where the data should be written, which data should move, and which locations require garbage collection.

Wear-levelling is the process of systematically erasing and writing all NAND flash locations to ensure all locations wear out at about the same time. It may require moving unchanging data (cold data) already stored so a block can be erased and rewritten with rapidly changing data (hot data).

Read Requests and Write Requests

Read and write requests arrive from the host over the **PCIe/NVMe** interface. The requests arrive in the form of descriptors that describe the parameters of a transfer. Write requests also carry the data to be written. These descriptors are written to either the **DDR** memory or a system RAM using a **DMA** operation from the host. For write requests, the user data is transferred to the controller DDR using a separate DMA operation.

Arm Storage Solution for SSD Controllers at 4.

A front-end processor is notified about the arrival of this descriptor. It then parses the descriptor to modify some of the parameters. In particular, the Logical Block Address (LBA) provided by the host must be modified to the physical address in the NAND where the read data is stored. With the correct physical address, the descriptor can be dispatched to the appropriate **flash channel** using a DMA.

For read requests, the flash channel issues a read operation to the appropriate NAND die. When the data is available, it is moved from the NAND die buffer to the flash channel and then to the LDPC engines, as space is available. The **LDPC engines** perform ECC correction and data recovery from the media. If necessary, erasure coding operations are performed to recover data lost due to a block or page failure. Encrypted data is decrypted and stored in controller DDR. Finally, the user data is moved to host memory using a DMA operation to write into the DDR of the host, and then the host is notified that the data transfer completed.

For write requests, the user data in DDR is sent to the SSD data path of encryption, erasure coding, and LDPC encoding. When the encoded user data reaches the flash channel, the flash channel issues an operation to write the data to the NAND buffers. When all the data is in the NAND buffer, a program operation is issued to program the NAND. A completion response is generated by the flash channel that is sent back to the host.

Arm Storage Solution for SSD Controllers at 5.

76. The Accused Products include a parity calculator coupled to the plurality of command processing units. The Accused Products rely on LDPC (Low Density Parity Calculation) Error Correcting Code (ECC) to quickly correct a limited number of random bit errors within an SSD. A RAID scheme is additionally used to recover more slowly when ECC does not allow data to be read reliably (for example, NAND defect protection or a sector or drive failure) by storing redundant data for a stripe. As shown in the text below, a parity calculator is used to create one or more redundant data blocks for RAID storage.

The Sentinel ECC&DSP error correction is based on state-of-the-art Low Density Parity Check (LDPC) coding and provides a full suite of NAND DSP (Digital Signal Processing) services, including data randomization or shaping, NAND health metering via Bit Error Rate (BER) estimation, ECC-based read thresholds calibration, and NAND defect protection and recovery via XOR based RAID scheme support.

ECC/DSP White Paper at 5.

Joint LDPC and RAID Decoding for Enterprise SSD

Certain flash applications, such as enterprise SSD, require an exceptionally high level of data reliability. In these applications, the probability of losing data – no matter what caused it – must be exceptionally low. More specifically, the system must be able to cope with flash failure modes that do not only include random errors, but failures of entire pages. These might include word line failures of various sorts, block failures, or even die-level faults.

One existing strategy is to use RAID parity. This method has many variants depending on reliability requirements (how many failed sectors can the system recover from), locality of recovery (how many read and transfer operations required for a single recovery), and other system tradeoffs. The simplest of such schemes is the plain-vanilla variant, where recovery from a single sector failure is possible. This requires one parity sector, which holds the RAID parity of all other sectors. Upon failure of one of the sectors, reconstruction is possible by performing RAID recovery on all the non-failing sectors. Of course, this simple scheme is equally effective if the failing sector is due to a large number of random errors, which is not correctable by the ECC. The following two are simple observations on this scheme:

ECC/DSP White Paper at 11.

Appendix – Case Study [FMS 2016, Patent Application US20170255518A1]

Let's assume a standard RAID stripe of length 32. This means that 31 sectors are protected by an additional parity sector. The over-provisioning used by this RAID scheme is around 3%. Next, let's assume an LDPC-based ECC solution that has approximately 10% over-provisioning.

If the LDPC and RAID layers were to be combined, one could observe that the result is a code with extra 3% overprovisioning and a code word size that is x32 larger than the original. A Tanner graph representing a unified LDPC-RAID code is shown in Figure 13 below. It shows a Tanner graph of a combined LDPC-RAID scheme. The smaller graphs G_0,G_1,...,G_30 represent the information sectors and G_31 is the parity sector.

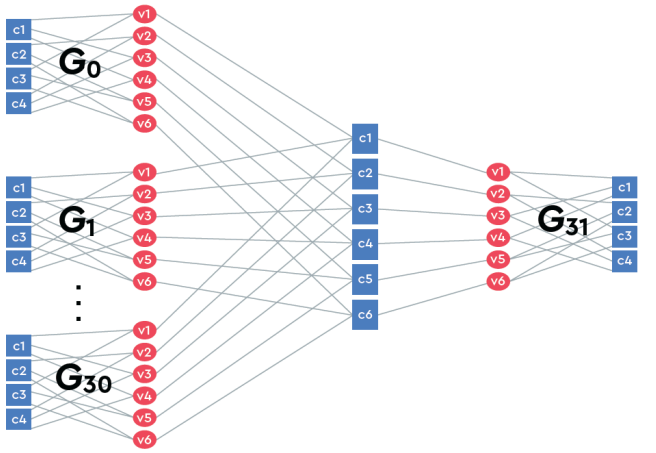

Figure 13: Tanner graph of a combined LDPC-RAID scheme

ECC/DSP White Paper at 12.

77. As Western Digital (WD) describes in a patent issued from the patent application cited in the ECC/DSP white paper:

> Error correction coding (ECC) is often used to correct errors that occur in data read from a memory device. Prior to storage, data may be encoded by an ECC encoder to generate redundant information (e.g., "parity bits") that may be stored with the data as an ECC codeword. As more parity bits are used, an error correction capacity of the ECC increases and a number of bits required to store the encoded data also increases. Using a sufficient number of parity bits to provide "worst-case" error correction capability for all data stored in a memory device reduces the storage

> density of the memory device in order to protect against an amount of data corruption that is statistically unlikely to occur before the memory device reaches the end of its useful life.
>
> SSD devices may also incorporate a redundant array of independent dies (RAID)-type storage scheme that may use parity bits to enable data recovery in case of memory defects and device failures, which cannot be recovered by the ECC which is aimed at handling random errors (e.g., due to program disturb, read disturb, charge loss due to data retention, etc.). ECC may not be able to recover the data in case of memory defects or complete failure, which may result in very high error rates that exceed the ECC capability. Hence, additional RAID-type protection may be required for protecting against such memory defects.

U.S. Patent No. 9,940,194 of Achtenberg et al. at 1:32-55, *ECC Decoding Using RAID-Type Parity* available at https://patentimages.storage.googleapis.com/bf/90/8b/dec1bb4115b581/US9940194.pdf [hereinafter Achtenberg].

> The two protection levels, ECC for random errors and RAID for memory defects and failures, may require memory overprovisioning for storing the ECC and RAID parity.

Achtenberg at 2:8-10.

> In a particular implementation, the stripe correction scheme corresponds to a redundant array of independent disks (RAID)-type exclusive-OR (XOR) scheme, and the first correction scheme corresponds to a low density parity check (LDPC) scheme.

Achtenberg at 14:47-51.

See also, WD Black 3D Review.

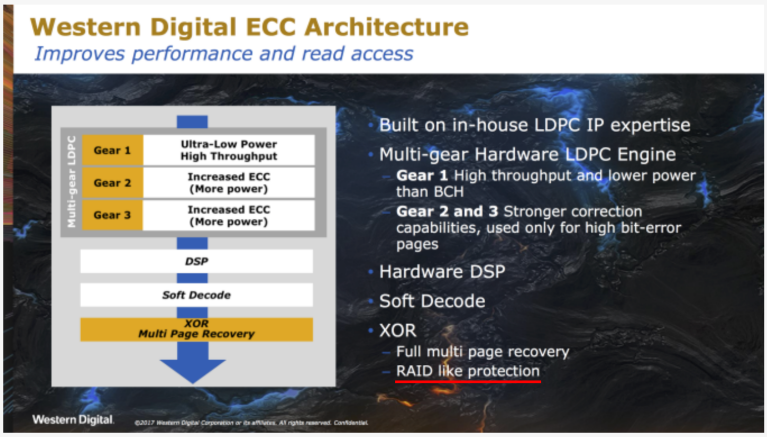

Like any good modern SSD controller, Western Digital's new architecture features multiple layers of error correction. The first three layers are different LDPC-style error correction codes for handling increasing bit error rates, which come at the cost of increasing power consumption and decreased performance. The base level of error correction is an LDPC code that is tuned to offer higher throughput and with lower power requirements than the BCH error correction that nearly all SSDs used before TLC NAND began to take over the market. This lowest layer of error correction is the only one needed during normal operation for most of the drive's lifespan, and this LDPC engine is responsible for less than 10% of the controller's power consumption. The second and third layers of error correction are intended to handle the increasing error rates of a drive that is nearing the end of its write endurance, and these codes are also entirely handled by dedicated hardware on the controller without taking the performance hit of involving the processor cores.

For handling severe data loss that cannot be recovered by the three layers of LDPC, the controller also performs traditional RAID5-like XOR parity. This can handle the failure of defects affecting multiple NAND pages, but consumer SSDs don't include enough excess flash to survive the complete failure of an entire NAND die. Data integrity is also protected by the use of ECC on all of the controller's SRAM and on the external DRAM.

WD Black 3D Review

78. Furthermore, other Defendants' patent filings indicate that its RAID parity calculations are performed by a parity calculator that performs a partial parity calculation on blocks received for writing to a stripe. The following image is from U.S. Patent No. 11,106, 534:

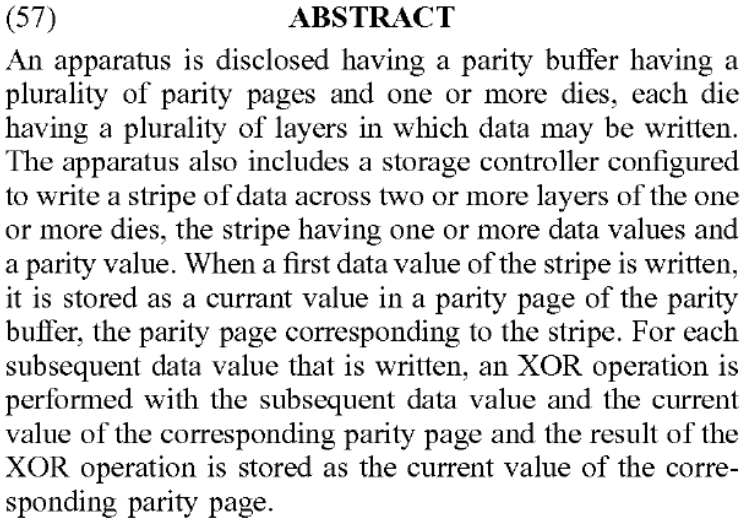

(57) **ABSTRACT**

An apparatus is disclosed having a parity buffer having a plurality of parity pages and one or more dies, each die having a plurality of layers in which data may be written. The apparatus also includes a storage controller configured to write a stripe of data across two or more layers of the one or more dies, the stripe having one or more data values and a parity value. When a first data value of the stripe is written, it is stored as a currant value in a parity page of the parity buffer, the parity page corresponding to the stripe. For each subsequent data value that is written, an XOR operation is performed with the subsequent data value and the current value of the corresponding parity page and the result of the XOR operation is stored as the current value of the corresponding parity page.

U.S. Patent No. 11,106,534, Sun et al., *3-Dimensional Nand Flash Layer Variation Aware SSD RAID*, https://patentimages.storage.googleapis.com/0a/c7/53/8923ad3f4269e8/US11106534.pdf [hereinafter Sun].

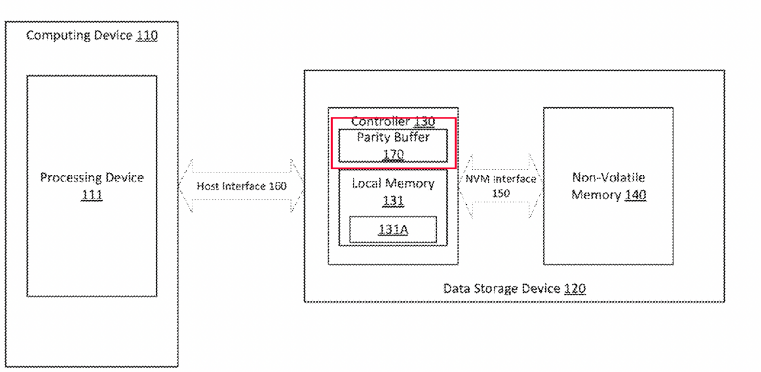

Sun at Fig. 1, annotation added.

79. The Accused Products include a context memory including a page frame. For example, Defendants' patent filings, such as Sun, describes a memory system for storing its partial parity RAID data calculation and associated metadata. This system includes features for indicating updated blocks within a data stripe and for converting logical pages to their physical counterparts.

(57) **ABSTRACT**

An apparatus is disclosed having a parity buffer having a plurality of parity pages and one or more dies, each die having a plurality of layers in which data may be written. The apparatus also includes a storage controller configured to write a stripe of data across two or more layers of the one or more dies, the stripe having one or more data values and a parity value. When a first data value of the stripe is written, it is stored as a currant value in a parity page of the parity buffer, the parity page corresponding to the stripe. For each subsequent data value that is written, an XOR operation is performed with the subsequent data value and the current value of the corresponding parity page and the result of the XOR operation is stored as the current value of the corresponding parity page.

Sun at Abstract.

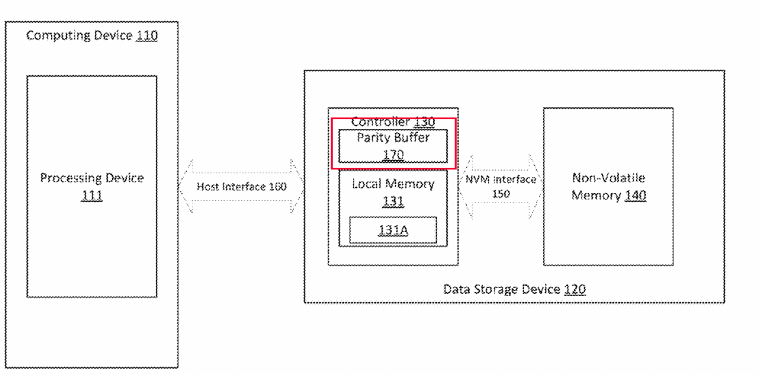

Sun at Fig. 1, annotation added.

> The logical interface may present to the computing device memory a set of logical addresses (e.g., sequential/contiguous addresses) where data may be stored. Internally, the controller 130 may map logical addresses to various physical memory addresses in the non volatile memory arrays and or other memory module(s).

Sun at 5:21-26.

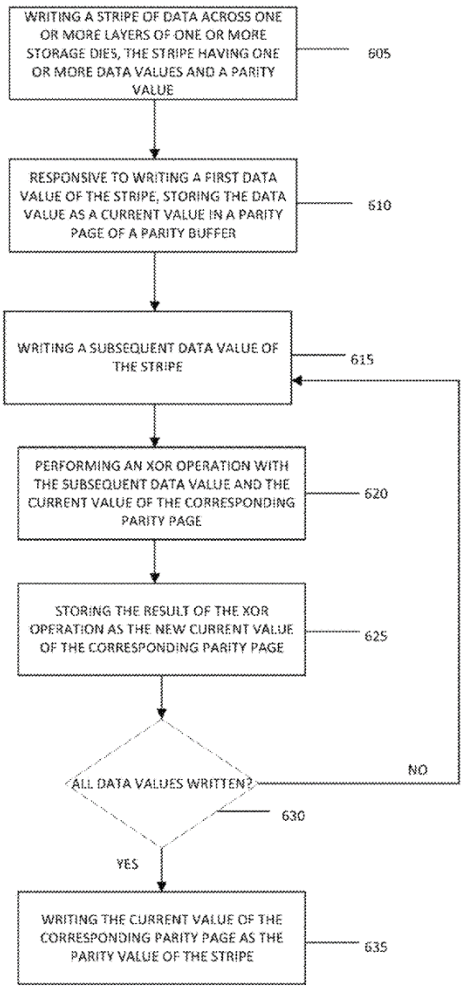

Sun Fig. 6.

> At block 630, the controller may determine whether all the data values are written. If there are still data values to be written, the controller returns to block 615. If all data values are written, then at block 635 the controller writes the current value of the corresponding parity page to the parity value of the stripe.

Sun at 9:52-58.

80. The Accused Products include a parity calculator configured to receive the plurality of data blocks as a sequence of data blocks. For example, U.S. Patent No. 9,940,194, assigned to WD and cited in WD's White Paper, The Application of ECC/DSP to Flash Memory, describes writing a sequence of data blocks for RAID parity calculation, as shown in the quote and image below:

> The controller 130 may be configured to transfer the codewords 160-164 from the Memory 138 for storage into the memory 104 of the memory device 103 to form the data structure 110 in the memory 104. For example, the controller 130 may be configured to sequentially write the codewords 160-164 to consecutively-addressed pages of the memory 104 so that the data structure 110 is aligned in a row-and-column format as depicted in FIG. 1, with the codewords 160-164 forming rows and the stripes 197-199 forming columns in the memory 104.

Achtenberg at 8:31-40.

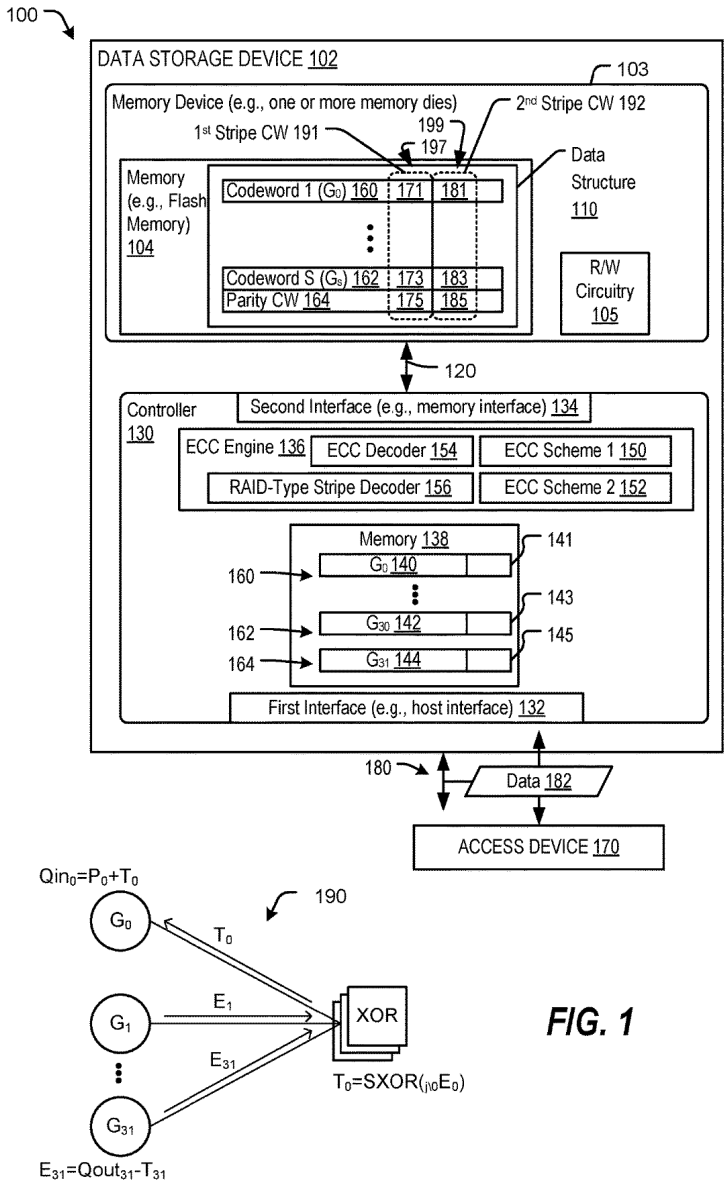

FIG. 1

Achtenberg at Fig. 1.

81. Furthermore, Sun describes a buffer for storing parity information, e.g., parity cache or "partial parity Cache (PPC)", which stores a parity block for a stripe that is updated for each data block written for the stripe.

(57) **ABSTRACT**

An apparatus is disclosed having a parity buffer having a plurality of parity pages and one or more dies, each die having a plurality of layers in which data may be written. The apparatus also includes a storage controller configured to write a stripe of data across two or more layers of the one or more dies, the stripe having one or more data values and a parity value. When a first data value of the stripe is written, it is stored as a currant value in a parity page of the parity buffer, the parity page corresponding to the stripe. For each subsequent data value that is written, an XOR operation is performed with the subsequent data value and the current value of the corresponding parity page and the result of the XOR operation is stored as the current value of the corresponding parity page.

Sun at Abstract.

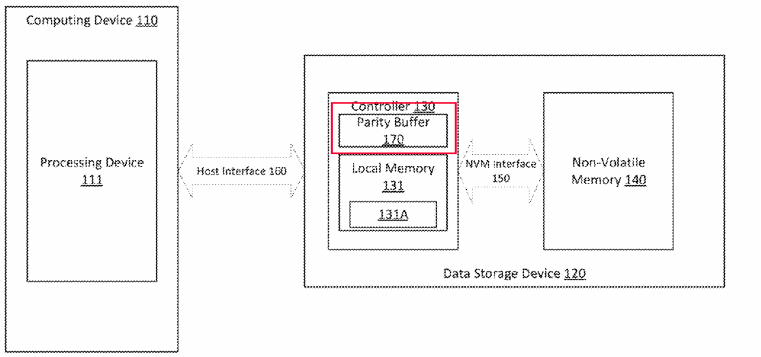

Sun at Fig. 1, annotation added, see also Fig. 6.

82. The Accused Products generate a parity block storing the first data block of the sequence of data blocks into the page frame and update the parity buffer with the first block written to the stripe. For example, in Sun, WD describes generating and updating the parity block.

(57) **ABSTRACT**

An apparatus is disclosed having a parity buffer having a plurality of parity pages and one or more dies, each die having a plurality of layers in which data may be written. The apparatus also includes a storage controller configured to write a stripe of data across two or more layers of the one or more dies, the stripe having one or more data values and a parity value. When a first data value of the stripe is written, it is stored as a currant value in a parity page of the parity buffer, the parity page corresponding to the stripe. For each subsequent data value that is written, an XOR operation is performed with the subsequent data value and the current value of the corresponding parity page and the result of the XOR operation is stored as the current value of the corresponding parity page.

Sun at Abstract.

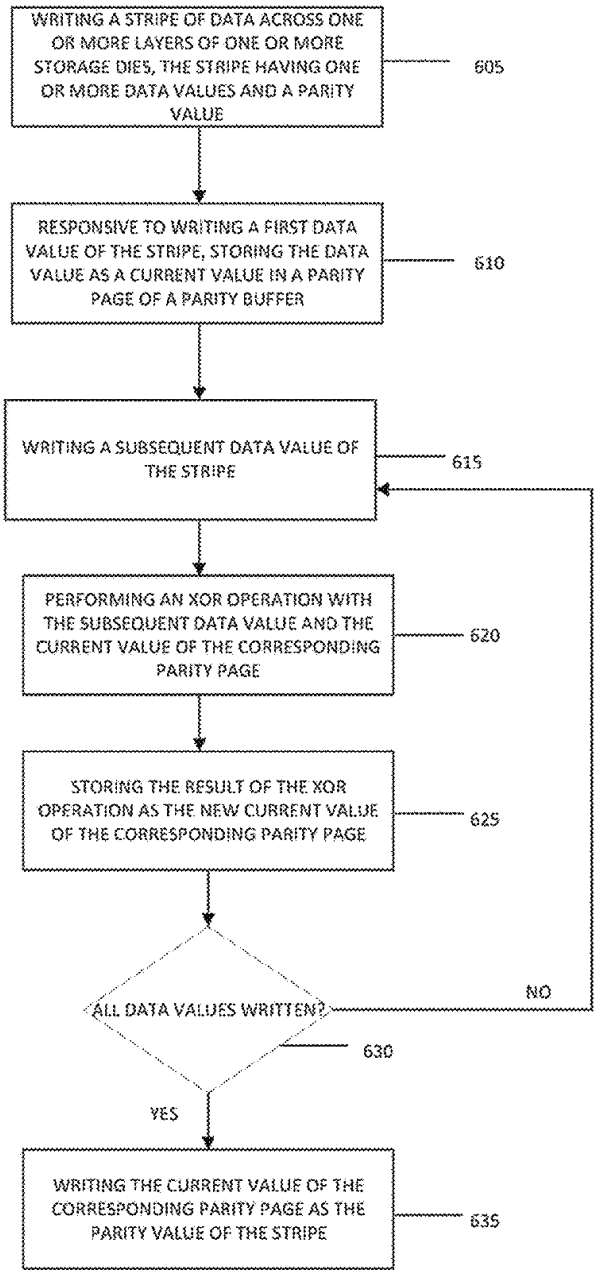

Sun at Fig. 6.

> At block 610, responsive to writing a first data value of the stripe, the controller may store the first data value as a current value in a parity page of a parity buffer. The parity page may correspond to the stripe being written. At block 615, the controller may write a subsequent data value. At block 620, the controller may perform an XOR operation with the subsequent data value and the current value of the corresponding parity page. At block 625, the controller may store the result of the XOR operation as the new current value of the corresponding parity page.

Sun at 9:43-52.

83. In The Accused Products, the command processing unit that receives the parity write command is configured to write the parity block to a nonvolatile memory based on the parity write command. For example, in Sun, WD describes writing the parity to nonvolatile memory.

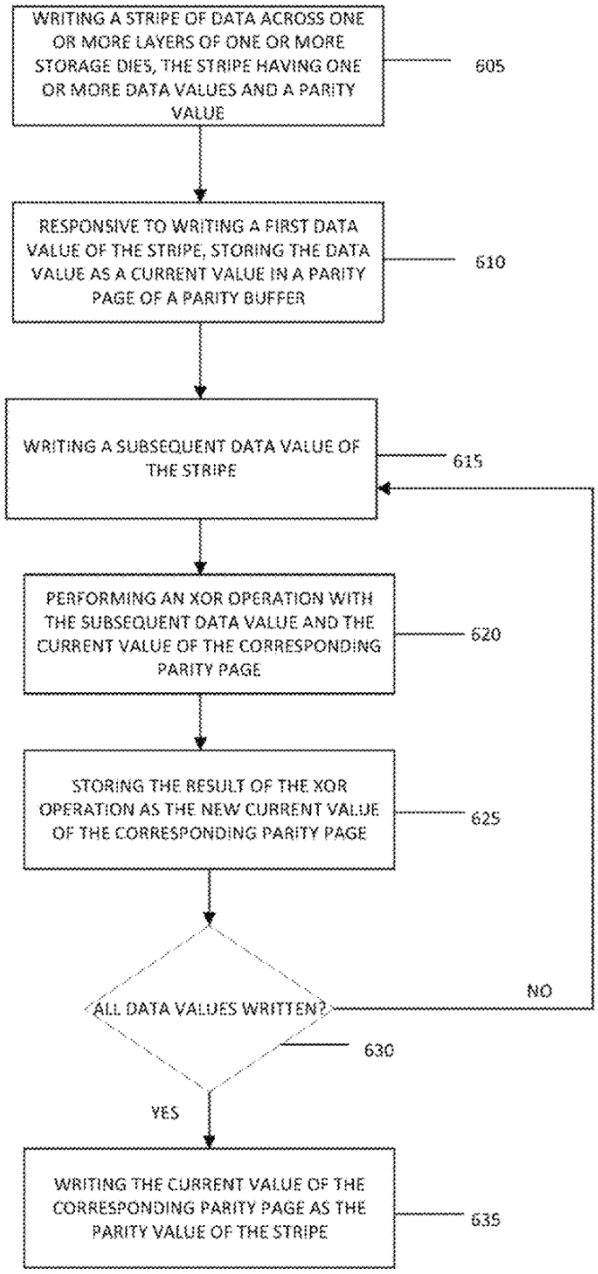

Sun at Fig. 6.

> At block 630, the controller may determine whether all the data values are written. If there are still data values to be written, the controller returns to block 615. If all

> data values are written, then at block 635 the controller writes the current value of the corresponding parity page to the parity value of the stripe.

Sun at 9:52-58.

84. For example, Defendants provide SSD products configured with the hardware and software that satisfy the limitations of at least claim 1. Defendants further directly infringe the '346 Patent when its SSD products with NVMe are installed and operated by its employees in a computer system, such as for gaming, personal computing, or data centers. Direct infringement further occurs when Defendants' employees use and test the hardware and software.

85. On January 23, 2024, Plaintiff sent a notice letter to Defendants, notifying them that they were infringing the Asserted Patents. Defendants' infringement of the '346 Patent has been willful and intentional under the standard announced in *Halo Elecs., Inc. v. Pulse Elecs., Inc.*, 579 U.S. 93, 136 S. Ct. 1923, 195 L. Ed 2d 278 (2016). Defendants have willfully infringed the '346 Patent by refusing to take a license and continuing to make, use, test, sell, license, and/or offer for sale/license the '346 Patent Accused Products. Instead of taking a license, Defendants have opted to make the business decision to "efficiently infringe" the '346 Patent. In doing so, Defendants willfully infringe the '346 Patent.

86. Defendants also have infringed, and continue to infringe the claims of the '346 Patent by actively inducing others to use the Accused Products. Defendants' users, customers, agents or other third parties who use the Accused Products in accordance with Defendants' instructions infringe the claims of the '346 Patent, in violation of 35 U.S.C. § 271(a). Defendants induce its customers to use its SSD products for the benefits of higher performance, reduced size and power, and increased reliability compared to other SSD products. Defendants are thereby liable for infringement of the '346 Patent under 35 U.S.C. § 271(b).

87. Defendants' users, customers, agents or other third parties who use the Accused

Products in accordance with Defendants' instructions infringe the claims of the '346 Patent, in violation of 35 U.S.C. § 271(b). Defendants intentionally instruct its customers to infringe through support information, demonstrations, brochures, videos, and user guides, such as those located at:

https://www.westerndigital.com/support;
https://support.wdc.com/contact.aspx?lang=en;
https://www.westerndigital.com/support/category-selection;
https://www.youtube.com/channel/UC_M0BlntDVSblWg6ggUa7UQ;
https://www.youtube.com/watch?v=L0zZtwRFj0E;
https://www.youtube.com/westerndigital; and
https://www.youtube.com/c/westerndigitalcorporation.

88. Defendants are liable as contributory infringers of the '346 Patent under 35 U.S.C. § 271(c) by having offered to sell, sold and imported and continuing to offer to sell, selling, and importing into the United States its SSD products, to be especially made or adapted for use in an infringement of the '346 Patent. Defendants' SSD products are key components in gaming consoles and gaming PCs, computers, everyday PCs, Network Attached Storage (NAS) devices, servers and data centers. These SSD products are material components for use in practicing the '346 Patent and are specifically made and are not a staple article of commerce suitable for substantial noninfringing use. Defendants supplied these components with knowledge of the '346 Patent and with knowledge that these components constitute material parts of the claimed inventions of the '346 Patent.

89. As a result of Defendants' infringement of the '346 Patent, Plaintiff has suffered monetary damages and is entitled to no less than a reasonable royalty for Defendants' use of the claimed inventions of the '346 Patent, together with interest and costs as determined by the Court. Plaintiff will continue to suffer damages in the future.

90. Defendants are on notice of their infringement by no later than the filing and service

of this Complaint. By the time of trial, Defendants will have known and intended (since receiving such notice) that its continued actions would actively induce to the infringement of the '346 Patent.

91. Defendants' acts of direct and indirect infringement have caused and continue to cause damage to Plaintiff. Plaintiff is entitled to damages in accordance with 35 U.S.C. §§ 271 and 281 sustained as a result of Defendants' wrongful acts in an amount to be proven at trial.

PRAYER FOR RELIEF

WHEREFORE, Plaintiff requests the following relief from this Court:

A. A judgment that each defendant is liable for infringement of one or more claims of the '968 Patent, the '346 Patent, and the '085 Patent;

B. Compensatory damages in an amount according to proof, and in any event no less than a reasonable royalty, including all pre-judgment and post-judgment interest at the maximum rate allowed by law;

C. A judgment granting Plaintiff such further relief as the Court may deem just and proper; and

D. That the Court declare this an exceptional case and award Plaintiff its attorneys' fees, as provided by 35 U.S.C. § 285 and that Plaintiff be awarded enhanced damages up to treble damages for willful infringement as provided by 35 U.S.C. § 284.

DEMAND FOR JURY TRIAL

Pursuant to Rule 38(b) of the Federal Rules of Civil Procedure, Plaintiff demands a trial by jury for all issues so triable.

Dated: January 30, 2024

/s/ *Deron R. Dacus*

Deron R. Dacus
Texas Bar No. 00790553
ddacus@dacusfirm.com
THE DACUS FIRM, P.C.

821 ESE Loop 323, Suite 430
Tyler, Texas 75701
Telephone: (903) 705-7233
Facsimile: (903) 581-2543

Robert F. Kramer (*pro hac vice* forthcoming)
rkramer@krameralberti.com
David Alberti (*pro hac vice* forthcoming)
dalberti@krameralberti.com
Sal Lim (*pro hac vice* forthcoming)
slim@krameralberti.com
Russell S. Tonkovich (*pro hac vice* forthcoming)
rtonkovich@krameralberti.com
KRAMER ALBERTI LIM & TONKOVICH LLP
577 Airport Blvd., Suite 250
Burlingame, California 94010
Telephone: (650) 825-4300
Facsimile: (650) 460-8443

Attorneys for Plaintiff
Polaris PowerLED Technologies, LLC

CERTIFICATE OF SERVICE

I hereby certify that a copy of the foregoing document was filed electronically in compliance with Local Rule CV-5(a). Therefore, this document was served on all counsel who are deemed to have consented to electronic service on this 30th day of January 2024.

By: */s/ Deron R. Dacus*
Deron R. Dacus