

**IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF TEXAS
WACO DIVISION**

BiTMICRO LLC,

Plaintiff,

v.

INTEL CORPORATION,

Defendant.

Civil Action No.: 6:22-cv-335

JURY TRIAL DEMANDED

COMPLAINT FOR PATENT INFRINGEMENT

Plaintiff BiTMICRO LLC (“Plaintiff” or “BiTMICRO”), through its attorneys, for its Complaint against Intel Corporation (“Intel” or “Defendant”), demands a trial by jury and alleges as follows:

FACTUAL INTRODUCTION

1. The novel inventions disclosed in the Asserted Patents in this matter were invented by BiTMICRO Networks, Inc. (“BNI”). BNI was founded in 1995 and was a leader in enterprise storage for mission-critical computing, particularly for military applications. BNI’s storage devices are best known for exceeding the extreme performance and data integrity required for enterprise, industrial, and military environments.

2. BNI made critical advances in the solid state drive (“SSD”) and integrated circuit technology that is embodied in the Asserted Patents. The Asserted Patents in this case are the result of the work from 15 different BNI engineers and developers, spanning a period of over a decade.

3. Innovation was one of the keys to the success at BNI. The company was involved with research and development projects in the SSD industry for about 20 years, over which time it accumulated over 50 U.S. patents, all of which are now owned by BiTMICRO.

THE PARTIES

4. BiTMICRO is the current owner and assignee of the Asserted Patents and holds all rights necessary to bring this action.

5. BiTMICRO is a Delaware limited liability company with its principal place of business located at 11921 Freedom Drive, Suite 550, Reston, Virginia 20190.

6. Defendant Intel Corporation is a Delaware corporation with its principal place of business at 2200 Mission College Boulevard, Santa Clara, California 95054.

7. Intel is registered to do business in the State of Texas and has been since at least April 1989. Defendant may be served by serving its registered agent, C T Corporation System, 1999 Bryan St., Ste. 900, Dallas, TX 75201-3136.

8. Intel has places of business at 9442 N. Capital of Texas Hwy., Bldg. 2, Suite 600, Austin, Texas 78759; and 1300 S. Mopac Expressway, Austin, Texas 78746.

JURISDICTION AND VENUE

9. This action arises under the patent laws of the United States, Title 35 of the United States Code. Subject matter jurisdiction is proper in this Court pursuant to 28 U.S.C. §§ 1331 and 1338(a).

10. This Court has personal jurisdiction over Intel in accordance with due process and/or the Texas Long Arm Statute because, in part, Intel “[r]ecruits Texas residents, directly or through an intermediary located in this state, for employment inside or outside this state.” *See* Tex. Civ. Prac. & Rem. Code § 17.042.

11. This Court also has personal jurisdiction over Intel because it has committed and continues to commit acts of direct and/or indirect infringement in this judicial district in violation of at least 35 U.S.C. §§ 271(a) and (b). In particular, on information and belief, Intel has made, used, offered to sell, and/or sold the accused products in this judicial district, and has induced others to use the accused products in this judicial district.

12. This Court also has personal jurisdiction over Intel because, *inter alia*, Intel (1) has substantial, continuous, and systematic contacts with this State and this judicial district; (2) owns, manages, and operates facilities in this State and this judicial district; (3) enjoys substantial income from its operations and sales in this State and this judicial district; (4) employs Texas residents in this State and this judicial district; and (5) solicits business and markets products, systems and/or services in this State and judicial district including, without limitation, those related to the infringing accused products.

13. Venue is proper in this District pursuant to at least 28 U.S.C. § 1319(b)-(c) and § 1400(b), at least because Intel, either directly or through its agents, has committed acts within this judicial district giving rise to this action, and continue to conduct business in this District, and/or has committed acts of patent infringement within this District giving rise to this action.

FACTUAL ALLEGATIONS

BiTMICRO Patents

14. The BiTMICRO inventions contained in the Asserted Patents in this case relate to groundbreaking improvements to memory controllers, mapping tables for memory devices, storage device security, memory saving during power loss, microchip configuration, and system booting as will be further described below.

U.S. Patent No. 7,826,243

15. On November 2, 2010, the U.S. Patent and Trademark Office duly and lawfully issued United States Patent No. 7,826,243 (“the ’243 Patent”), entitled “Multiple chip module and package stacking for storage devices.” A true and correct copy of the ’243 Patent is attached hereto as **Exhibit A**.


16. BiTMICRO is the owner and assignee of all right, title, and interest in and to the ’243 Patent, including the right to assert all causes of action arising under said patent and the right to any remedies for infringement of it.

17. The ’243 Patent describes, among other things, module stacking and packaging for electronic devices. Specifically, multiple modules, such as semiconductor dies, are stacked together and connected to create a stacked module. A stacked module creates a package with specific functions or a range of memory capacity. Multiple packages can be stacked to create a desired memory capacity or different packages can be stacked to create a desired functionality. The stacking techniques described in the ’243 Patent allow large capacity storage and high functionality devices to be implemented in a smaller package, saving space on an integrated circuit board and thus allowing smaller devices to include memory capacity or functionality that may otherwise be unachievable. The ’243 Patent further describes interconnections and signal routing between two or more of the plurality of modules. The signal paths described in the ’243 Patent may be used to test stacked modules to ensure proper function. Furthermore, the signal paths may be used in the operation of a stacked module.

18. Prior to the inventions in the ’243 Patent, processing and memory chips in computer systems were typically arranged as separate packages within a device, with electrical connections between the packages to facilitate communications between the packages. For example, in an

electronic storage device, separate packages were necessary for an interface controller, a DMA controller, a processor, and the memory modules. This type of design required a significant amount of physical space within the device and therefore presented an obstacle to miniaturization of the device. The inventions in the '243 Patent solved this problem by providing a way for processing and memory chips to be stacked on top of each other to minimize the amount of space that the chips occupy within the device while maintaining reliable interconnectivity between the chips. (See Ex. A at 1:6-2:2).

19. An exemplary embodiment of the '243 inventions is depicted in Figure 21a of the '243 Patent as follows:



20. As shown in Figure 21a, a main board 2101 has multiple chip modules 2103, 2106, 2130, and 2108 stacked on top of it. The modules are electrically connected to each other through a series of ball and pad connections. For example, an input signal 2102 can be passed from the main

board through the IN ball 2104 to the OUT pad 2105 of base module 2103. From there the input signal can be passed through the IN balls and OUT pads of the other modules stacked above in a serial chain until the signal reaches the end module 2108. In the end module, the signal can be routed internally through the module via connection 2112 to another pad 2110 on the end module. The signal can then be passed back down through the stack of modules to the main board through a similar serial chain of ball and pad connections between the modules. (See Ex. A at 9:61-10:35).

21. Figure 21b in the '243 Patent provides an example of how the end module can be configured to receive the incoming serial chain signal and route the signal to other pads:




Figure 21b

(See also Ex. A at 10:8-35).

22. The novel features of the '243 inventions are recited in the claims. For example, claim 1 of the '243 Patent recites:

A stacked module comprising a plurality of modules each comprising:

one or more active ports for carrying one or more active signal;

one or more passive ports for passing through the one or more active signals;

a first serial chain route that includes at least one serial chain connection, the serial chain connection including: a serial chain circuit, a serial chain input, and serial chain output; said serial chain input coupled to said serial chain output through said serial chain circuit;

a second serial chain route and a control circuit for enabling a routing path that connects the first serial chain route with the second serial chain route within an end module;

and said control circuit is disposed to enable the routing path in response to a control input signal received from another module from the plurality of modules when the end module is coupled to the another module.

(Ex. A at 10:42-60). Claim 1 of the '243 Patent describes claim elements individually or as an ordered combination, that were non-routine and unconventional at the time of the inventions in 2005 and an improvement over prior art, as it provided a semiconductor stacking and packaging design (not previously available) that allowed signals to be carried and routed through multiple stacked modules via serial chain routes. (*See id.* at Abstract, 1:6-2:2).

U.S. Patent No. 9,135,190

23. On September 15, 2015, the U.S. Patent and Trademark Office duly and lawfully issued United States Patent No. 9,135,190 (“the '190 Patent”), entitled “Multi-profile memory controller for computing devices.” A true and correct copy of the '190 Patent is attached hereto as **Exhibit B.**

24. BiTMICRO is the owner and assignee of all right, title, and interest in and to the '190 Patent, including the right to assert all causes of action arising under said patent and the right to any remedies for infringement of it.

25. The '190 Patent describes, among other things, a multi-profile memory controller for computing devices. Specifically, the '190 Patent describes a memory controller that can operate with memory locations, memory devices, or both which are associated with different memory attributes, different attribute qualifiers, or the like. For example, a non-volatile memory storage device may be portioned to allow a memory controller to treat a portion of the memory device as a temporary cache memory to store data prior to writing the data to a permanent storage location. This eliminates the need for a separate memory cache, often composed of volatile memory. This has the additional advantage of maintaining the temporary data in the non-volatile cache partition in the event of an unexpected power loss.

26. Prior to the inventions in the '190 Patent, memory controllers were designed to operate with memory locations and memory devices that all shared the same set of memory device characteristics, such as block size. Due to this limitation, there was no way of varying how a memory controller performed read and write operations on different memory locations or memory devices. The '190 Patent overcame this limitation by disclosing a novel multi-profile memory controller with the ability to operate differently with memory locations and memory devices based on differences between the attributes of particular memory locations and memory devices. (*See* Ex. B, at 1:20-60).

27. As described in the '190 Patent, a memory store includes multiple addressable memory locations, and each location is associated with a set of memory attributes, which can include, for example, the type of memory device in which the memory location is located, the data

size used by the memory device, or the memory protocol of the device. (*See* Ex. B, at 2:63-3:13). These attributes are organized into device profiles that can be used by a memory controller connected to the memory store to determine how memory transactions are to be performed with each memory location. (*See id.* at 3:14-4:25). By analyzing the requirements of the requested memory transaction and comparing those requirements to the device profiles, the memory controller selects the appropriate memory location for the memory transaction. The criteria used by the controller to select the optimal memory location for the memory transaction based on the stored attributes within the device profiles can be programmed in any number of ways. (*See id.* at 7:60-9:48).

28. The novel features of the '190 inventions are recited in the claims. For example, claim 59 of the '190 Patent recites:

A memory controller comprising:

an interface controller coupled to a memory device interface and an input/output (IO) device interface;

a memory store;

wherein the memory device interface is directly coupled to the memory store;

said interface controller disposed to perform a memory transaction by addressing a first memory location in the memory store,

said first memory location and a second memory location respectively associated with a first device profile and a second device profile;

wherein said first device profile is optimal for a data type subject to the memory transaction, wherein said data type comprises one of a random data type or a sequential data type;

said interface controller identifies command details for causing the memory transaction to be performed, wherein said command details comprising the first memory device;

said device profile representing a first set of attributes of said first memory location, and said second device profile representing a second set of attributes of said second memory location, and a difference exists between said first and second device profiles;

said interface controller obtaining the first set of attributes after identifying the command details; and said addressing of said first memory location includes using said attributes from said first device profile;

and said addressing of said first memory location includes selecting a transfer size for the memory transaction, wherein the transfer size is a function of a data size of the memory transaction and the first set of attributes.

(Ex. B at 18:34-65). Claim 59 of the '190 Patent describes claim elements individually or as an ordered combination, that were non-routine and unconventional at the time of the inventions in 2009 and an improvement over prior art, as it provided a memory controller (not previously available) with an interface controller capable of performing memory transactions with different transfer sizes on different memory locations based on attributes associated with the different memory locations as defined in differing device profiles for those memory locations. (*See id.* at Abstract, 1:20-60, 2:41-62).

U.S. Patent No. 8,010,740

29. On August 30, 2011, the U.S. Patent and Trademark Office duly and lawfully issued United States Patent No. 8,010,740 (“the ’740 Patent”), entitled “Optimizing memory operations in an electronic storage device.” A true and correct copy of the ’740 Patent is attached hereto as **Exhibit C**.

30. BiTMICRO is the owner and assignee of all right, title, and interest in and to the ’740 Patent, including the right to assert all causes of action arising under said patent and the right to any remedies for infringement of it.

31. The ’740 Patent describes, among other things, a mapping table for optimizing memory operations in an electronic storage device. Prior to the inventions in the ’740 Patent,

memory operations in solid state storage devices were subject to a number of inefficiencies. As described in the '740 specification, SSDs such as those that include NAND flash memory “suffer from write cycle limitations and to a certain degree, bad blocks. In addition, flash drives use block addressing rather than byte addressing, and these flash drives use block addresses that are usually much larger than the block address used by the host. Block addressing may impose an additional level of complexity and additional processing cycles when performing a write operation, and which in turn, may increase write operation latency. This additional level of complexity may include performing a read-modify-write transaction to complete the write operation.” (Ex. C at 1:42-53).

32. To address these issues and increase the speed and efficiency of memory operations in their products, SSD manufacturers tried solutions such as adding complex algorithms to handle the management of memory operations and adding more powerful processing devices to run these complex algorithms. (*See id.* at 1:54-2:10). These solutions, however, increased both the cost and design complexity of the SSDs. (*See id.*).

33. The '740 Patent overcame this problem by providing a solution that optimizes memory operations in a solid state storage device while minimizing the amount of additional cost and complexity to the design of the device. (*See id.* at 2:11-14). The '740 Patent achieves this through an improved mapping table that “increas[es] the likelihood that, in response to an I/O transaction initiated by a host, the operational load imposed on the storage device by these memory operations will be optimally distributed across different storage device resources, such as by interleaving or parallel memory operations, reducing memory operation latency, increasing operational device efficiency, or both.” (*Id.* at 2:14-21; *see also id.* at 3:12-31). For example, the '740 Patent describes a mapping table that includes a set of logical fields that represent a plurality of logical block address (LBA) sets. The mapping table also includes a set of physical block

address (PBA) fields that represent a set of PBAs and access parameters for the PBAs, as well as information that associates the LBA sets with the PBA sets in a highly efficient manner. The mapping table enables the storage device to perform optimized memory operations on memory locations based on the information in the table regarding the relationship between the LBA and PBA sets and the access parameters. (*See id.* at 2:27-49).

34. The novel features of the '740 inventions are recited in the claims. For example, claim 1 of the '740 Patent recites:

A mapping table for optimizing memory operations performed by an electronic storage device in response to receiving an I/O transaction request initiated by a host, said mapping table comprising:

a set of logical fields, including a first logical field and a second logical field, and said logical fields respectively disposed for representing a plurality of LBA sets, including said first logical field disposed for representing a first LBA set and said second logical field disposed for representing a second LBA set, said first and second LBA sets each representing a set of consecutive LBAs;

a set of PBA fields, including a first PBA field and a second PBA field, said set of PBA fields respectively disposed for representing a set of PBAs, including a first PBA disposed for representing a first set of access parameters and a second PBA disposed for representing a second set of access parameters, said PBAs each associated with a physical memory location in a memory store, said set of logical fields and said set of PBA fields disposed to associate said first and second LBA sets with said first and second PBAs;

and wherein, in response to receiving the I/O transaction request, said mapping table causes the electronic storage device to perform optimized memory operations on memory locations respectively associated with said first PBA and said second PBA, if the I/O transaction request is associated with said first and second LBA sets.

(Ex. C at 9:64-10:24). Claim 1 of the '740 Patent describes claim elements individually or as an ordered combination, that were non-routine and unconventional at the time of the inventions in 2006 and an improvement over prior art, as it provided a mapping table (not previously available) that enables optimized memory operations in an electronic storage device through information

stored in a mapping table regarding logical fields, PBA fields, access parameters, and relationships between LBA and PBA sets. (*See id.* at Abstract, 2:27-49).

U.S. Patent No. 9,858,084 and U.S. Patent No. 10,120,694

35. On January 2, 2018, the U.S. Patent and Trademark Office duly and lawfully issued United States Patent No. 9,858,084 (“the ’084 Patent”), entitled “Copying of power-on reset sequencer descriptor from nonvolatile memory to random access memory.” A true and correct copy of the ’084 Patent is attached hereto as **Exhibit D**.

36. BiTMICRO is the owner and assignee of all right, title, and interest in and to the ’084 Patent, including the right to assert all causes of action arising under said patent and the right to any remedies for infringement of it.

37. On November 6, 2018, the U.S. Patent and Trademark Office duly and lawfully issued United States Patent No. 10,120,694 (“the ’694 Patent”), entitled “Embedded system boot from a storage device.” A true and correct copy of the ’694 Patent is attached hereto as **Exhibit E**.

38. BiTMICRO is the owner and assignee of all right, title, and interest in and to the ’694 Patent, including the right to assert all causes of action arising under said patent and the right to any remedies for infringement of it.

39. The ’084 Patent and ’694 Patent share the same specification, and both patents claim priority to U.S. Provisional Application No. 61/801,952, filed with the U.S. Patent and Trademark Office on March 15, 2013.

40. The ’084 and ’694 Patents describe, among other things, methods and systems for booting up and initializing an embedded system directly from a storage device.

41. As described in the ’084 and ’694 Patents, “[a]n embedded system refers to an apparatus that is made up of a single or multiple controller chips interconnected to each other on a

printed circuit board” and can also refer to “a single CPU or a multiprocessor system.” (Ex. D at 1:22-30; Ex. E at 1:20-28). Prior to the inventions in the ’084 and ’694 Patents, embedded systems relied on firmware (such as a boot code loader, a basic input/output system (BIOS), an operating system (OS), or application firmware) in order to boot up and initialize the system. As part of this process, the boot code loader would need to copy the BIOS, OS, or application firmware from an external mass storage device to the random access memory (RAM) of the embedded system, whereupon the system’s central processing unit (CPU) would read and execute the copied BIOS, OS, or application firmware to perform further initialization and diagnostic tasks before the system could proceed to normal operations. The boot up process for embedded systems therefore required the transfer and loading of large amounts of data before the system could begin normal operations, and there was no ability for the system to self-check the integrity of the BIOS, OS, or application firmware transferred into the RAM. (*See* Ex. D at 1:31-55; Ex. E at 1:29-53).

42. The inventions in the ’084 and ’694 Patents addressed this problem by providing more efficient and reliable ways of handling the boot up process of an embedded system using a novel power-on reset (POR) sequencer stored in the nonvolatile memory of the embedded system. By way of example and without limitation, a POR sequencer descriptor may contain data necessary for setting up the configuration registers and managing other initialization processes in the system. Upon power-up of the system, the POR sequencer in the system attempts to retrieve a POR sequencer descriptor from the nonvolatile memory. If a POR sequencer descriptor is detected, the POR sequencer transfers it to the system RAM. The POR sequencer may then perform an integrity check of the POR sequencer descriptor to determine if there are any errors. If no error is detected, the POR sequencer can use the POR sequencer descriptor to, among other things, initialize the

configuration registers as well as coordinate further data transfers necessary for the initialization process with any DMA controllers in the system. (*See* Ex. D at 2:13-4:4; Ex. E at 2:10-4:2).

43. The novel features of the inventions of the '084 and '694 Patents are recited in the claims. For example, claim 19 of the '084 Patent recites:

An article of manufacture, comprising:

a non-transitory computer-readable medium having stored thereon instructions operable to permit an apparatus to perform a method comprising:

releasing components of an embedded system from reset; detecting a power-on reset (POR) sequencer descriptor in a nonvolatile memory;

wherein the POR sequencer descriptor comprises information to initialize configuration registers of the embedded system, Direct Memory Access (DMA) descriptors used to fetch other POR sequencer descriptor fragments, and a system firmware;

copying the POR sequencer descriptor from the nonvolatile memory to a random access memory;

verifying an integrity of the POR sequencer descriptor; and

detecting any error in the POR sequencer descriptor.

(Ex. D at 10:49-65). Claim 19 of the '084 Patent describes claim elements individually or as an ordered combination, that were non-routine and unconventional at the time of the inventions in 2013 and an improvement over prior art, as it provided methods of starting up an embedded system (not previously available) using a novel POR sequencer descriptor stored in nonvolatile memory whose integrity is verified and is checked for any errors before start up. (*See id.* at Abstract, 1:16-2:60).

44. As another example, claim 6 of the '694 Patent recites:

An apparatus, comprising:

an embedded system comprising one or more processors, a reset controller, a storage device controller, one or more direct memory access (DMA) controllers, a

random access memory (RAM) and a memory controller, a nonvolatile memory and a nonvolatile memory controller, a debug interface, and

a power-on reset (POR) sequencer, wherein the POR sequencer uses a POR sequencer descriptor which is a preassembled descriptor that is stored in the nonvolatile memory,

wherein the POR sequencer descriptor includes register information and DMA controller descriptors, and

wherein the POR sequencer reads and uses the register information to update configuration registers of the embedded system.

(Ex. E at 9:64-10:12). Claim 6 of the '694 Patent describes claim elements individually or as an ordered combination, that were non-routine and unconventional at the time of the inventions in 2013 and an improvement over prior art, as it provided an apparatus capable of starting up an embedded system (not previously available) using a minimal amount of data that is used by a novel POR sequencer to update the configuration registers of the system. (*See id.* at Abstract, 1:14-2:58).

U.S. Patent No. 6,496,939

45. On December 17, 2002, the U.S. Patent and Trademark Office duly and lawfully issued United States Patent No. 6,496,939 (“the '939 Patent”), entitled “Method and system for controlling data in a computer system in the event of a power failure.” A true and correct copy of the '939 Patent is attached hereto as **Exhibit F**.

46. BiTMICRO is the owner and assignee of all right, title, and interest in and to the '939 Patent, including the right to assert all causes of action arising under said patent and the right to any remedies for infringement of it.

47. The '939 Patent describes, among other things, systems and methods for controlling data in a computer or memory system when the system loses external power. (*See Ex. F at Abstract*). Specifically, the '939 Patent describes a plurality of capacitors configured to supply power to the computer or memory system, activating those capacitors to supply power in the event

of a sudden loss of external power, and reconfiguring the data in the computer or memory system while that auxiliary power is supplied. (*Id.*). Through the use of the systems and methods described by the '939 Patent, newly written and/or modified data can be stored (or flushed) from the volatile memory to the non-volatile memory in the event of a sudden power loss, and/or enable the data to be erased automatically or manually—in all cases, preserving the security of the stored data. (*Id.*).

48. Prior to the inventions in the '939 Patent, computing systems utilized various types of memory for storing and managing data, including volatile memory for main memory and cache, as well as mechanical disk drives for longer term storage. (*Id.* at 1:14-33). One disadvantage of volatile memory, however, is that it requires a constant source of power, or the data stored thereon will be lost. Consequently, when there is a sudden loss of power, there is insufficient time to safely write all the newly written and/or modified data from the volatile memory to the hard disk drive before the system shuts down – resulting in data loss and insecurity. (*Id.* at 1:34-45). Batteries could be coupled to the internal power system, but those have a number of disadvantages, including a limited number of charge-drain cycles, large internal resistance, size, and weight. A battery-backed uninterruptable power supply (UPS) could be coupled to the external power supply of the computer system, but those systems suffer from the same or similar disadvantages. (*Id.* at 1:46-58). In addition, prior to the inventions of the '939 Patent, computing systems did not have a means for ensuring that the data to be securely erased when the system loses power. (*Id.* at 1:59-2:21).

49. The '939 Patent overcame these various limitations by disclosing novel systems and methods for controlling data in a computer system when the system loses power, comprised of a plurality of capacitors configured to supply power to the computer system, activating those capacitors to supply power in the event of a sudden loss of external power, and reconfiguring the

data in the computer system while that auxiliary power is supplied. (*Id.* at 2:24-37). In addition, these systems and methods allow the computing system to rapidly erase data from a non-volatile memory automatically, in the event of a sudden loss of power, and/or manually. (*Id.* at 2:37-43).

50. The novel features of the inventions of the '939 Patent are recited in the claims. For example, claim 10 of the '939 Patent recites:

1. A system for controlling data in a computer system when the computer system loses power, the computer system comprising a computer engine, comprising:

means for activating a plurality of super capacitors to supply power to the computing engine based upon power being removed from the computer system;

means for reconfiguring the data in the computing engine; and

means for deactivating the plurality of super capacitors to cut off power to the computing engine based upon the plurality of super capacitors discharging to a predetermined level.

(Ex. F at 7:20-31). Claim 10 of the '939 Patent describes claim elements individually or as an ordered combination, that were non-routine and unconventional at the time of the inventions in or before 1999 and an improvement over prior art, as it used a plurality of capacitors configured to supply power to the computer system when the system loses power, thereby preserving the security of the data stored in volatile memory therein. (*See* Ex. E at Abstract, 1:14-2:21, 2:24-43).

Intel's Use of the Patented Technology

51. Intel is a multinational technology company and the world's largest semiconductor chip manufacturer. Founded in 1968, Intel's products include processors, flash memory, SSDs, field-programmable gate arrays (FPGAs) and other products that are incorporated into servers, laptops, desktops, smart phones, and tablet computers, which infringe the Asserted Patents. The infringing products also include stacked electronic components such as those found in SDRAM memory, processor/DRAM packages, application processors, high-bandwidth memory, and

products (e.g., mobile phones, tablets, and watches) that incorporate those components or other stacked electronics components. Specific examples of infringing products made, sold, and/or offered for sale in the United States, and/or imported into the United States are discussed below.

52. Intel makes, uses, sells, and/or offers to sell in the United States, and/or imports into the United States (or has made, used, sold, offered for sale, and/or imported into the United States) vertically stacked integrated circuits under the Foveros brand, which infringe one or more claims of the '243 Patent. Such products include the Lakefield Foveros product.

53. Intel makes, uses, sells, and/or offers to sell in the United States, and/or imports into the United States (or has made, used, sold, offered for sale, and/or imported into the United States) SSDs with SLC caching capabilities, which infringe one or more claims of the '190 Patent. Such SSDs include, for example, Intel's 545s Series, Pro 5450s Series, 600p Series, Pro 6000p Series, 660p Series, 665p Series, 670p Series, 760p Series, and Pro 7600p Series SSDs.

54. Intel makes, uses, sells, and/or offers to sell in the United States, and/or imports into the United States (or has made, used, sold, offered for sale, and/or imported into the United States) memory products under the term "persistent memory," which infringe one or more claims of the '190 Patent. Such products include Intel's Optane Persistent Memory and Persistent Memory 200 Series (PMem) modules.

55. Intel makes, uses, sells, and/or offers to sell in the United States, and/or imports into the United States (or has made, used, sold, offered for sale, and/or imported into the United States) Non-Volatile Memory Express (NVMe) SSDs, Optane SSDs, and Optane memory products, which infringe one or more claims of the '740 Patent. Such products include Intel's DC P3500, DC P3600, DC P3608, DC P3700, 600p, 660p, 665p, 670p, 760p, Pro 6000p, and Pro 7600p Series SSDs; Optane 800P, 900P, 905P, DC P4800X, DC P4801X, DC D4800X, DC P5800X, and

P1600X Series SSDs; and the Optane Persistent Memory and Persistent Memory 200 Series modules, Optane Memory M10, Optane Memory H10 with Solid State Storage, and Optane Memory H20 with Solid State Storage.

56. Intel makes, uses, sells, and/or offers to sell in the United States, and/or imports into the United States (or has made, used, sold, offered for sale, and/or imported into the United States) FPGA products, which infringe one or more claims of the '084 and '694 Patents. Such products include, for example, Intel's Stratix 10 line of FPGA products.

57. Intel makes, uses, sells, and/or offers to sell in the United States, and/or imports into the United States (or has made, used, sold, offered for sale, and/or imported into the United States) RAID Controllers, which include modules and adapters, along with RAID Maintenance Free Backup Units (RMFBU), which infringe one or more claims of the '939 Patent. Such products include, for example, Intel's Raid Controller RMSP3AD160F and the RMFBU AXCRMFBU.

58. Intel makes, uses, sells, and/or offers to sell in the United States, and/or imports into the United States (or has made, used, sold, offered for sale, and/or imported into the United States) SSDs with power loss protection, which infringe one or more claims of the '939 Patent. Such SSDs include, for example, at least the Intel SSD DC S3500 Series; SSD DC S3700 Series; Optane™ Memory H10 with Solid State Storage; Optane™ SSD DC D4800X Series; SSD D3-S4510 Series; SSD D3-S4610 Series; SSD D5-P4320 Series; SSD D5-P4326 Series; SSD D5-P4420 Series; SSD DC P4510 Series; SSD DC P4511 Series; SSD DC P4610 Series; SSD DC P4618 Series; SSD 750 Series; SSD DC P3500 Series; SSD DC P3520 Series; SSD DC P3600 Series; SSD DC P3608 Series; SSD DC P3700 Series; SSD DC P4500 Series; SSD DC P4501 Series; SSD DC P4600 Series; SSD DC P4608 Series; SSD DC S3320 Series; SSD DC S3510

Series; SSD DC S3520 Series; SSD DC S3610 Series; SSD DC S3710 Series; SSD DC S4500 Series; SSD DC S4600 Series; SSD E 5410s Series; and SSD E 7000s Series.

59. On information and belief, in December 2021, Solidigm, a subsidiary of SK hynix Inc., acquired a portion of Intel's business relating to NAND SSDs, and as a result of that acquisition, some of the SSD products identified above are no longer being sold by Intel but instead are now being sold by Solidigm. To the extent Intel is no longer selling such SSD products, Intel is still liable for its past infringing manufacture, use, sale, offer for sale, and/or importation into the United States of such SSDs.

Notice and Marking

60. As set forth below, Intel has been on constructive and/or actual notice of the Asserted Patents.

61. BiTMICRO has complied with 35 U.S.C. § 287 with respect to the Asserted Patents.

62. The previous owner of the Asserted Patents, BNI, also complied with 35 U.S.C. § 287, and thereby provided notice to the public, including but not limited to Intel, of the Asserted Patents. Specifically, to the extent BNI made, offered for sale, sold, or imported into the United States products covered by the Asserted Patents, BNI marked substantially all of such products with those patent numbers and provided an internet address at which BNI posted information associating the patented products with their corresponding patent numbers in compliance with 35 U.S.C. § 287.

63. For example, BNI's Ace Drive II products, which BNI contended were covered by the '190 and '939 Patents, were sold by BNI with a label affixed on the products listing the '190 and '939 Patents, as well as an internet address at which BNI posted a listing of additional patents it contended were practiced by that product, as shown below:



64. BNI's other products similarly included product labels identifying specific patents by number and/or an internet address at which BNI posted a listing of the patents it contended were associated with each product, including but not limited to the '243, '190, '740, and '939 Patents.

65. BNI's product documentation, which was provided to customers and available to the general public, also identified by number specific patents that BNI contended were practiced by the products and included an internet address at which BNI posted a listing of other patents it contended were associated with the products, pursuant to 35 U.S.C. § 287.

66. On information and belief, BNI has never made, offered for sale, sold, or imported into the United States any products that are covered by the '084 or '694 Patents. Thus, there were no BNI products that would require marking of those patent numbers under 35 U.S.C. § 287.

67. BiTMICRO has not made, offered for sale, sold, or imported into the United States any products that are covered by any of the Asserted Patents. Thus, there are no BiTMICRO products that would require marking under 35 U.S.C. § 287.

68. The Asserted Patents have been widely cited by the industry and by the USPTO during the prosecution of other patents. For example, the '243 Patent has been cited in patents and/or applications by Apple, Mosaid, and other well-known industry participants. The '190 Patent has been cited in patents and/or applications by Western Digital, SanDisk, Micron, Huawei, Microsoft, and other well-known industry participants. The '740 Patent has been cited in patents and/or patent applications by Seagate and other well-known industry participants. The '084 and/or '694 Patents have been cited in patents and/or applications by NGD Systems, and other well-known industry participants. And the '939 Patent has been cited in patents and/or applications by Western Digital, Qualcomm, Samsung, SK hynix, HTC, IBM, Google, Avaya, and other well-known industry participants.

69. Indeed, the Asserted Patents have been used by the USPTO as a basis to reject patent applications filed by well-known industry participants under 35 U.S.C. § 102 and/or § 103. For example, the '243 Patent and its family members have served as the basis for § 102/103 rejections

at least ten times, including against IBM, Mosaid, Micron. The '190 Patent has served as the basis for § 102/103 rejections at least three times, including against SanDisk. The '740 Patent has served as the basis for § 102/103 rejections at least four times, including against Seagate and Samsung. The '084 Patent and/or '694 Patents have served as the basis for § 102/103 rejections at least twice, including against Samsung and SK hynix. And the '939 Patent has served as the basis for § 102/103 rejections at least ten times, including against Amazon, Marvell, and other well-known industry participants.

70. In addition, as set forth in greater detail below, Intel has had actual notice of the '243 and '190 Patents and the rest of the Asserted Patents since at least September 5, 2018 by virtue of its participation in a proceeding before the International Trade Commission involving those patents.

FIRST COUNT

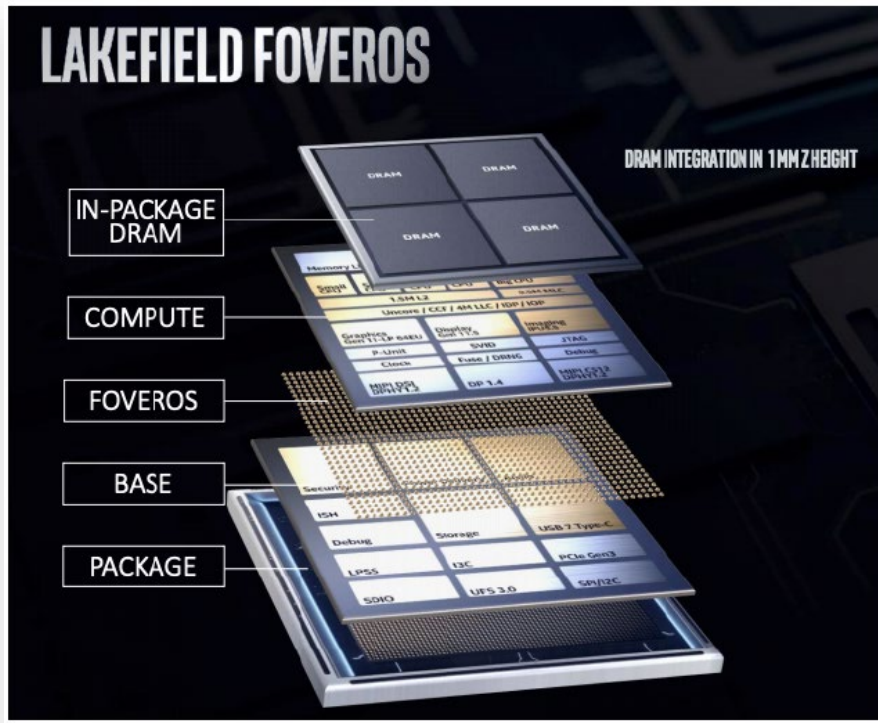
(INFRINGEMENT OF U.S. PATENT NO. 7,826,243)

71. BiTMICRO incorporates by reference the allegations set forth in paragraphs 1-70 as though fully set forth herein.

72. On information and belief, Intel has directly infringed and continues to directly infringe one or more claims of the '243 Patent, including at least claim 1 of the '243 Patent, in the state of Texas, in this judicial district, and elsewhere in the United States by, among other things, making, using, selling, offering for sale, and/or importing into the United States products that embody one or more of the inventions claimed in the '243 Patent, including but not limited to the above-identified Lakefield Foveros product, and all reasonably similar products (“the '243 Accused Products”), in violation of 35 U.S.C. § 271(a).


73. As an example, the '243 Accused Products, such as the Lakefield Foveros product, are “a stacked module comprising a plurality of modules.” Specifically, the Lakefield Foveros

product consists of two layers of DRAM memory modules stacked on top of a compute module, which is in turn stacked on top of a base module, which is in turn stacked on top of a package module:



(Intel presentation, “Lakefield: Hybrid Cores in 3D Package,” p. 13).¹

¹ Available at <https://newsroom.intel.com/wp-content/uploads/sites/11/2019/08/Intel-Lakefield-HotChips-presentation.pdf>.



(Intel video, “Intel Previews New Hybrid CPU Architecture with Foveros 3D Packaging”).²




(Intel 2018 Architecture Day, Intel Architecture Directions presentation, p.19).³

² Available at <https://www.youtube.com/watch?v=-besHp8HLxo>.


³ Available at <https://wccftech.com/intel-architecture-day-2018-raja-koduris-commentary-and-vision-for-intel/>.

74. Multiple layers of the Lakefield Foveros product are connected by a ball and pad array, with a ball sandwiched between a top and bottom pad:



(Intel 2018 Architecture Day, Intel Architecture Directions presentation, p.19).⁴


⁴ Available at <https://wccfttech.com/intel-architecture-day-2018-raja-koduris-commentary-and-vision-for-intel/>.




(WikiChip article, “A Look at Intel Lakefield: A 3D-Stacked Single-ISA Heterogeneous Pentacore SOC”).⁵

⁵ Available at <https://fuse.wikichip.org/news/3417/a-look-at-intel-lakefield-a-3d-stacked-single-isa-heterogeneous-penta-core-soc/>.

75. Intel's Lakefield Foveros product includes "one or more active ports for carrying one or more active signal." For example, the Lakefield Foveros product has pads that carry one or more active signals (and therefore constitute active ports) from the compute chip to the SoC chip:




76. The Lakefield Foveros product includes “one or more passive ports for passing through the one or more active signals.” For example, as shown below, the Lakefield Foveros product has passive ports consisting of a passive ball on the bottom side of the compute chip and passive pad on the top of the base SoC chip. The passive ports pass through the active signals (*e.g.*, the signal depicted by the yellow arrow):




77. Intel’s Lakefield Foveros product includes “a first serial chain route that includes at least one serial chain connection, the serial chain connection including: a serial chain circuit, a serial chain input, and serial chain output; said serial chain input coupled to said serial chain output through said serial chain circuit.” For example, as shown in the figure below, the Lakefield Foveros product includes a serial chain circuit (the signal that couples the ball and pads between the compute module and the SoC module, identified by the top red arrow), a serial chain input (*e.g.*, an input for a signal on the pad on the bottom side of the compute module), and a serial chain output (*e.g.*, an output for a signal on the ball between the compute and SoC modules); said serial chain

input coupled to said serial chain output through said serial chain circuit (*e.g.*, the signal couples the serial chain input and the serial chain output). Thus, as depicted in the top red arrow, the serial chain circuit, serial chain input, and serial chain output form a serial chain connection.




78. Similarly, another serial chain connection exists between the SoC module and the package module, and is depicted in the bottom red arrow in the figure above. The bottom red arrow has corresponding components to the top red arrow (*e.g.*, a serial chain circuit, serial chain input, and serial chain output). The two serial chain connections designated by the top and bottom red arrows are connected to each other by a routing path (depicted in yellow above). The serial chain route includes both of the serial chain connections and their connection to each other through the routing path (*e.g.*, the route shown in the red rectangular label).



79. Intel's Lakefield Foveros product includes "a second serial chain route and a control circuit for enabling a routing path that connects the first serial chain route with the second serial chain route within an end module." For example, as shown in the figure below, the Lakefield Foveros product includes a second serial chain route, which includes at least two serial chain connections (*e.g.*, two serial chain connections depicted by the blue arrows). These serial chain connections include corresponding components to the serial chain connections described with respect to the first serial chain route. The second serial chain route follows the signal path depicted in the blue rectangle (*e.g.*, the signal path from the ball and pads between the package and SoC modules through the ball and pads between the SoC and compute modules). A routing path (depicted in yellow) through the PoP memory module, which constitutes an end module, connects

the first serial chain route with the second serial chain route. A control circuit within the compute module enables the routing path.

80. In Intel's Lakefield Foveros products, the "said control circuit is disposed to enable the routing path in response to a control input signal received from another module from the plurality of modules when the end module is coupled to the another module." For example, the POP memory module, compute module, SoC module, and package module are all coupled together through the stacking of the modules as discussed above. The control circuit in the compute module can enable the routing path in response to a control input received, for example, from the SoC module.



81. By making, using, offering for sale, and/or selling products in the United States and/or importing products into the United States, including but not limited to the '243 Accused Products, Intel has injured BiTMICRO and is liable to BiTMICRO for directly infringing one or more claims of the '243 Patent, including without limitation claim 1 pursuant to 35 U.S.C. § 271(a).

82. On information and belief, Intel has had knowledge of the '243 Patent since at least September 5, 2018, when BiTMICRO served a subpoena on Intel in connection with a proceeding before the International Trade Commission, *In the Matter of Certain Solid State Storage Drives, Stacked Electronics Components, and Products Containing Same*, Inv. No. 337-TA-1097 (“the ITC Action”). In the ITC Action, BiTMICRO alleged that various solid state computer drives (“SSDs”) and electronic devices that incorporate stacked electronics components sold by Samsung Electronics, SK Hynix, and other electronic device manufacturers infringed four BiTMICRO patents, including the '243 Patent.

83. Intel responded to BiTMICRO’s subpoena in the ITC Action by serving objections and responses to the subpoena, conferring with BiTMICRO’s ITC counsel regarding the subpoena, and filing a motion in the ITC on September 14, 2018 for an extension of time to respond to the subpoena.

84. On information and belief, through its participation in the ITC Action as described above, Intel had knowledge of the '243 Patent and its relevance to the '243 Accused Products. Despite this knowledge, Intel has continued to directly infringe one or more claims of the '243 Patent as described above. Thus, on information and belief, Intel’s infringement of the '243 Patent has been willful.

85. On information and belief, Intel is also inducing and/or has induced infringement of one or more claims of the '243 Patent, including at least claim 1, as a result of, amongst other activities, instructing, encouraging, and directing its customers on the use of the '243 Accused Products in an infringing manner in violation of 35 U.S.C. § 271(b). Through its website, instructional guides, and manuals, Intel provides its customers with detailed explanations, instructions, and information on how to use and implement the '243 Accused Products which

demonstrate active steps taken to encourage direct infringement. (*See, e.g.*, Lakefield: Hybrid Cores in 3D Package).⁶ On information and belief, Intel has had knowledge of the '243 Patent since at least September 5, 2018 as set forth above. Despite this knowledge, Intel has continued to engage in activities to encourage and assist its customers in the use of the '243 Accused Products. Thus, on information and belief, Intel (1) had actual knowledge of the patent; (2) knowingly induced its customers to infringe the patent; and (3) had specific intent to induce the patent infringement.

86. On information and belief, by using the '243 Accused Products as encouraged and assisted by Intel, Intel's customers have directly infringed and continue to directly infringe one or more claims of the '243 Patent, including at least claim 1. On information and belief, Intel knew or was willfully blind to the fact that its actions would induce its customers' direct infringement of the '243 Patent.

87. Intel's infringement of the '243 Patent has been and continues to be deliberate and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284-285.

88. On information and belief, Intel will continue to infringe the '243 Patent unless enjoined by this Court.

89. As a result of Intel's infringement of the '243 Patent, BiTMICRO has suffered monetary damages, and seeks recovery, in an amount to be proven at trial, adequate to compensate for Intel's infringement, but in no event less than a reasonable royalty with interest and costs. Intel's infringement of BiTMICRO's rights under the '243 Patent will continue to damage BiTMICRO,

⁶ Available at <https://newsroom.intel.com/wp-content/uploads/sites/11/2019/08/Intel-Lakefield-HotChips-presentation.pdf>.

causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

SECOND COUNT

(INFRINGEMENT OF U.S. PATENT NO. 9,135,190)

90. BiTMICRO incorporates by reference the allegations set forth in paragraphs 1-89 as though fully set forth herein.

91. On information and belief, Intel has directly infringed and continues to directly infringe one or more claims of the '190 Patent, including at least claim 59 of the '190 Patent, in the state of Texas, in this judicial district, and elsewhere in the United States by, among other things, making, using, selling, offering for sale, and/or importing into the United States products that embody one or more of the inventions claimed in the '190 Patent, including but not limited to the above-identified SSDs with SLC caching and Optane Persistent Memory products, and all reasonably similar products ("the '190 Accused Products"), in violation of 35 U.S.C. § 271(a).


92. As an example, the 660p Series SSDs include "a memory controller." Specifically, the 660p Series SSDs includes a memory controller for handling read and write operations on the NAND memory cells on the drive. "Intel® SSD 660p, built using Intel® QLC NAND technology, enables 4-bits of storage per cell. The QLC NAND allows 33% more storage in the same area compared to TLC NAND. Compared to TLC NAND cell, the QLC NAND cell has a higher write and read latency. To mitigate the impact of higher latency of the QLC NAND in Intel® SSD 660p, Intel has designed and implemented an innovative SLC cache, which is variable based on available

unused capacity on the SSD, that delivers faster sequential and random data writes.” (See Intel Solid State Drive 660p Series Evaluation Guide, at 6).⁷

93. The 660p Series SSDs include “an interface controller coupled to a memory device interface and an input/output (IO) device interface.” For example, the 660p Series SSDs include an SMI 2263 controller chip that serves as an interface controller:

Specifications			
Product	Intel SSD 660P 512GB	Intel SSD 660P 1TB	Intel SSD 660P 2TB
Pricing	\$64	\$97	\$204
Capacity (User / Raw)	512GB / 512GB	1024GB / 1024GB	2048GB / 2048GB
Form Factor	M.2 2280 (single-sided)	M.2 2280 (single-sided)	M.2 2280 (single-sided)
Interface / Protocol	PCIe 3.0 x4 / NVMe 1.3	PCIe 3.0 x4 / NVMe 1.3	PCIe 3.0 x4 / NVMe 1.3
Controller	SMI 2263	SMI 2263	SMI 2263
DRAM	NANYA DDR3L	NANYA DDR3L	NANYA DDR3L
Memory	IMFT 64L 3D QLC	IMFT 64L 3D QLC	IMFT 64L 3D QLC
Sequential Read	1,500 MB/s	1,800 MB/s	1,800 MB/s
Sequential Write	1,000 MB/s	1,800 MB/s	1,800 MB/s
Random Read	90,000 IOPS	150,000 IOPS	220,000 IOPS
Random Write	220,000 IOPS	220,000 IOPS	220,000 IOPS

⁷ Available at https://www.intel.com/content/dam/support/us/en/documents/memory-and-storage/consumer-ssds/Intel_SSD_660p_EvaluationGuide337971.pdf.



(“Intel SSD 660p 2TB Review: a QLC Bargain (Update)”)⁸.

⁸ Available at <https://www.tomshardware.com/reviews/intel-ssd-660p-qlc-nvme.5719.html>.

94. The SMI 2263 controller chip is connected to at least one memory bus, which serves as a memory device interface. The controller chip is also connected to a PCIe interface, which serves as an input/output interface with a computing device:

Specifications			
Product	Intel SSD 660P 512GB	Intel SSD 660P 1TB	Intel SSD 660P 2TB
Pricing	\$64	\$97	\$204
Capacity (User / Raw)	512GB / 512GB	1024GB / 1024GB	2048GB / 2048GB
Form Factor	M.2 2280 (single-sided)	M.2 2280 (single-sided)	M.2 2280 (single-sided)
Interface / Protocol	PCIe 3.0 x4 / NVMe 1.3	PCIe 3.0 x4 / NVMe 1.3	PCIe 3.0 x4 / NVMe 1.3
Controller	SMI 2263	SMI 2263	SMI 2263
DRAM	NANYA DDR3L	NANYA DDR3L	NANYA DDR3L
Memory	IMFT 64L 3D QLC	IMFT 64L 3D QLC	IMFT 64L 3D QLC
Sequential Read	1,500 MB/s	1,800 MB/s	1,800 MB/s
Sequential Write	1,000 MB/s	1,800 MB/s	1,800 MB/s
Random Read	90,000 IOPS	150,000 IOPS	220,000 IOPS
Random Write	220,000 IOPS	220,000 IOPS	220,000 IOPS

(“Intel SSD 660p 2TB Review: a QLC Bargain (Update)”⁹).

95. The 660p Series SSDs include “a memory store.” For example, the 660 Series SSDs include 64-layer QLC NAND memory, which constitutes a memory store. As advertised by Intel, “Intel® SSD 660p, built using Intel® QLC NAND technology, enables 4-bits of storage per cell.” (See Intel Solid State Drive 660p Series Evaluation Guide, at 6).

⁹ Available at <https://www.tomshardware.com/reviews/intel-ssd-660p-qlc-nvme.5719.html>.

96. In the 660p Series SSDs, “the memory device interface is directly coupled to the memory store.” For example, the memory bus connected to the controller chip is directly coupled to the QLC NAND memory store, thereby enabling the controller to handle read and write operations to the memory store.

97. In the 660 Series SSDs, “said interface controller [is] disposed to perform a memory transaction by addressing a first memory location in the memory store.” For example, in the 660 Series SSDs, a portion of the QLC NAND memory store is reserved to act as an SLC cache. Data can be written to the SLC cache at a faster rate than to other portions of the QLC NAND memory store. As advertised by Intel, “[t]o mitigate the impact of higher latency of the QLC NAND in Intel® SSD 660p, Intel has designed and implemented an innovative SLC cache, which is variable based on available unused capacity on the SSD, that delivers faster sequential and random data writes. . . . Data writes from the host will be directed through the variable size, high speed SLC cache and offer faster data writes through the full span of the SLC cache. When the SLC cache is reaching its full capacity, the drive firmware moves the contents of the SLC cache into the available QLC cells.” (See Intel Solid State Drive 660p Series Evaluation Guide, at 6).¹⁰ For write operations to the SLC cache, the controller chip performs a memory transaction by addressing a first memory location within the SLC cache.

98. In the 660 Series SSDs, “said first memory location and a second memory location [are] respectively associated with a first device profile and a second device profile.” For example, a second memory location within the QLC NAND memory store is a QLC cell that is not a part of the SLC cache. The first memory location and second memory location are associated with a first

¹⁰ Available at https://www.intel.com/content/dam/support/us/en/documents/memory-and-storage/consumer-ssds/Intel_SSD_660p_EvaluationGuide337971.pdf.

device profile and a second device profile, respectively, that define how data is to be stored in those locations.

99. In the 660 Series SSDs, “said first device profile is optimal for a data type subject to the memory transaction, wherein said data type comprises one of a random data type or a sequential data type.” For example, as advertised by Intel, the SLC cache “delivers faster sequential and random data writes.” (See Intel Solid State Drive 660p Series Evaluation Guide, at 6).¹¹

100. In the 660 Series SSDs, “said interface controller identifies command details for causing the memory transaction to be performed, wherein said command details comprising the first memory device.” For example, for memory transactions with the first memory device such as writing data to the SLC cache, the controller chip identifies command details for causing the memory transaction to be performed in the first memory device.

101. In the 660 Series SSDs, “said device profile represent[s] a first set of attributes of said first memory location, and said second device profile represent[s] a second set of attributes of said second memory location, and a difference exists between said first and second device profiles.” For example, the first device profile (*e.g.*, the profile for the SLC cache) represents a first set of attributes (*e.g.*, data size, memory protocol, device type) associated with the first memory location (*e.g.*, a location within the SLC cache). The second device profile (*e.g.*, the profile for the non-cache portion of the QLC NAND memory store) represents a second set of attributes (*e.g.*, data size, memory protocol, device type) associated with the second memory location (*e.g.*, a location within the non-cache portion of the QLC NAND memory store). The first and second device profiles are different because the SLC cache has attributes associated with a write protocol of one

¹¹ Available at https://www.intel.com/content/dam/support/us/en/documents/memory-and-storage/consumer-ssds/Intel_SSD_660p_EvaluationGuide337971.pdf.

bit per cell, whereas the non-cache portion of the QLC NAND memory store has attributes associated with a write protocol of four bits per cell.

102. In the 660 Series SSDs, “said interface controller obtain[s] the first set of attributes after identifying the command details; and said addressing of said first memory location includes using said attributes from said first device profile.” For example, after identifying command details specifying that data is to be written to the SLC cache, the controller chip obtains the first set of attributes (*e.g.*, data size, memory protocol, device type) associated with the memory location in the SLC cache.

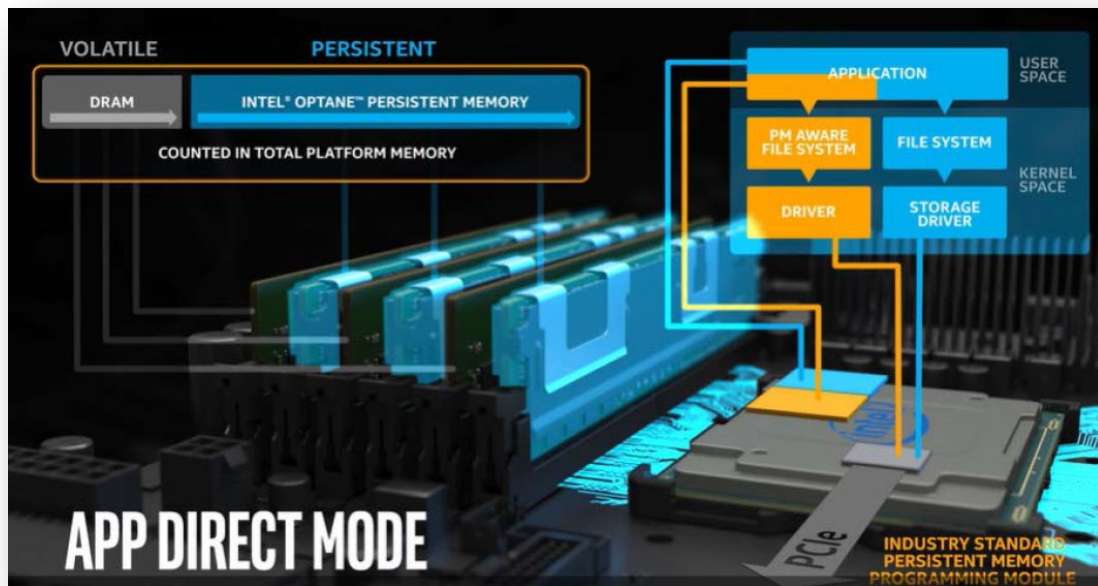
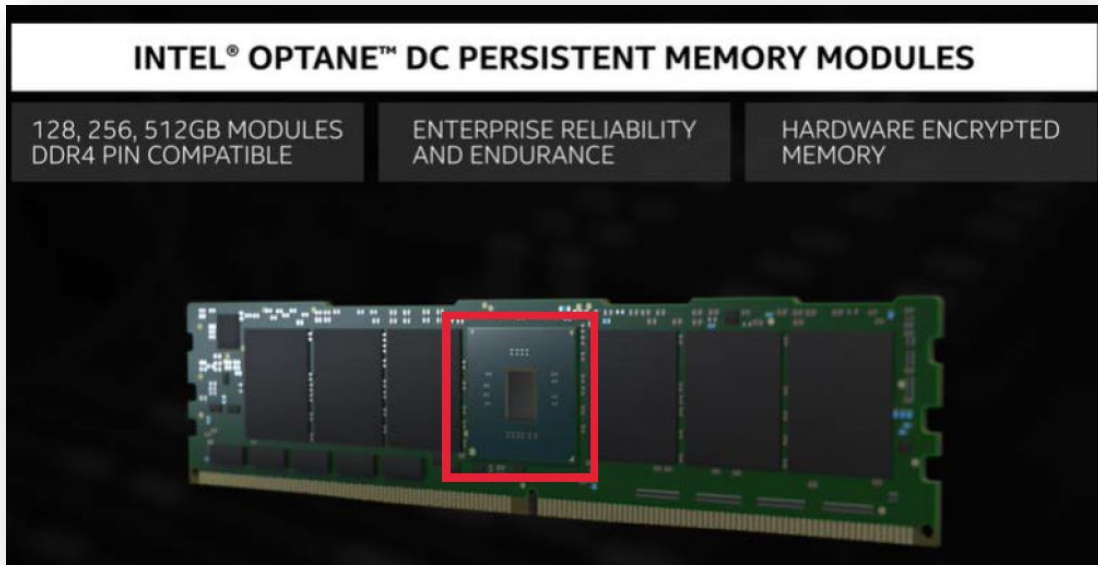
103. In the 660 Series SSDs, “said addressing of said first memory location includes selecting a transfer size for the memory transaction, wherein the transfer size is a function of a data size of the memory transaction and the first set of attributes.” As advertised by Intel, “[w]hen the SLC cache is reaching its full capacity, the drive firmware moves the contents of the SLC cache into the available QLC cells. The firmware also has an intelligent algorithm to opportunistically move the data from SLC cache to QLC cells during the drive idles. This feature allows the user to experience a smoother and refreshed experience from the SSD. The variable SLC adjusts in size bi-directionally based on the available capacity of drive.” (*See Intel Solid State Drive 660p Series Evaluation Guide*, at 6).¹² Thus, for example, after receiving a write transaction command, the controller chip obtains the attributes of the SLC cache profile and determines whether there is sufficient capacity within the SLC cache to write the data to the cache. The controller will select a transfer size to the SLC cache based on an analysis of the size of the data to be written, the remaining memory capacity within the SLC cache, and the attributes of the SLC cache profile.

¹² Available at https://www.intel.com/content/dam/support/us/en/documents/memory-and-storage/consumer-ssds/Intel_SSD_660p_EvaluationGuide337971.pdf.

104. As another example, Intel’s Optane Persistent Memory product includes “a memory controller.” Specifically, the Optane Persistent Memory product includes a memory controller for handling read and write operations on the DRAM memory and persistent memory modules in the product. The persistent memory modules have properties similar to flash memory. As advertised by Intel, for example, “NVDIMM uses a dual in-line memory module (DIMM) package compatible with a standard DIMM slot, and communicates through a standard double data rate (DDR) bus. Considering the fact that it is nonvolatile and compatible with a traditional dynamic random-access memory (DRAM) interface, it is also called persistent memory (PMEM). . . . NVDIMM-P combines the features of DRAM and flash, supporting both block addressing and traditional DRAM-like byte addressing, with a possible terabyte capacity, like NAND flash. . . . Intel Corporation released Intel Optane DC persistent memory in April 2019. It can be seen as an instance of an NVDIMM-P implementation.” (*See Enabling Persistent Memory in the Storage Performance Development Kit (SPDK)*).¹³

105. The Optane Persistent Memory product includes “an interface controller coupled to a memory device interface and an input/output (IO) device interface.” For example, the Optane Persistent Memory product includes a controller that is connected to at least one memory bus, which serves as an interface with a memory device such as DRAM memory or persistent memory. The controller is also connected to an input/output interface which is designed to connect to a computing device via a standard DIMM slot.

¹³ Available at <https://www.intel.com/content/www/us/en/developer/articles/technical/enabling-persistent-memory-in-the-storage-performance-development-kit-spdk.html>.



(Intel Optane DC Persistent Memory Technical Video).¹⁴

¹⁴ Available at <https://www.intel.com/content/www/us/en/design/products-and-solutions/memory-and-storage/pmem/optane-dc-persistent-technical-video.html>.

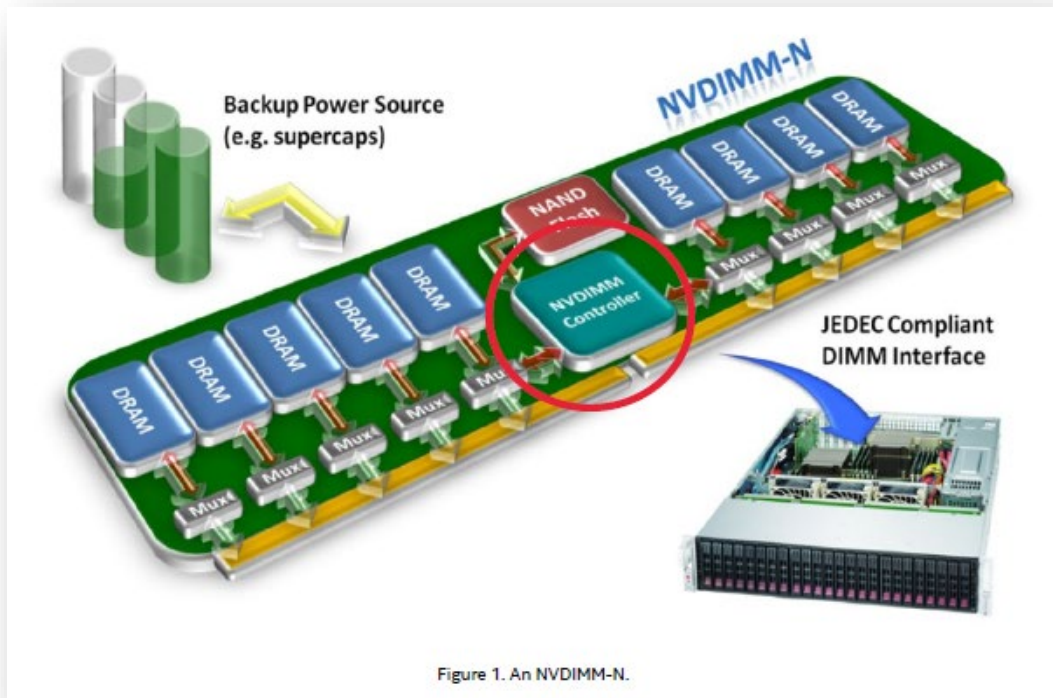
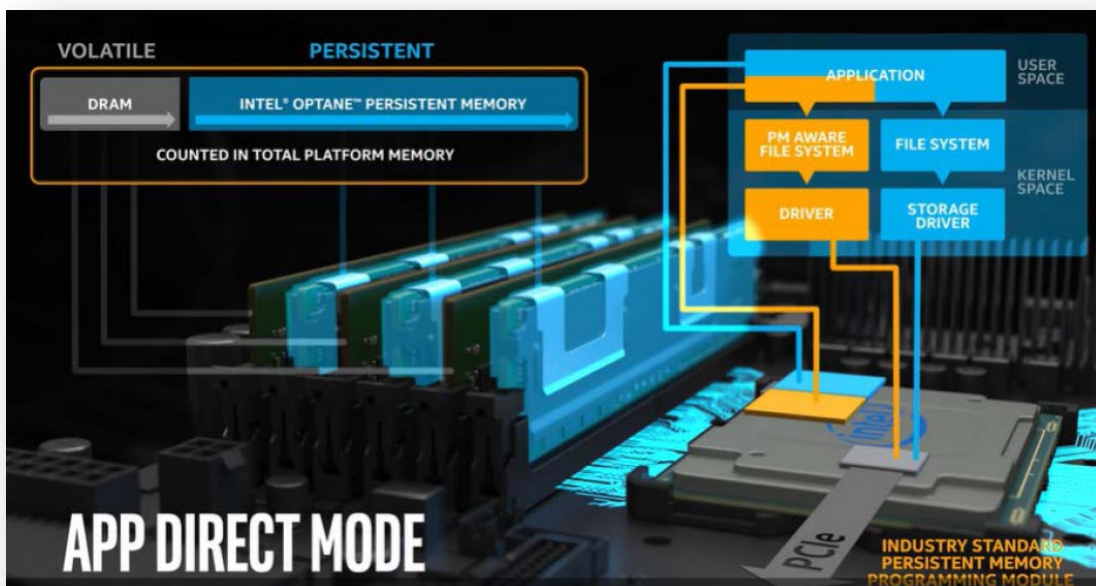
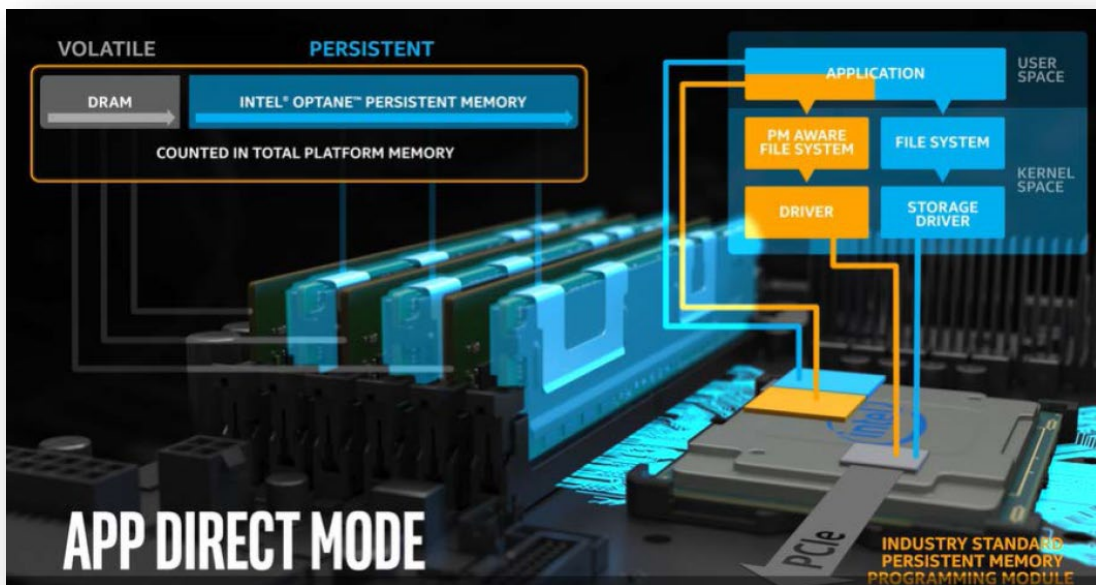


Figure 1. An NVDIMM-N.

(“Enabling Persistent Memory in the Storage Performance Development Kit (SPDK)”)¹⁵

106. The Optane Persistent Memory product includes “a memory store.” For example, the Optane Persistent Memory product includes both DRAM memory and persistent memory modules, which separately or in combination are a memory store:

¹⁵ Available at <https://www.intel.com/content/www/us/en/developer/articles/technical/enabling-persistent-memory-in-the-storage-performance-development-kit-spdk.html>.



(Intel Optane DC Persistent Memory Technical Video).¹⁶

¹⁶ Available at <https://www.intel.com/content/www/us/en/design/products-and-solutions/memory-and-storage/pmem/optane-dc-persistent-technical-video.html>.

107. In the Optane Persistent Memory product, “the memory device interface is directly coupled to the memory store.” For example, in the Intel Optane Persistent Memory product, the memory bus connected to the controller is directly coupled to the DRAM memory and persistent memory, thereby enabling the controller to handle read and write operations to the DRAM memory and persistent memory.

108. In the Optane Persistent Memory product, “said interface controller [is] disposed to perform a memory transaction by addressing a first memory location in the memory store.” For example, a first memory location in the memory store is a location within the DRAM memory. The controller performs a memory transaction (*e.g.*, a write operation) by addressing a first memory location (*e.g.*, a location within the DRAM memory) in the memory store.

109. In the Optane Persistent Memory product, “said first memory location and a second memory location [are] respectively associated with a first device profile and a second device profile.” For example, in addition to the first memory location within the DRAM memory, the Intel Optane Persistent Memory product also includes a second memory location within the persistent memory. The first memory location (*e.g.*, a location within the DRAM memory) is associated with a first device profile (*e.g.*, the profile for the DRAM memory). Similarly, a second memory location (*e.g.*, a location within the persistent memory) is associated with a second device profile (*e.g.*, the profile for the persistent memory). Separate profiles exist for the DRAM memory and the persistent memory in the Intel Optane Persistent Memory product. This is due to the fact that the Optane Persistent Memory product uses a different write protocol when writing to the DRAM memory than it does when writing to the persistent memory.

110. For example, as advertised by Intel, as an NVDIMM implementation the Optane Persistent Memory product “combines the features of DRAM and flash, supporting both block

addressing and traditional DRAM-like byte addressing, with a possible terabyte capacity, like NAND flash. . . .” (“Enabling Persistent Memory in the Storage Performance Development Kit (SPDK)”).¹⁷ As also advertised by Intel, “Intel Optane PMem 200 series has multiple operating modes: Memory Mode delivers large memory capacity without application changes and with performance close to that of DRAM, depending on the workload. In Memory Mode, the CPU memory controller sees all of the Intel Optane PMem 200 series as volatile system memory (without persistence). The CPU uses DRAM as a fast cache to the Intel Optane PMem. . . . App Direct Mode enables large memory capacity and data persistence for software to access DRAM and persistent memory as two separate pools of memory. . . . In App Direct Mode, data is encrypted using a key stored on the module in a security metadata region, which can only be accessed by the Intel Optane PMem 200 series controller.” (Data Center Intel Optane Persistent Memory 200 Series Product Brief).¹⁸ Different device profiles are therefore necessary in order for the Intel Optane Persistent Memory to be able to know which protocol to use with the DRAM and persistent memory devices.

111. In the Optane Persistent Memory product, “said first device profile is optimal for a data type subject to the memory transaction, wherein said data type comprises one of a random data type or a sequential data type.” For example, as advertised by Intel, as an NVDIMM implementation the Optane Persistent Memory product includes “memory that is nonvolatile and

¹⁷ Available at <https://www.intel.com/content/www/us/en/developer/articles/technical/enabling-persistent-memory-in-the-storage-performance-development-kit-spdk.html>.

¹⁸ Available at <https://www.intel.com/content/dam/www/public/us/en/documents/product-briefs/optane-persistent-memory-200-series-brief.pdf>.

allows for random access.” (“Enabling Persistent Memory in the Storage Performance Development Kit (SPDK)”)¹⁹.

112. In the Optane Persistent Memory product, “said interface controller identifies command details for causing the memory transaction to be performed, wherein said command details comprising the first memory device.” For example, for memory transactions with the first memory device (*e.g.*, the DRAM memory), the controller identifies command details for causing the memory transaction to be performed in the first memory device.

113. In the Optane Persistent Memory product, “said device profile represent[s] a first set of attributes of said first memory location, and said second device profile represent[s] a second set of attributes of said second memory location, and a difference exists between said first and second device profiles.” For example, the first device profile (*e.g.*, the profile for the DRAM memory) represents a first set of attributes (*e.g.*, data size, memory protocol, device type, operating mode) associated with the first memory location (*e.g.*, a location within the DRAM memory). The second device profile (*e.g.*, the profile for the persistent memory) represents a second set of attributes (*e.g.*, data size, memory protocol, device type, operating mode) associated with the second memory location (*e.g.*, a location within the persistent memory). The first and second device profiles are different because DRAM and persistent memory are different types of memory with different parameters for reading and writing data.

114. In the Optane Persistent Memory product, “said interface controller obtain[s] the first set of attributes after identifying the command details; and said addressing of said first memory location includes using said attributes from said first device profile.” For example, after identifying

¹⁹ Available at <https://www.intel.com/content/www/us/en/developer/articles/technical/enabling-persistent-memory-in-the-storage-performance-development-kit-spdk.html>.

the command details specifying that data is to be written to the DRAM memory, the controller obtains the first set of attributes (*e.g.*, data size, memory protocol, device type) associated with the memory location in the DRAM memory.

115. In the Optane Persistent Memory product, “said addressing of said first memory location includes selecting a transfer size for the memory transaction, wherein the transfer size is a function of a data size of the memory transaction and the first set of attributes.” For example, after receiving a write transaction command, the controller obtains the mode attributes and/or available data attributes associated with the first and second device profiles. As advertised by Intel, “[i]n Memory Mode, the DRAM acts as a cache for the most frequently accessed data, while the Intel Optane DC persistent memory provides large memory capacity. . . . In App Direct Mode, applications and the Operating System are explicitly aware there are two types of direct load/store memory in the platform and can direct which type of data read or write is suitable for DRAM or Intel Optane DC persistent memory. Operations that require the lowest latency and don’t need permanent data storage can be executed on DRAM, such as database ‘scratch pads’. Data that needs to be made persistent or structures that are very large can be routed to the Intel Optane DC persistent memory.” (Intel Optane DC Persistent Memory Quick Start Guide, Revision 1.1 (June 2020), p. 3).²⁰ The Intel Optane Persistent Memory product can therefore receive a data size of the memory transaction from the write transaction command and select a transfer size to the DRAM memory based on the mode and available data attributes associated with the DRAM memory and persistent memory.

²⁰ Available at <https://www.intel.com/content/dam/support/us/en/documents/memory-and-storage/data-center-persistent-mem/Intel-Optane-DC-Persistent-Memory-Quick-Start-Guide.pdf>.

116. By making, using, offering for sale, and/or selling products in the United States and/or importing products into the United States, including but not limited to the '190 Accused Products, Intel has injured BiTMICRO and is liable to BiTMICRO for directly infringing one or more claims of the '190 Patent, including without limitation claim 59 pursuant to 35 U.S.C. § 271(a).

117. On information and belief, Intel has had knowledge of the '190 Patent since at least September 5, 2018, when BiTMICRO served a subpoena on Intel in connection with the ITC Action. In the ITC Action, BiTMICRO alleged that various solid state computer drives (“SSDs”) and electronic devices that incorporate stacked electronics components sold by Samsung Electronics, SK Hynix, and other electronic device manufacturers infringed four BiTMICRO patents, including the '190 Patent.

118. Intel responded to BiTMICRO’s subpoena in the ITC Action by serving objections and responses to the subpoena, conferring with BiTMICRO’s ITC counsel regarding the subpoena, and filing a motion in the ITC on September 14, 2018 for an extension of time to respond to the subpoena.

119. On information and belief, through its participation in the ITC Action as described above, Intel had knowledge of the '190 Patent and its relevance to the '190 Accused Products. Despite this knowledge, Intel has continued to directly infringe one or more claims of the '190 Patent as described above. Thus, on information and belief, Intel’s infringement of the '190 Patent has been willful.

120. On information and belief, Intel is also inducing and/or has induced infringement of one or more claims of the '190 Patent, including at least claim 59, as a result of, amongst other activities, instructing, encouraging, and directing its customers on the use of the '190 Accused

Products in an infringing manner in violation of 35 U.S.C. § 271(b). Through its website, instructional guides, and manuals, Intel provides its customers with detailed explanations, instructions, and information on how to use and implement the '190 Accused Products which demonstrate active steps taken to encourage direct infringement. (*See, e.g.*, Enabling Persistent Memory in the Storage Performance Development Kit (SPDK)²¹; Intel Solid State Drive 660p Series Evaluation Guide²²). On information and belief, Intel has had knowledge of the '190 Patent since at least September 5, 2018 as set forth above, and also as of the filing and/or service of this Complaint. Despite this knowledge, Intel has continued to engage in activities to encourage and assist its customers in the use of the '190 Accused Products. Thus, on information and belief, Intel (1) had actual knowledge of the patent; (2) knowingly induced its customers to infringe the patent; and (3) had specific intent to induce the patent infringement.

121. On information and belief, by using the '190 Accused Products as encouraged and assisted by Intel, Intel's customers have directly infringed and continue to directly infringe one or more claims of the '190 Patent, including at least claim 59. On information and belief, Intel knew or was willfully blind to the fact that its actions would induce its customers' direct infringement of the '190 Patent.

122. Intel's infringement of the '190 Patent has been and continues to be deliberate and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284-285.

²¹ Available at <https://www.intel.com/content/www/us/en/developer/articles/technical/enabling-persistent-memory-in-the-storage-performance-development-kit-spdk.html>.

²² Available at https://www.intel.com/content/dam/support/us/en/documents/memory-and-storage/consumer-ssds/Intel_SSD_660p_EvaluationGuide337971.pdf.

123. On information and belief, Intel will continue to infringe the '190 Patent unless enjoined by this Court.

124. As a result of Intel's infringement of the '190 Patent, BiTMICRO has suffered monetary damages, and seeks recovery, in an amount to be proven at trial, adequate to compensate for Intel's infringement, but in no event less than a reasonable royalty with interest and costs. Intel's infringement of BiTMICRO's rights under the '190 Patent will continue to damage BiTMICRO, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

THIRD COUNT

(INFRINGEMENT OF U.S. PATENT NO. 8,010,740)


125. BiTMICRO incorporates by reference the allegations set forth in paragraphs 1-124 as though fully set forth herein.

126. On information and belief, Intel has directly infringed and continues to directly infringe one or more claims of the '740 Patent, including at least claim 1 of the '740 Patent, in the state of Texas, in this judicial district, and elsewhere in the United States by, among other things, making, using, selling, offering for sale, and/or importing into the United States products that embody one or more of the inventions claimed in the '740 Patent, including but not limited to the above-identified NVMe and Optane SSDs and Optane memory products, and all reasonably similar products ("the '740 Accused Products"), in violation of 35 U.S.C. § 271(a).

127. As an example, Intel's Optane SSDs and Optane Persistent Memory (PMem) modules include "a mapping table for optimizing memory operations performed by an electronic storage device in response to receiving an I/O transaction request initiated by a host." Specifically, the Optane SSDs and Pmem modules include a mapping table, which Intel calls an "address

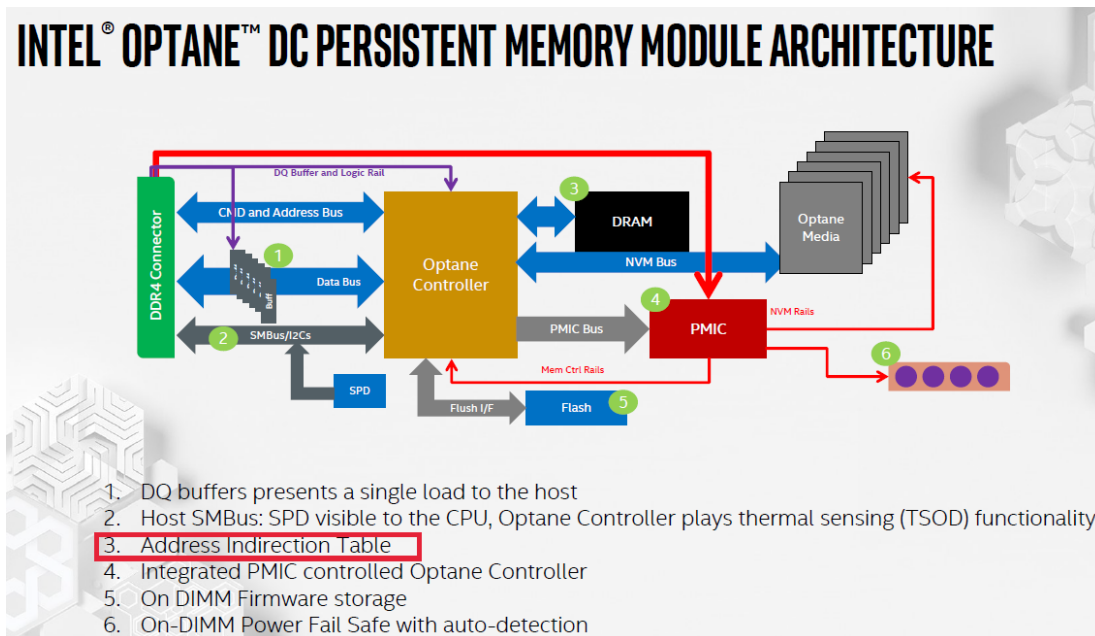
indirection table,” located in DRAM or other memory to map logical block addresses (LBAs) to physical block addresses (PBAs) on the memory devices within the SSD or Optane MemoryPMem products.

128. Intel describes the address indirection table in its products, for example, as follows:



(Intel presentation “Completing the Memory and Storage Hierarchy,” p. 59).²³

²³ Available at https://newsroom.intel.com/wp-content/uploads/sites/11/2019/09/Intel-2019_MemoryStorageDay_FrankHady_MohamedArafa_PranavKalavade_final.pdf



(Intel presentation “Optane Data Center Persistent Memory,” p. 7).²⁴ See also “An Empirical Guide to the Behavior and Use of Scalable Persistent Memory,”²⁵ at 170 (“The Optane DIMM is the first scalable, commercially available NVDIMM. . . . Memory accesses to the NVDIMM (Figure 1(b)) arrive first at the on-DIMM controller . . . which coordinates access to the Optane media. Similar to SSDs, the Optane DIMM performs an internal address translation for wear-leveling and bad-block management, and maintains an address indirection table (AIT) for this translation.”); “The Intel SSD DC S3700: Intel’s 3rd Generation Controller Analyzed”²⁶ (“All controllers have to map logical block addresses to physical locations in NAND. This map is stored on the NAND itself (and wear leveled so it actually moves locations), but it’s cached in DRAM for fast access. Intel calls this map its indirection table.”).

²⁴ Available at <https://newsroom.intel.com/wp-content/uploads/sites/11/2019/08/Intel-Optane-HotChips-presentation.pdf>

²⁵ Available at <https://www.usenix.org/system/files/fast20-yang.pdf>

²⁶ Available at <https://www.anandtech.com/show/6432/the-intel-ssd-dc-s3700-intels-3rd-generation-controller-analyzed/2>

129. Information in the indirection table is used by a controller in the Optane SSD or Pmem module to optimize memory operations, by for example allowing the controller to perform interleaving to optimally distribute data across multiple non-volatile memory channels, buses, devices, dies, planes, etc., thereby improving the speed and efficiency of storing and accessing the data. *See, e.g.*, “Quick Start Guide Part 1: Persistent Memory Provisioning Introduction”²⁷ (“A region is a group of one or more Pmem modules, also known as an interleaved set. The Regions are created in interleaved or non-interleaved configurations. Interleaving is a technique that makes multiple Pmem devices appear as a single logical virtual address space. It allows spreading adjacent virtual addresses within a page across multiple memory devices. This hardware-level parallelism increases the available bandwidth from the devices.”); “Exploring Performance Characteristics of the Optane 3D Xpoint Storage Technology,”²⁸ at 4.2 (“To provide high bandwidth, NAND SSD manufacturers normally use an interleaving technique to access pages from multiple channels in parallel. The Optane SSD inherits a similar channel-based interleaving technique to that of the NAND SSD, but with advanced implementation.”); “Platform Storage Performance With 3D Xpoint Technology,” *Proceedings of the IEEE*, Vol. 105, No. 9 (Sept. 2017) at 1823 (“The Optane SSD has been optimized throughput to deliver low IO latency. . . . Within the SSD itself, the controller spreads single 4-kB Ios across multiple 3D Xpoint memory channels to harness the throughput of multiple memory dies to deliver low latency for a single IO.”).

130. The indirection table of Intel’s Optane SSDs and Optane Pmem modules also includes “a set of logical fields, including a first logical field and a second logical field, and said

²⁷ Available at <https://www.intel.com/content/www/us/en/developer/articles/technical/qsg-intro-to-provisioning-pmem.html>

²⁸ Available at <https://dl.acm.org/doi/fullHtml/10.1145/3372783#Bib0018>

logical fields respectively disposed for representing a plurality of LBA sets, including said first logical field disposed for representing a first LBA set and said second logical field disposed for representing a second LBA set, said first and second LBA sets each representing a set of consecutive LBAs.” For example, the indirection table includes a set of logical fields that represent sets of consecutive LBAs. *See, e.g.*, “The Intel SSD DC S3700: Intel’s 3rd Generation Controller Analyzed”²⁹ (“All controllers have to map logical block addresses to physical locations in NAND. This map is stored on the NAND itself (and wear leveled so it actually moves locations), but it’s cached in DRAM for fast access. Intel calls this map its indirection table.”); “An Empirical Guide to the Behavior and Use of Scalable Persistent Memory,”³⁰ at 170 (“The Optane DIMM is the first scalable, commercially available NVDIMM. . . . Memory accesses to the NVDIMM (Figure 1(b)) arrive first at the on-DIMM controller . . . which coordinates access to the Optane media. Similar to SSDs, the Optane DIMM performs an internal address translation for wear-leveling and bad-block management, and maintains an address indirection table (AIT) for this translation.”); “Towards an Unwritten Contract of Intel Optane SSD,”³¹ at 5 (“Finally, we study the mapping policy (LBA→PBA) in Optane SSD by comparing three workload variations. The first is the same workload as for the interleaving experiments in Figure 3: blocks are first written in logical address order and then read back in that same LBA order (LBA-order write:LBA-order read). The second workload preconditions the working zone with random writes (random write:LBA-order read). The third workload preconditions with random writes, but then reads in the order in which the chunks were written (random write:written-order read). Figure 11 shows the throughput of the three

²⁹ Available at <https://www.anandtech.com/show/6432/the-intel-ssd-dc-s3700-intels-3rd-generation-controller-analyzed/2>

³⁰ Available at <https://www.usenix.org/system/files/fast20-yang.pdf>

³¹ Available at <https://www.usenix.org/system/files/hotstorage19-paper-wu-kan.pdf>

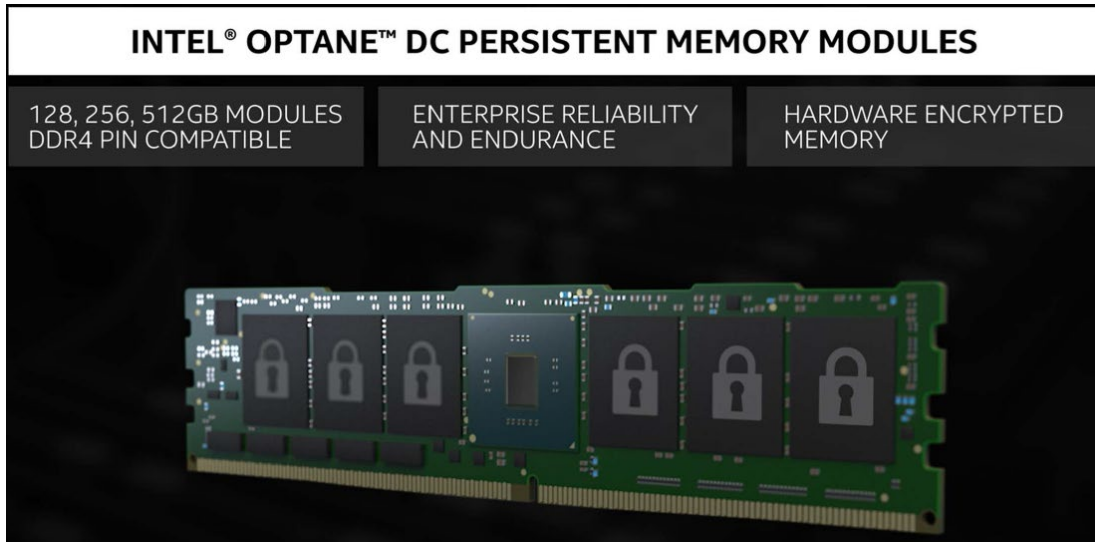
workloads. . . . Optane SSD behaves quite differently; no matter how we precondition the device, the pattern occurs when reading according to LBA. Hence, Optane SSD likely adopts LBA-based mapping.”); “Quick Start Guide Part 1: Persistent Memory Provisioning Introduction”³² (“A region is a group of one or more pmem modules, also known as an interleaved set. The Regions are created in interleaved or non-interleaved configurations. Interleaving is a technique that makes multiple Pmem devices appear as a single logical virtual address space. It allows spreading adjacent virtual addresses within a page across multiple memory devices. This hardware-level parallelism increases the available bandwidth from the devices.”).

131. The indirection table in Intel’s Optane SSDs and Optane Pmem modules also includes “a set of PBA fields, including a first PBA field and a second PBA field, said set of PBA fields respectively disposed for representing a set of PBAs, including a first PBA disposed for representing a first set of access parameters and a second PBA disposed for representing a second set of access parameters, said PBAs each associated with a physical memory location in a memory store, said set of logical fields and said set of PBA fields disposed to associate said first and second LBA sets with said first and second PBAs.” For example, the first and second PBAs can be associated with different non-volatile memory channels, buses, devices, dies, planes, etc. within the SSD or Pmem module. And in the case where multiple Pmem modules are used together in a system, the first and second PBAs can also be associated with different non-volatile memory channels, buses, devices, dies, planes, etc. across multiple Pmem modules. *See, e.g.*, “The Intel Optane SSD DC P4800X (375GB) Review: Testing 3D Xpoint Performance”³³ (“The first capacity

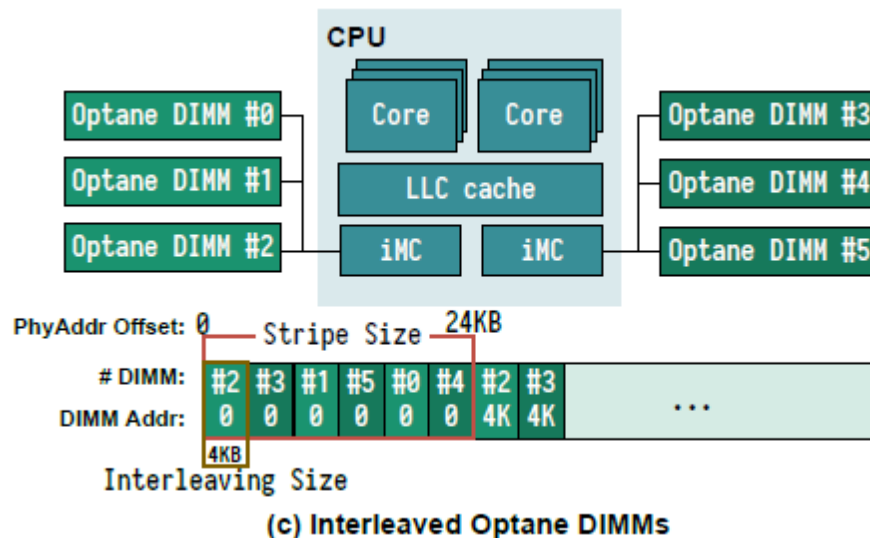
³² Available at <https://www.intel.com/content/www/us/en/developer/articles/technical/qsg-intro-to-provisioning-pmem.html>

³³ Available at <https://www.anandtech.com/show/11209/intel-optane-ssd-dc-p4800x-review-a-deep-dive-into-3d-xpoint-enterprise-performance>

of the Optane SSD DC P4800X to ship and the model we’ve tested here offers a usable capacity of 375GB from a total of 28 3D Xpoint memory dies (four per channel) for a raw capacity of 448GB.”); *See also* “Intel Optane DC Persistent Memory Technical Video,”³⁴ at 2:50:



See also “An Empirical Guide to the Behavior and Use of Scalable Persistent Memory,”³⁵ at 170:



³⁴ Available at <https://www.intel.com/content/www/us/en/design/products-and-solutions/memory-and-storage/pmem/optane-dc-persistent-technical-video.html>

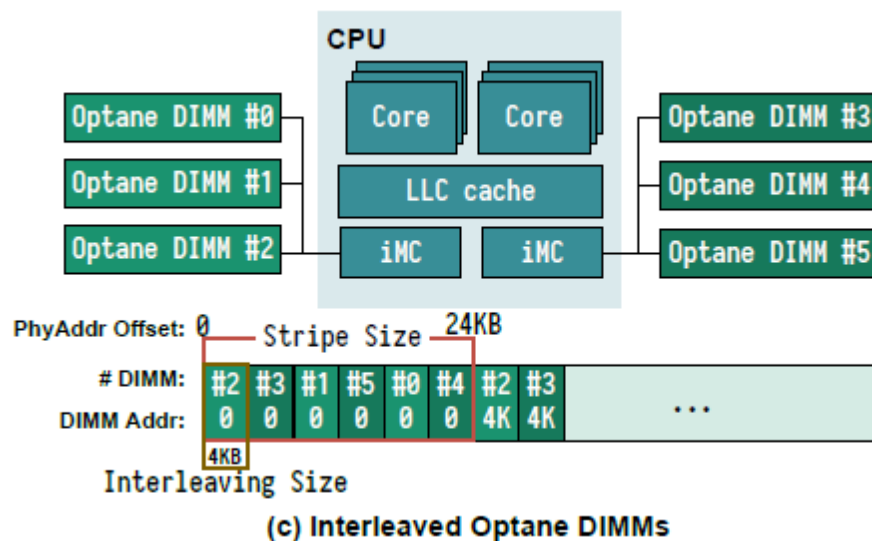
³⁵ Available at <https://www.usenix.org/system/files/fast20-yang.pdf>

The first and second PBA fields represent these first and second PBAs, respectively, in the indirection table. Each of the PBAs has a set of access parameters defined in the indirection table, such as identifying information that allows the PBA to be associated with one or more LBAs, and the first and second LBA sets are associated with the first and second PBAs, respectively. *See, e.g.*, “Towards an Unwritten Contract of Intel Optane SSD,”³⁶ at 5 (“Finally, we study the mapping policy (LBA→PBA) in Optane SSD by comparing three workload variations. The first is the same workload as for the interleaving experiments in Figure 3: blocks are first written in logical address order and then read back in that same LBA order (LBA-order write:LBA-order read). The second workload preconditions the working zone with random writes (random write:LBA-order read). The third workload preconditions with random writes, but then reads in the order in which the chunks were written (random write:written-order read). Figure 11 shows the throughput of the three workloads. . . . Optane SSD behaves quite differently; no matter how we precondition the device, the pattern occurs when reading according to LBA. Hence, Optane SSD likely adopts LBA-based mapping.”).

132. Also, the indirection table in Intel’s Optane SSDs and Optane Pmem modules, “in response to receiving the I/O transaction request, [] causes the electronic storage device to perform optimized memory operations on memory locations respectively associated with said first PBA and said second PBA, if the I/O transaction request is associated with said first and second LBA sets.” For example, the interleaving function in the Optane SSDs and Pmem modules allows the products to perform optimized memory operations on memory locations associated with the first PBA and second PBA by distributing the reading and writing of data across those memory locations to

³⁶ Available at <https://www.usenix.org/system/files/hotstorage19-paper-wu-kan.pdf>

increase the speed and efficiency of storing and accessing the data. This interleaving requires the use of the LBA-PBA mapping information in the indirection table by the memory controller within the SSD or Pmem product. *See, e.g.*, “Quick Start Guide Part 1: Persistent Memory Provisioning Introduction”³⁷ (“A region is a group of one or more Pmem modules, also known as an interleaved set. The Regions are created in interleaved or non-interleaved configurations. Interleaving is a technique that makes multiple Pmem devices appear as a single logical virtual address space. It allows spreading adjacent virtual addresses within a page across multiple memory devices. This hardware-level parallelism increases the available bandwidth from the devices.”); “An Empirical Guide to the Behavior and Use of Scalable Persistent Memory,”³⁸ at 170:



See also “Exploring Performance Characteristics of the Optane 3D Xpoint Storage Technology,”³⁹ at 4.2 (“To provide high bandwidth, NAND SSD manufacturers normally use an interleaving

³⁷ Available at <https://www.intel.com/content/www/us/en/developer/articles/technical/qsg-intro-to-provisioning-pmem.html>

³⁸ Available at <https://www.usenix.org/system/files/fast20-yang.pdf>

³⁹ Available at <https://dl.acm.org/doi/fullHtml/10.1145/3372783#Bib0018>

technique to access pages from multiple channels in parallel. The Optane SSD inherits a similar channel-based interleaving technique to that of the NAND SSD, but with advanced implementation.”); “Platform Storage Performance With 3D Xpoint Technology,” Proceedings of the IEEE, Vol. 105, No. 9 (Sept. 2017) at 1823 (“The Optane SSD has been optimized throughout to deliver low IO latency. We will show the impact of the SSD’s NVMe interface later in the paper. Within the SSD itself, the controller spreads single 4-kB Ios across multiple 3D Xpoint memory channels to harness the throughput of multiple memory dies to deliver low latency for a single IO.”).

133. By making, using, offering for sale, and/or selling products in the United States and/or importing products into the United States, including but not limited to the ’740 Accused Products, Intel has injured BiTMICRO and is liable to BiTMICRO for directly infringing one or more claims of the ’740 Patent, including without limitation claim 1 pursuant to 35 U.S.C. § 271(a).

134. On information and belief, Intel is also inducing and/or has induced infringement of one or more claims of the ’740 Patent, including at least claim 1, as a result of, amongst other activities, instructing, encouraging, and directing its customers on the use of the ’740 Accused Products in an infringing manner in violation of 35 U.S.C. § 271(b). Through its website, instructional guides, and manuals, Intel provides its customers with detailed explanations, instructions, and information on how to use and implement the ’740 Accused Products which demonstrate active steps taken to encourage direct infringement. (*See, e.g.*, “Quick Start Guide Part

1: Persistent Memory Provisioning Introduction;”⁴⁰ “Intel Optane DC Persistent Memory Technical Video,”⁴¹ Intel presentation “Completing the Memory and Storage Hierarchy”⁴²).

135. On information and belief, Intel has had knowledge of the ’740 Patent at least as of the filing and/or service of this Complaint. Despite this knowledge, Intel has continued to engage in activities to encourage and assist its customers in the use of the ’740 Accused Products. Thus, on information and belief, Intel (1) had actual knowledge of the patent; (2) knowingly induced its customers to infringe the patent; and (3) had specific intent to induce the patent infringement.

136. On information and belief, by using the ’740 Accused Products as encouraged and assisted by Intel, Intel’s customers have directly infringed and continue to directly infringe one or more claims of the ’740 Patent, including at least claim 1. On information and belief, Intel knew or was willfully blind to the fact that its actions would induce its customers’ direct infringement of the ’740 Patent.

137. Intel’s infringement of the ’740 Patent has been and continues to be deliberate and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys’ fees and costs pursuant to 35 U.S.C. §§ 284-285.

138. On information and belief, Intel will continue to infringe the ’740 Patent unless enjoined by this Court.

139. As a result of Intel’s infringement of the ’740 Patent, BiTMICRO has suffered monetary damages, and seeks recovery, in an amount to be proven at trial, adequate to compensate

⁴⁰ Available at <https://www.intel.com/content/www/us/en/developer/articles/technical/qsg-intro-to-provisioning-pmem.html>

⁴¹ Available at <https://www.intel.com/content/www/us/en/design/products-and-solutions/memory-and-storage/pmem/optane-dc-persistent-technical-video.html>

⁴² Available at https://newsroom.intel.com/wp-content/uploads/sites/11/2019/09/Intel-2019_MemoryStorageDay_FrankHady_MohamedArafa_PranavKalavade_final.pdf

for Intel's infringement, but in no event less than a reasonable royalty with interest and costs. Intel's infringement of BiTMICRO's rights under the '740 Patent will continue to damage BiTMICRO, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

FOURTH COUNT

(INFRINGEMENT OF U.S. PATENT NO. 9,858,084)

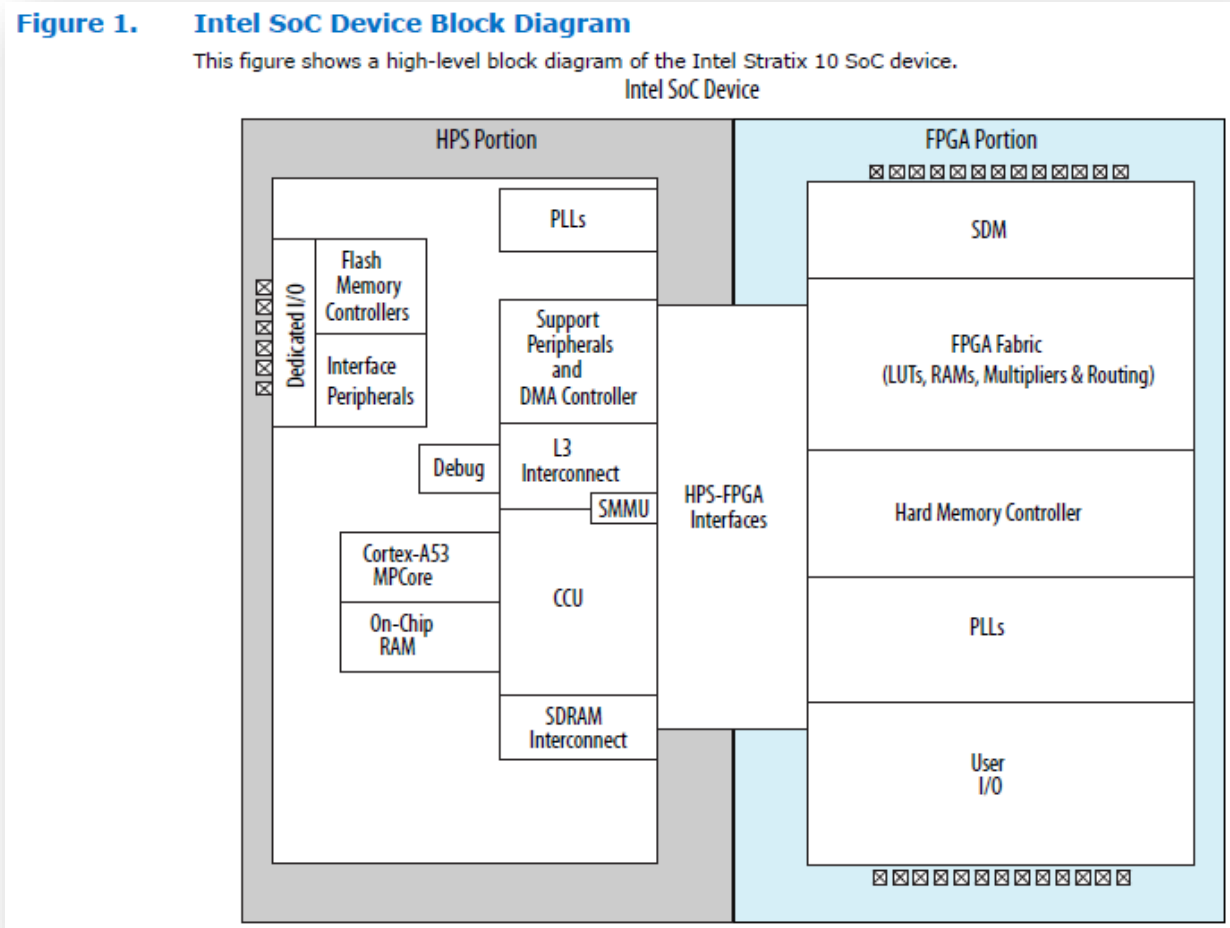
140. BiTMICRO incorporates by reference the allegations set forth in paragraphs 1-139 as though fully set forth herein.

141. Intel has directly infringed and continues to directly infringe one or more claims of the '084 Patent, including at least claim 19 of the '084 Patent, in the state of Texas, in this judicial district, and elsewhere in the United States by, among other things, making, using, selling, offering for sale, and/or importing into the United States products that embody one or more of the inventions claimed in the '084 Patent, including but not limited to the Intel Stratix 10 FPGA products and all reasonably similar products ("the '084 Accused Products"), in violation of 35 U.S.C. § 271(a).

142. As an example, the Intel Stratix 10 products are "an article of manufacture." Specifically, the Stratix 10 is a system-on-a-chip (SoC) that includes multiple processing components. As advertised by Intel, "[t]he Intel Stratix 10 system-on-a-chip (SoC) is composed of two distinct portions; a 64-bit quad core Arm Cortex-A53 hard processor system (HPS) and an FPGA. The HPS architecture integrates a wide set of peripherals that reduce board size and increase performance within a system." (Intel Stratix 10 Hard Processor System Technical Reference Manual, p. 26).⁴³

⁴³ Available at https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/stratix-10/s10_5v4.pdf.

143. Intel has published the following high-level block diagram of the Intel Stratix 10 product:



(Intel Stratix 10 Hard Processor System Technical Reference Manual, p. 26).⁴⁴

144. The Intel Stratix 10 products include “a non-transitory computer-readable medium having stored thereon instructions operable to permit an apparatus to perform a method.” For example, the Intel Stratix 10 includes a Secure Device Manager (“SDM”) that enables the Stratix 10 to perform a system startup method. The SDM “is a triple-redundant processor-based module that

⁴⁴ Available at https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/stratix-10/s10_5v4.pdf.

manages configuration and the security features of the Intel Stratix 10 devices. The SDM is available on all Intel Stratix 10 FPGA and SoC devices.” (Intel Stratix 10 Configuration User Guide, p. 12).⁴⁵

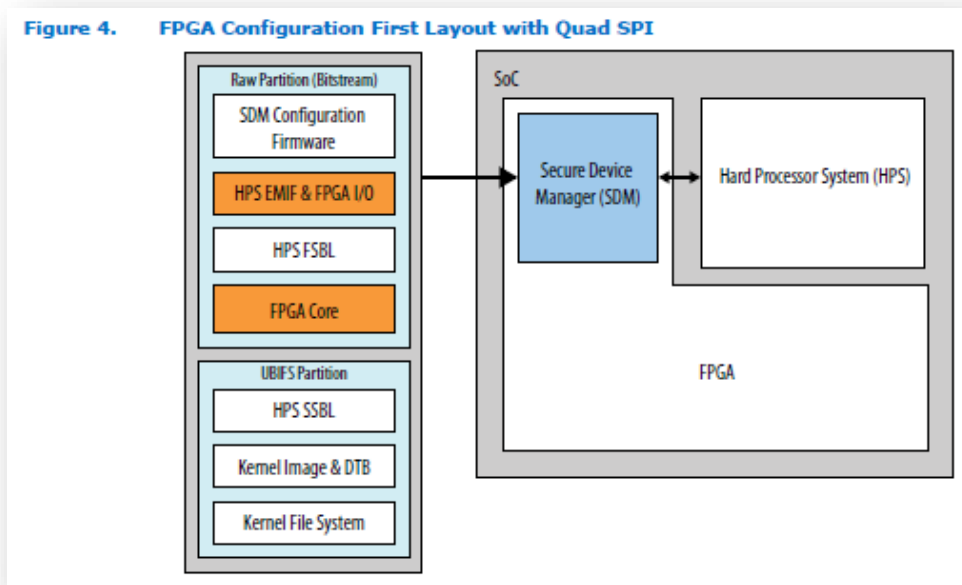
145. The startup method performed by the Intel Stratix 10 includes “releasing components of an embedded system from reset.” For example, in the FPGA First configuration, “the SDM fully configures the FPGA, then configures the HPS SDRAM pins, loads the HPS first stage boot loader (FSBL) and takes the HPS out of reset. In this mode the fabric begins functioning just before the HPS exits reset.” (Intel Stratix 10 Configuration User Guide, p. 18;⁴⁶ *see also* Intel Stratix 10 Hard Processor System Technical Reference Manual, p. 215⁴⁷).

⁴⁵ Available at <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/stratix-10/ug-s10-config.pdf>.

⁴⁶ Available at <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/stratix-10/ug-s10-config.pdf>.

⁴⁷ Available at https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/stratix-10/s10_5v4.pdf.

146. The startup method performed by the Intel Stratix 10 includes “detecting a power-on reset (POR) sequencer descriptor in a nonvolatile memory.” For example, as described by Intel, “[o]nce the Intel Stratix 10 SoC FPGA exits POR, the SDM samples the MSEL [2:0] pins to determine the boot source. Next, the device configures the SDM I/Os according to the selected boot source interface and the SDM retrieves the configuration bitstream through the interface.” (Intel Stratix 10 SoC FPGA Boot User Guide, p. 8).⁴⁸ The configuration bitstream is stored in nonvolatile memory. For example, in the FPGA First mode “the Quad SPI flash connected to the SDM contains all the data required for configuring and booting the system, including the configuration bitstream.” (*Id.* at p. 11).



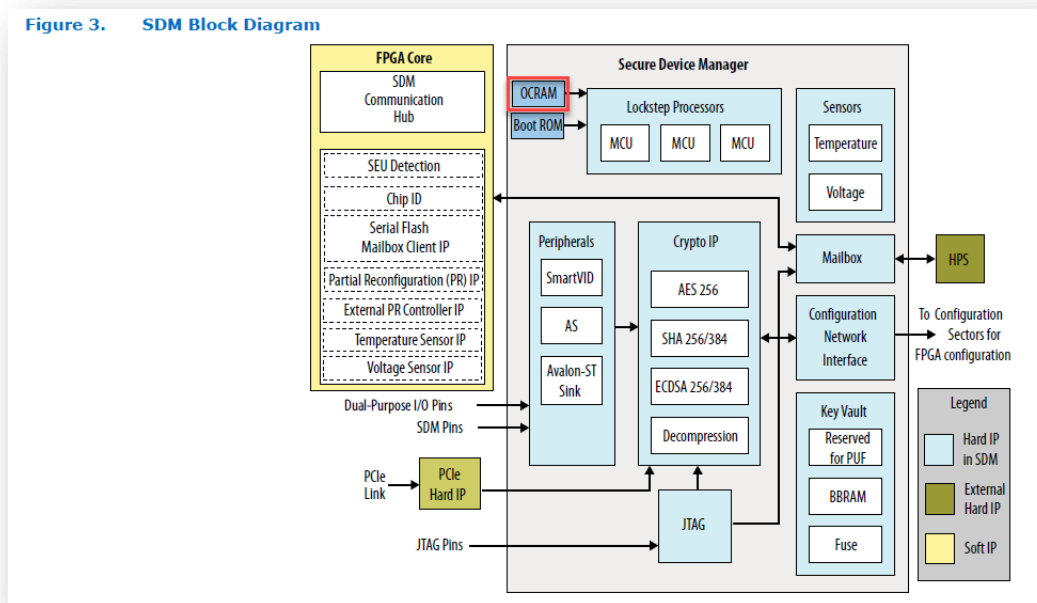
(Intel Stratix 10 SoC FPGA Boot User Guide, p. 12).⁴⁹

⁴⁸ Available at <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/ug/ug-s10-soc-boot.pdf>

⁴⁹ Available at <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/ug/ug-s10-soc-boot.pdf>.

147. In the Stratix 10, “the POR sequencer descriptor comprises information to initialize configuration registers of the embedded system, Direct Memory Access (DMA) descriptors used to fetch other POR sequencer descriptor fragments, and a system firmware.” For example, as shown in the Figures 4 and 5 above, the configuration bitstream includes information to initialize the FPGA I/O and FPGA Core, first-stage bootloader (FSBL) for the HPS, and SDM configuration firmware. This initialization process also includes “initialization of the FPGA fabric, including registers and state machines.” (Intel Stratix 10 Configuration User Guide, p. 20).⁵⁰

148. The startup method performed by the Stratix 10 includes “copying the POR sequencer descriptor from the nonvolatile memory to a random access memory.” For example, the SDM has an on-chip RAM (OCRAM) to which the configuration bitstream is copied:



⁵⁰ Available at <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/stratix-10/ug-s10-config.pdf>.

(Intel Stratix 10 Configuration User Guide, p. 15).⁵¹ Copying of the configuration bitstream from the nonvolatile memory to the OCRAM occurs after the Stratix 10 exits POR, when “the SDM retrieves the configuration bitstream through the interface,” such as from the Quad SPI nonvolatile flash memory. (Intel Stratix 10 SoC FPGA Boot User Guide, p. 8).⁵²

149. The startup method performed by the Stratix 10 includes “verifying an integrity of the POR sequencer descriptor” and “detecting any error in the POR sequencer descriptor.” For example, the SDM performs the function of “[c]onfiguration bitstream authentication: During the configuration state, the SDM authenticates the Intel-generated configuration firmware and configuration bitstream, ensuring that configuration bitstream is from a trusted source. All Intel Stratix 10 support authentication.” (Intel Stratix 10 Configuration User Guide, p. 13).⁵³ The SDM also performs the function of “[i]ntegrity checking: Integrity checking verifies that an accidental event has not corrupted the configuration bitstream. This function is active, even if you do not enable authentication.” (*Id.*).

150. By making, using, offering for sale, and/or selling products in the United States and/or importing products into the United States, including but not limited to the '084 Accused Products, Intel has injured BiTMICRO and is liable to BiTMICRO for directly infringing one or more claims of the '084 Patent, including without limitation claim 19 pursuant to 35 U.S.C. § 271(a).

⁵¹ Available at <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/stratix-10/ug-s10-config.pdf>

⁵² Available at <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/ug/ug-s10-soc-boot.pdf>

⁵³ Available at <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/stratix-10/ug-s10-config.pdf>

151. On information and belief, Intel is also inducing and/or has induced infringement of one or more claims of the '084 Patent, including at least claim 19, as a result of, amongst other activities, instructing, encouraging, and directing its customers on the use of the '084 Accused Products in an infringing manner in violation of 35 U.S.C. § 271(b). Through its website, instructional guides, and manuals, Intel provides its customers with detailed explanations, instructions, and information on how to use and implement the '084 Accused Products which demonstrate active steps taken to encourage direct infringement. (*See, e.g.*, Intel Stratix 10 SoC FPGA Boot User Guide⁵⁴; Intel Stratix 10 Configuration User Guide;⁵⁵ Intel Stratix 10 Hard Processor System Technical Reference Manual⁵⁶).

152. On information and belief, Intel has had knowledge of the '084 Patent at least as of the filing and/or service of this Complaint. Despite this knowledge, Intel has continued to engage in activities to encourage and assist its customers in the use of the '084 Accused Products. Thus, on information and belief, Intel (1) had actual knowledge of the patent; (2) knowingly induced its customers to infringe the patent; and (3) had specific intent to induce the patent infringement.

153. On information and belief, by using the '084 Accused Products as encouraged and assisted by Intel, Intel's customers have directly infringed and continue to directly infringe one or more claims of the '084 Patent, including at least claim 19. On information and belief, Intel knew

⁵⁴ Available at <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/ug/ug-s10-soc-boot.pdf>.

⁵⁵ Available at <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/stratix-10/ug-s10-config.pdf>

⁵⁶ Available at https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/stratix-10/s10_5v4.pdf.

or was willfully blind to the fact that its actions would induce its customers' direct infringement of the '084 Patent.

154. Intel's infringement of the '084 Patent has been and continues to be deliberate and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284-285.

155. On information and belief, Intel will continue to infringe the '084 Patent unless enjoined by this Court.

156. As a result of Intel's infringement of the '084 Patent, BiTMICRO has suffered monetary damages, and seeks recovery, in an amount to be proven at trial, adequate to compensate for Intel's infringement, but in no event less than a reasonable royalty with interest and costs. Intel's infringement of BiTMICRO's rights under the '084 Patent will continue to damage BiTMICRO, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

FIFTH COUNT

(INFRINGEMENT OF U.S. PATENT NO. 10,120,694)

157. BiTMICRO incorporates by reference the allegations set forth in paragraphs 1-156 as though fully set forth herein.

158. On information and belief, Intel has directly infringed and continues to directly infringe one or more claims of the '694 Patent, including at least claim 6 of the '694 Patent, in the state of Texas, in this judicial district, and elsewhere in the United States by, among other things, making, using, selling, offering for sale, and/or importing into the United States products that embody one or more of the inventions claimed in the '694 Patent, including but not limited to the

Intel Stratix 10 FPGA products and all reasonably similar products (“the ’694 Accused Products”), in violation of 35 U.S.C. § 271(a).

159. As an example, the Intel Stratix 10 products are “an apparatus.” Specifically, the Stratix 10 is a system-on-a-chip (SoC) that includes multiple processing components. As advertised by Intel, “[t]he Intel Stratix 10 system-on-a-chip (SoC) is composed of two distinct portions; a 64-bit quad core Arm Cortex-A53 hard processor system (HPS) and an FPGA. The HPS architecture integrates a wide set of peripherals that reduce board size and increase performance within a system.” (Intel Stratix 10 Hard Processor System Technical Reference Manual, p. 26).⁵⁷

160. The Stratix 10 is “an embedded system comprising one or more processors, a reset controller, a storage device controller, one or more direct memory access (DMA) controllers, a random access memory (RAM) and a memory controller, a nonvolatile memory and a nonvolatile memory controller, a debug interface.” For example, Intel’s product documentation for the Stratix 10 shows that the product includes at least one processor, reset controller, storage device controller,

⁵⁷ Available at https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/stratix-10/s10_5v4.pdf

DMA controller, RAM, memory controller, nonvolatile memory and nonvolatile memory controller, and a debug interface:


2.1. Features of the HPS

The main modules of the HPS are:

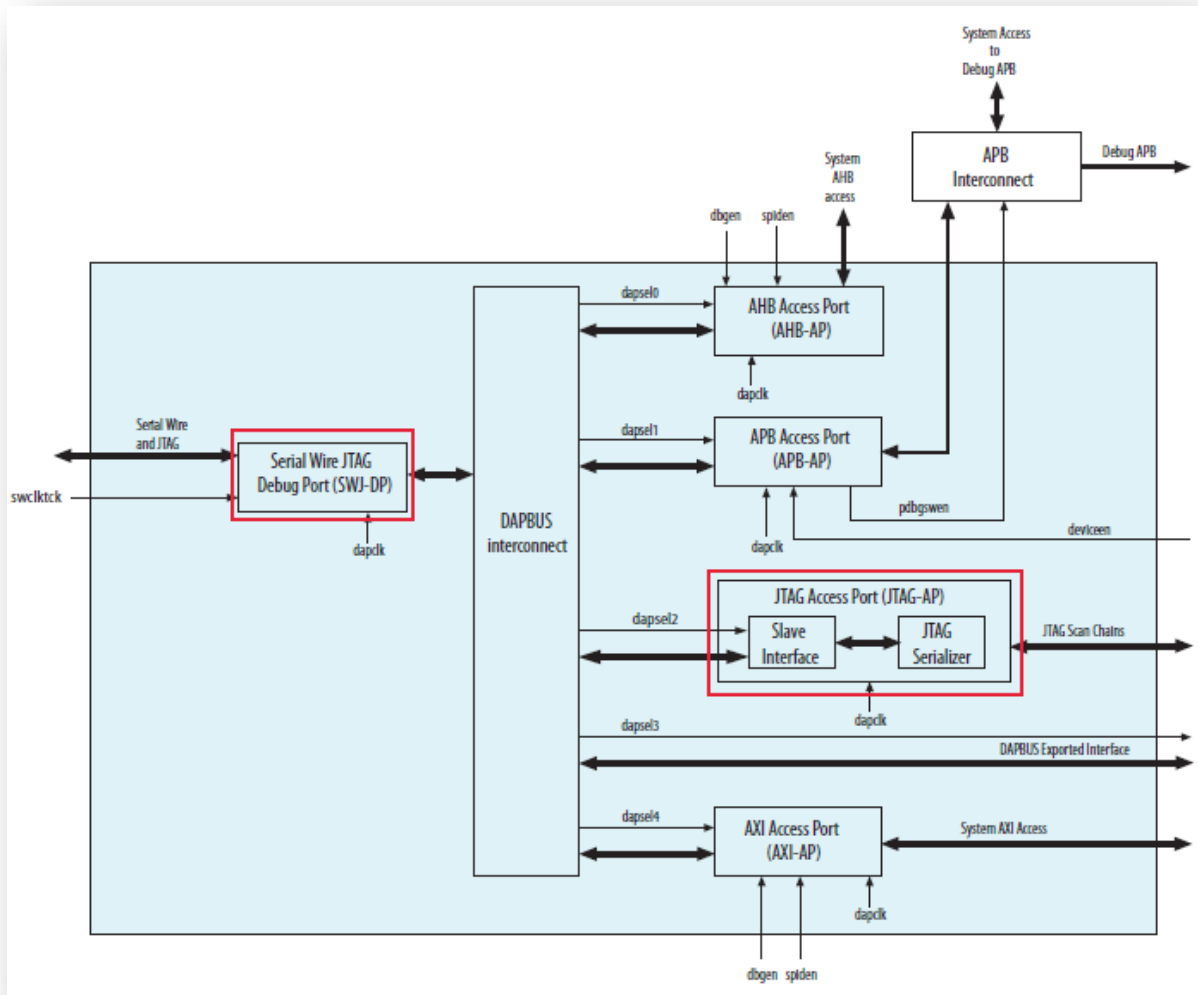
- Quad-core Arm Cortex-A53 MPCore processor
- Cache Coherency Unit (CCU)
- System Memory Management Unit (SMMU)
- System interconnect that includes:
 - Three memory-mapped interfaces between the HPS and FPGA:
 - HPS-to-FPGA bridge: 32-, 64-, or 128-bit wide Arm Advanced Microcontroller Bus Architecture (AMBA) Advanced eXtensible Interface (AXI)-4
 - Lightweight HPS-to-FPGA bridge: 32-bit wide AXI-4
 - FPGA-to-HPS bridge: 128-bit wide AXI Coherency Extensions-Lite (ACE-Lite)
 - Three memory-mapped FPGA-to-SDRAM AXI-4 interfaces, 32, 64, or 128 bits wide, allow the FPGA to directly share the HPS-connected SDRAM
- General-purpose direct memory access (DMA) controller
- 256 KB on-chip RAM
- Error checking and correction controllers for on-chip RAM and peripheral RAMs
- Clock manager
- Reset manager
- System manager
- Dedicated I/O pin multiplexer (MUX)
- NAND flash controller
- Secure digital/multimedia card (SD/MMC) controller

(Intel Stratix 10 Hard Processor System Technical Reference Manual, p. 27).⁵⁸

⁵⁸ Available at https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/stratix-10/s10_5v4.pdf



(*Id.*, p.26).



(*Id.*, p. 551).

25.4.1.1. JTAG Interface Options

The JTAG can be interfaced in two ways: through the HPS shared I/O or dedicated JTAG pins that are part of the device configuration pins. You can choose the method you want to use to connect to the DAP through Intel Quartus Prime Pro Edition. The HPS JTAG signals are multiplexed with HPS GPIO. The table below details which GPIO pin is multiplexed with each HPS JTAG pin.

(*Id.*, p. 552).

12. Reset Manager

The reset manager generates module reset signals based on reset requests from the various sources in the HPS, and software writing to the module-reset control registers.

The HPS contains multiple reset domains. Each reset domain can be reset independently. A reset may be initiated externally, internally or through software.

Table 99. HPS Reset Domains

Reset Domain	Reset Source	Description
POR (Power-on Reset)	Secure Device Manager (SDM)	SDM requests reset manager to assert POR reset. During a voltage tampering or out-of-range event, the SDM asserts POR. When voltage returns to operating range, the POR is de-asserted. During POR, the entire HPS and FPGA is reset. When the device is released from POR, SDM begins initialization.
System Cold Reset ⁽²¹⁾	<ul style="list-style-type: none"> SDM (HPS mailbox message)⁽²²⁾ HPS_COLD_NRESET pin⁽²³⁾ 	SDM requests reset manager to assert or de-assert cold reset.
System Warm Reset	<ul style="list-style-type: none"> Software requests a warm reset through the EL3 register 	Reset manager asserts warm reset provided that the Cortex-A53 MPCore is idle. An L2 reset must be performed before requesting a warm reset. Before you request L2 reset via software, you must flush L2 using the <code>l2flushen</code> bit of the <code>hdsken</code> register. <i>Note:</i>
Watchdog Reset	Watchdog Timeout Event	Reset manager asserts watchdog reset based on the watchdog timer register. As the CoreSight logic is not reset, the debug/trace can continue immediately after reset manager de-asserts watchdog reset.
MPU Cold Reset	Software requests a cold reset through the <code>COLDMODRST</code> register	Reset manager asserts cold reset to the MPU provided that all four cores are idle. Before you request MPU cold reset via software, you must idle all four cores using a WFI instruction and flush L2 using the <code>l2flushen</code> bit of the <code>hdsken</code> register. <i>Note:</i>

continued...

(*Id.* at p. 214).

2.2.9.6. DMA Controller

The DMA controller provides high-bandwidth data transfers for modules without integrated DMA controllers. The DMA controller is based on the Arm CoreLink* DMA Controller (DMA-330) and offers the following features:

- Micro-coded to support flexible transfer types
 - Memory-to-memory
 - Memory-to-peripheral
 - Peripheral-to-memory
 - Scatter-gather
- Supports up to eight channels
- Supports up to 32 peripheral request interfaces

(*Id.* at p. 36).

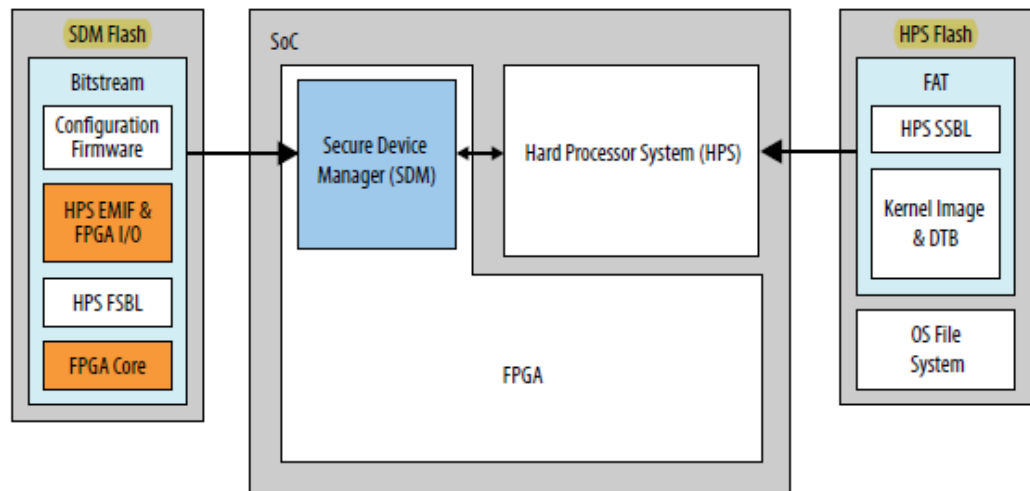
16. SD/MMC Controller

The hard processor system (HPS) provides a Secure Digital/Multimedia Card (SD/MMC) controller for interfacing to external SD and MMC flash cards, secure digital I/O (SDIO) devices, and Consumer Electronics Advanced Transport Architecture (CE-ATA) hard drives.

The SD/MMC flash controller enables you to use the flash card to expand the on-board storage capacity for larger applications or user data. Other applications include interfacing to embedded SD (eSD) and embedded MMC (eMMC) non-removable flash devices.

(*Id.*, p. 276).

Figure 5. FPGA Configuration First Dual SDM and HPS Flash



(Intel Stratix 10 SoC FPGA Boot User Guide, p. 12).⁵⁹

161. The Stratix 10 includes “a power-on reset (POR) sequencer, wherein the POR sequencer uses a POR sequencer descriptor which is a preassembled descriptor that is stored in the nonvolatile memory.” For example, the Secure Device Manager (“SDM”) in the Stratix 10 is a power-on reset (POR) sequencer. The SDM “is a triple-redundant processor-based module that

⁵⁹ Available at <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/ug/ug-s10-soc-boot.pdf>


manages configuration and the security features of the Intel Stratix 10 devices. The SDM is available on all Intel Stratix 10 FPGA and SoC devices.” (Intel Stratix 10 Configuration User Guide, p. 12).⁶⁰ For startup operations, the SDM uses a configuration bitstream. For example, as described by Intel, “the SDM samples the MSEL [2:0] pins to determine the boot source. Next, the device configures the SDM I/Os according to the selected boot source interface and the SDM retrieves the configuration bitstream through the interface.” (Intel Stratix 10 SoC FPGA Boot User Guide, p. 8).⁶¹ The configuration bitstream constitutes a POR sequencer descriptor, which is a preassembled descriptor that is stored in nonvolatile memory such as a Quad SPI flash memory or flash memory within the SDM. (See Intel Stratix 10 SoC FPGA Boot User Guide, pp. 11-12).⁶²

162. In the Stratix 10, “the POR sequencer descriptor includes register information and DMA controller descriptors.” For example, the configuration bitstream includes information to initialize the FPGA I/O and FPGA Core, first-stage bootloader (FSBL) for the HPS, and SDM configuration firmware:

⁶⁰ Available at <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/stratix-10/ug-s10-config.pdf>

⁶¹ Available at <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/ug/ug-s10-soc-boot.pdf>

⁶² Available at <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/ug/ug-s10-soc-boot.pdf>



(Intel Stratix 10 SoC FPGA Boot User Guide, p. 12).⁶³ This initialization process also includes “initialization of the FPGA fabric, including registers and state machines.” (Intel Stratix 10 Configuration User Guide, p. 20).⁶⁴

163. In the Stratix 10, “the POR sequencer reads and uses the register information to update configuration registers of the embedded system.” For example, as discussed above, the SDM uses the information from the configuration bitstream to initialize the registers in the system.

164. By making, using, offering for sale, and/or selling products in the United States and/or importing products into the United States, including but not limited to the ’694 Accused Products, Intel has injured BiTMICRO and is liable to BiTMICRO for directly infringing one or more claims of the ’694 Patent, including without limitation claim 6 pursuant to 35 U.S.C. § 271(a).

⁶³ Available at <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/ug/ug-s10-soc-boot.pdf>

⁶⁴ Available at <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/stratix-10/ug-s10-config.pdf>

165. On information and belief, Intel is also inducing and/or has induced infringement of one or more claims of the '694 Patent, including at least claim 6, as a result of, amongst other activities, instructing, encouraging, and directing its customers on the use of the '084 Accused Products in an infringing manner in violation of 35 U.S.C. § 271(b). Through its website, instructional guides, and manuals, Intel provides its customers with detailed explanations, instructions, and information on how to use and implement the '694 Accused Products which demonstrate active steps taken to encourage direct infringement. (*See, e.g.*, Intel Stratix 10 SoC FPGA Boot User Guide⁶⁵; Intel Stratix 10 Configuration User Guide;⁶⁶ Intel Stratix 10 Hard Processor System Technical Reference Manual⁶⁷).

166. On information and belief, Intel has had knowledge of the '694 Patent at least as of the filing and/or service of this Complaint. Despite this knowledge, Intel has continued to engage in activities to encourage and assist its customers in the use of the '694 Accused Products. Thus, on information and belief, Intel (1) had actual knowledge of the patent; (2) knowingly induced its customers to infringe the patent; and (3) had specific intent to induce the patent infringement.

167. On information and belief, by using the '694 Accused Products as encouraged and assisted by Intel, Intel's customers have directly infringed and continue to directly infringe one or more claims of the '694 Patent, including at least claim 6. On information and belief, Intel knew or

⁶⁵ Available at <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/ug/ug-s10-soc-boot.pdf>.

⁶⁶ Available at <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/stratix-10/ug-s10-config.pdf>

⁶⁷ Available at https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/stratix-10/s10_5v4.pdf.

was willfully blind to the fact that its actions would induce its customers' direct infringement of the '694 Patent.

168. Intel's infringement of the '694 Patent has been and continues to be deliberate and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284-285.

169. On information and belief, Intel will continue to infringe the '694 Patent unless enjoined by this Court.

170. As a result of Intel's infringement of the '694 Patent, BiTMICRO has suffered monetary damages, and seeks recovery, in an amount to be proven at trial, adequate to compensate for Intel's infringement, but in no event less than a reasonable royalty with interest and costs. Intel's infringement of BiTMICRO's rights under the '694 Patent will continue to damage BiTMICRO, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

SIXTH COUNT

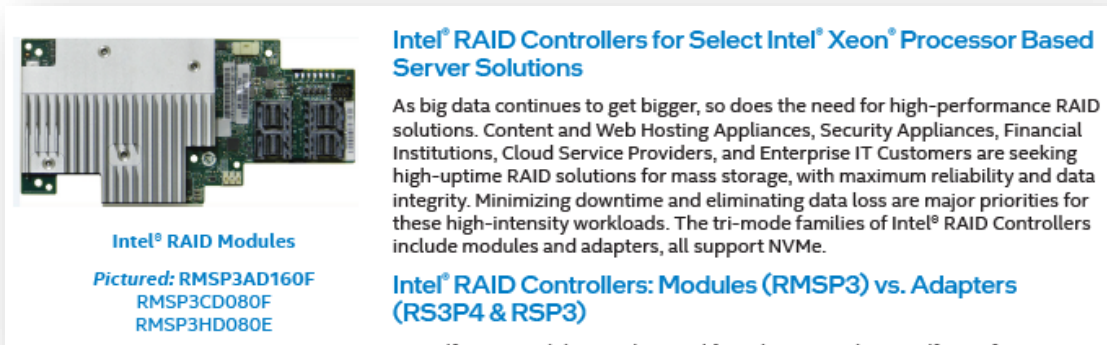
(INFRINGEMENT OF U.S. PATENT NO. 6,496,939)

171. BiTMICRO incorporates by reference the allegations set forth in paragraphs 1-170 as though fully set forth herein.


172. Intel has directly infringed one or more claims of the '939 Patent, including at least claim 10 of the '939 Patent, in the state of Texas, in this judicial district, and elsewhere in the United States by, among other things, making, using, selling, offering for sale, and/or importing into the United States products that embody one or more of the inventions claimed in the '939 Patent, including but not limited to the above-identified RAID Controllers / RMFBU, as well as

SSDs with Power Loss Protection, and all reasonably similar products (“the ’939 Accused Products”), in violation of 35 U.S.C. § 271(a).

173. As an example, the RAID Controllers / RMFBU comprise a “computer system” with a “computing engine” to receive, transmit, and manage the storage of data, as demonstrated below.



(Intel, Tri-mode Intel Storage Controllers supporting PCIe NVMe: Product Brief, at 1).⁶⁸



⁶⁸ Available at <https://cdrdv2.intel.com/v1/dl/getContent/639394?explicitVersion=true>.

(Intel RAID Modules, Hardware User Guide, Rev. 1.4 (Sept. 2020), at 42).⁶⁹

174. The RAID Controllers / RMFBUs include “a system for controlling data in a computer system when the computer system loses power” that includes “means for activating a plurality of super capacitors to supply power to the computing engine based upon power being removed from the computer system.” The system further includes “means for reconfiguring the data in the computing engine,” for example, by flushing the data from the volatile memory to non-volatile flash memory when power is removed. Specifically, the RMFBUs “monitor the voltage level of the DRAM” in the RAID Controller and “[i]f the voltage drops below a predefined level due to an [external] AC power failure or brief power outage, the RMFBUs protect the integrity of the cached data by providing sufficient back-up power to offload the data from the RAID RAM to the NAND Flash.”

The cache memory available on Intel® RAID Modules and Intel® RAID Adapters can improve the overall system performance. Writing data to the module’s cache memory is much faster than writing data to a storage device. Write operations appear to complete very quickly at the software application level. The Intel® RAID Module writes the cached data to the storage device when system activity is low or when the cache is getting full. The risk of using write-back cache is that the cached data can be lost if the AC power fails before it is written to the storage device. This risk factor is eliminated when the Intel® RAID Module has an RMFBUs installed. In addition,

(Intel RAID Modules, Hardware User Guide, Rev. 1.4 (Sept. 2020), at 38).⁷⁰

During normal system operation, the RMFBUs monitor the voltage level of the DRAM of the Intel® RAID Module or Intel® RAID Module. If the voltage drops below a predefined level due to an AC power failure or brief power outage, the RMFBUs protect the integrity of the cached data by providing sufficient back-up power to offload the data from the RAID RAM to the NAND flash. When the voltage level returns to an acceptable level, the RAID RAM is recovered from flash, and all pending writes to storage devices are completed without losing any data.

⁶⁹ Available at https://www.intel.com/content/dam/support/us/en/documents/server-products/raid-products/Tri_Mode_Modules_Hardware_User_Guide.pdf.

⁷⁰ Available at https://www.intel.com/content/dam/support/us/en/documents/server-products/raid-products/Tri_Mode_Modules_Hardware_User_Guide.pdf.

(*Id.*).

175. The system of the RAID Controllers / RMFBU also includes “means for deactivating the plurality of super capacitors to cut off power to the computing engine based upon the plurality of super capacitors discharging to a predetermined level,” for example, when the capacitors are spent.

The AXRMFBU7 RMFBU is a kit consisting of a FBU345 super capacitor bank, a plastic bracket to attach the super capacitor bank to the chassis, and two (2) extender cables, one 605 mm long and the other 930 mm long. The FBU345 consists of a bank of Electric Double-Layer Capacitors (EDLC) or super capacitors capable of storing a high amount of electric energy while active. In the event of a power failure, the FBU345 provides the power needed for the data offload. It also has an over-temperature detection circuitry and a discharge circuitry that discharges the capacitors while disconnected from the RAID module or when the system on which it is installed is turned off.

(*Id.*).

176. As another example, Intel’s SSDs with Power Loss Data Protection, also known as Enhanced Power Loss Data Protection and/or Power Loss Imminent (PLI) technology, include “a system for controlling data in a computer system when the computer system loses power.” For example, as described by Intel’s documentation for the SSD DC-S3700, this feature “reduces potential data loss by detecting and protecting data from an unexpected system power loss” by, for example, “sav[ing] all cached data in the process of being written before shutting down.”

Enhanced Power-Loss Data Protection. Reduces potential data loss by detecting and protecting data from an unexpected system power loss. The drive saves all cached data in the process of being written before shutting down, thereby minimizing potential data loss

(Intel® SSD DC S3700 Series, Product Brief).⁷¹

177. The system of Intel's SSDs with Power Loss Data Protection further includes "means for activating a plurality of super capacitors to supply power to the computing engine based upon power being removed from the computer system." Specifically, as described by Intel's marketing and product documentation, "SSDs with PLI technology contain energy storing capacitors. The energy storing capacitors provide enough energy (power) to complete any commands in progress and to make sure that any data in the temporary buffers is committed to the non-volatile NAND Flash media. These capacitors act as backup batteries for the drive." As further illustrated by Intel:

⁷¹ Available at <https://www.intel.com/content/dam/support/us/en/documents/solid-state-drives/ssd-dc-s3700-brief.pdf>.

PLI Technology Details

Not all SSDs contain PLI technology. SSDs with PLI technology contain energy storing capacitors. The energy storing capacitors provide enough energy (power) to complete any commands in progress and to make sure that any data in the temporary buffers is committed to the non-volatile NAND Flash media. These capacitors act as backup batteries for the drive.

Figure 1 highlights (yellow box) a PLI capacitor in an SSD. The SSD supply voltage detector circuit constantly monitors the drive's supply voltage. If voltage falls below a predefined level, indicating a power loss event is imminent, the SSD will use the backup energy in the capacitors to write any data in temporary buffers to the NAND Flash media. After power is restored, the PLI capacitors will be charged again, ready for a future power loss event.




Figure 1: 1 of 2 PLI Capacitors shown from Intel® SSD DC 3700 Series Product

(Intel, “Power Loss Imminent (PLI) Technology,” at 2).⁷²

178. As described above and below, Intel’s SSDs with Power Loss Data Protection also include means for “reconfiguring the data in the computing engine” in the event of a power loss, and “deactivating the plurality of super capacitors to cut off power to the computing engine based upon the plurality of super capacitors discharging to a predetermined level.” For example, the figure below “shows an operational architectural diagram depicting the key PLI components, the SSD voltage detector, the charge storing capacitor bank, and the SMART (Self-Monitoring Analysis and Monitoring Technology) attribute extraction module.” “The SSD voltage detector circuit constantly monitors the drive’s supply voltage. If the voltage falls below a predefined level, indicating a power loss event is imminent, the SSD will use the backup energy supply from the capacitor bank to write any data in temporary buffers to the NAND Flash media. After the power is restored, the PLI capacitors will be charged up again and ready for a future power loss event.”

⁷² Available at <https://www.intel.com/content/dam/www/public/us/en/documents/technology-briefs/ssd-power-loss-imminent-technology-brief.pdf>.

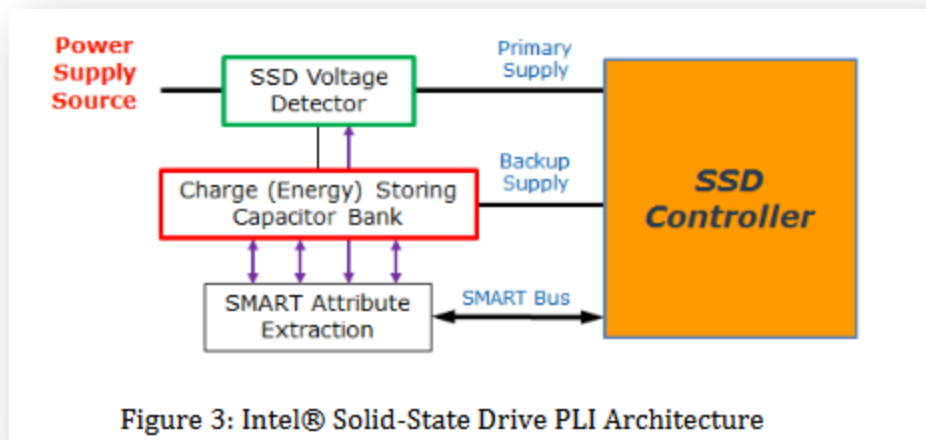


Figure 3: Intel® Solid-State Drive PLI Architecture

(Intel, “Power Loss Imminent (PLI) Technology,” at 2).⁷³

179. Intel explains the benefits of Enhanced Power-Loss Data Protection a/k/a Power Loss Imminent Technology. This includes, but is not limited to improving data integrity during power loss events, help speed the data center recover after a power loss event, and much more.

What Do PLI-Enabled SSDs Do for Me?

PLI technology significantly improves data integrity during power loss events whether from acts of nature, careless mistakes, or even inadvertent hot removal of the SSD storage device. Improved data integrity can help speed the data center recovery after a power loss event with fewer data path checks and validation. Bringing the data center back online as fast as possible will minimize users’ delayed access to their data, media, or web materials.

Do I Need PLI in My Storage Device Platforms?

PLI technology offers many advantages by adding higher levels of integrity to data storage environments. Before deciding, it is critical to evaluate the geographical location and its susceptibility to power outages. In addition, evaluate the robustness of the data center regarding power outage risks and mitigating the loss of critical data. Another critical factor in deciding if PLI is the right data center choice is evaluating the workload. If it is more write than read intensive, what is the mix? Is the nature of the data being read or written more random or sequential? If it data is long sequential writes, how long are they? What file system is deployed and does it coalesce the writes? Obviously there is a lot to analyze and consider, including whether the cost of time to rebuild backup data is less than the cost of buying a drive with PLI technology. Down time is lost time, and will impact the total cost of ownership of the entire data center.

(Intel, “Power Loss Imminent (PLI) Technology,” at 7).⁷⁴

⁷³ Available at <https://www.intel.com/content/dam/www/public/us/en/documents/technology-briefs/ssd-power-loss-imminent-technology-brief.pdf>.

⁷⁴ Available at <https://www.intel.com/content/dam/www/public/us/en/documents/technology-briefs/ssd-power-loss-imminent-technology-brief.pdf>.

180. As a result of Intel's infringement of the '939 Patent, BiTMICRO has suffered monetary damages, and seeks recovery, in an amount to be proven at trial, adequate to compensate for Intel's infringement, but in no event less than a reasonable royalty with interest and costs.

PRAYER FOR RELIEF

WHEREFORE, BiTMICRO prays for judgment and seeks relief against Intel as follows:

- A. For judgment that Intel has infringed and/or continues to infringe one or more claims of the Asserted Patents, directly, and/or indirectly by way of inducement;
- B. For a permanent injunction against Intel and its respective officers, directors, agents, servants, affiliates, employees, divisions, branches, subsidiaries, parents, and all other acting in active concert therewith from infringement of the Asserted Patents;
- C. For an accounting of all damages sustained by BiTMICRO as the result of Intel's acts of infringement;
- D. For a mandatory future royalty payable on each and every future sale by Intel of a product that is found to infringe one or more of the Asserted Patents and on all future products which are reasonably similar to those products found to infringe;
- E. For a judgment and order finding that Intel's infringement is willful and awarding to BiTMICRO enhanced damages pursuant to 35 U.S.C. § 284;
- F. For a judgment and order requiring Intel to pay BiTMICRO's damages, costs, expenses, and pre- and post-judgment interest for its infringement of the Asserted Patents as provided under 35 U.S.C. § 284;
- G. For a judgment and order finding that this is an exceptional case within the meaning of 35 U.S.C. § 285 and awarding to BiTMICRO its reasonable attorneys' fees; and

H. For such other and further relief in law and in equity as the Court may deem just and proper.

DEMAND FOR JURY TRIAL

Pursuant to Rule 38(b) of the Federal Rules of Civil Procedure, BiTMICRO hereby demands a trial by jury of this action.

Dated: March 30, 2022

Respectfully submitted,

/s/ B. Russell Horton

Denise M. De Mory (*Pro Hac Pending*)
California State Bar No. 168076
Corey Johanningmeier (*Pro Hac Pending*)
California State Bar No. 251297
Richard C. Lin (*Pro Hac Pending*)
California State Bar No. 209233
Michael Flynn-O'Brien (*Pro Hac Pending*)
California State Bar No. 291301
Bunsow De Mory LLP
701 El Camino Real
Redwood City, CA 94063
(650) 351-7241 Telephone
(415) 426-4744 Facsimile
ddemory@bdiplaw.com
cjohanningmeier@bdiplaw.com
rlin@bdiplaw.com
mflynno'Brien@bdiplaw.com

Attorney in Charge for Plaintiff BiTMICRO LLC

B. Russell Horton
State Bar No. 10014450
George Brothers Kincaid & Horton, L.L.P.
114 West 7th Street, Ste. 1100
Austin, Texas 78701
(512) 495-1400 Telephone
(512) 499-0094 Facsimile
rhorton@gbkh.com

Attorney for Plaintiff BiTMICRO LLC