



- (51) **International Patent Classification:**  
*G06F 1/32* (2006.01) *G06F 13/16* (2006.01)
- (21) **International Application Number:**  
PCT/US2016/036354
- (22) **International Filing Date:**  
8 June 2016 (08.06.2016)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**  
14/803,732 20 July 2015 (20.07.2015) US
- (71) **Applicant:** SANDISK TECHNOLOGIES LLC  
[US/US]; 6900 Dallas Parkway, Suite 325, Plano, TX  
75024 (US).
- (72) **Inventors:** TZAFRIR, Yonatan; Yael Rom 13, 49062 Pe-  
tah Tikva (IL). BORUKHOV, Hannon Aharon; Miler St  
5, 76284 Rehovot (IL).
- (74) **Agent:** HETZ, Joseph, F.; Brinks Gilson & Lione, P.O.  
Box 10087, Chicago, IL 60610 (US).

(81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) **Designated States** (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

**Published:**

— with international search report (Art. 21(3))

(54) **Title:** MEMORY SYSTEM AND METHOD FOR ADAPTIVE AUTO-SLEEP AND BACKGROUND OPERATIONS




Figure 1

(57) **Abstract:** A memory system and method are provided for adaptive auto-sleep and background operations. In one embodiment, a controller of a memory system measures an amount of time between when the memory completes an operation and when the controller receives a command to perform another operation in the memory. The controller adjusts a time period after which the controller enters an auto-sleep mode and/or starts a background operation based on the measured amount of time. Other embodiments are disclosed

## **Memory System and Method for Adaptive Auto-Sleep and Background Operations**

### **Background**

[0001] Some memory systems with non-volatile memories (such as flash memory) have an auto-sleep function. Typically, the controller of such a memory system is programmed with a predetermined time threshold, and if the memory has been idle for longer than the predetermined time threshold (e.g., due to inactivity of a host), the controller places the memory system in an auto-sleep mode to reduce the current and power consumed by the memory system. Before the memory system is put to sleep, the controller typically performs various tasks to prepare for sleep mode, such as saving user and management data from volatile memory to non-volatile memory and powering off hardware components. When the controller senses a new command from the host, the controller typically performs various tasks to exit from sleep mode, such as copying user and management data from non-volatile memory to volatile memory and powering on hardware components.

### **Brief Description of the Drawings**

[0002] Figure 1 is a block diagram of an exemplary memory system of an embodiment.

[0003] Figure 2A is a block diagram of a memory system of an embodiment removably connected to a host.

[0004] Figure 2B is a block diagram of a memory system of an embodiment embedded in a host.

[0005] Figure 3 is a flow chart of a method of an embodiment for adaptive auto-sleep.

[0006] Figure 4 is a flow chart of a method of an embodiment for adaptive auto-sleep using interrupt and self-learning circuitry.

[0007] Figure 5 is a graph illustrating adaptive auto-sleep of an embodiment.

[0008] Figure 6 is a flow chart of a method of an embodiment for adaptive background operations.

## Detailed Description

[0009] By way of introduction, the below embodiments relate to a memory system and method for adaptive auto-sleep and background operations is provided. In one embodiment, a controller of a memory system measures an amount of time between when the memory completes an operation and when the controller receives a command to perform another operation in the memory. The controller adjusts a time period after which the controller enters an auto-sleep mode based on the measured amount of time.

[0010] In some embodiments, the time period is adjusted based on the measured amount of time and on at least one additional amount of time for at least one additional operation.

[0011] In some embodiments, the time period can be reduced when the measured amount of time is less than the time period, while, in other embodiments, the time period is reduced when the measured amount of time is less than a sum of the time period and a time it takes the controller to enter and exit the auto-sleep mode.

[0012] In some embodiments, the time period is increased when the measured amount of time is greater than the time period, while, in other embodiments, the time period is increased when the measured amount of time is greater than a sum of the time period and a time it takes the controller to enter and exit the auto-sleep mode.

[0013] In some embodiments, the controller measures the amount of time by measuring time between when the controller receives an end signal via a memory interface and when the controller received a start signal via a host interface.

[0014] In another embodiment, a memory system is disclosed having a memory and a controller. The controller is configured to measure idle time of the memory and determine if the measured idle time of the memory is less than or greater than an auto-sleep threshold. In response to determining that the measured idle time is less than the auto-sleep threshold, the controller decreases the auto-sleep threshold. In response to determining that the measured idle time is greater than the auto-sleep threshold, the controller increases the auto-sleep threshold.

[0015] In some embodiments, the controller is configured to measure idle time by measuring an amount of time between when the controller completes an operation in the

memory and when the controller receives a command to perform another operation in the memory.

[0016] In some embodiments, the controller is configured to measure the amount of time by measuring time between when the controller receives an end signal via a memory interface and when the controller received a start signal via a host interface.

[0017] In some embodiments, the controller is configured to use a plurality of measured idle times to determine if the auto-sleep threshold should be increased or decreased.

[0018] In some embodiments, the auto-sleep threshold includes a time it takes the controller to enter and exit the auto-sleep mode.

[0019] In some embodiments, the controller is configured to measure idle time using interrupt circuitry and a clock, and the controller is further configured to determine if the measured idle time of the controller is less than or greater than the auto-sleep threshold using self-learning circuitry.

[0020] In another embodiment, the above method and/or memory system is used to performing adaptive background operations, in addition to or instead of performing adaptive auto-sleep.

[0021] In some embodiments, the memory is a three-dimensional memory.

[0022] In some embodiments, the memory system is embedded in a host, while, in other embodiments, the memory system is removably connected to a host.

[0023] Other embodiments are possible, and each of the embodiments can be used alone or together in combination.

[0024] As mentioned above, these embodiments generally relate to a memory system and method for adaptive auto-sleep and background operations. Before turning to several examples, the follow paragraphs present an exemplary memory system that can be used with these embodiments.

[0025] Turning now to the drawings, Figure 1 is a block diagram of a memory system 100 of an embodiment. As shown in Figure 1, the memory system 100 of this embodiment comprises a controller 110 in communication with a memory die 120. As used herein, the phrase "in communication with" could mean directly in communication

with or indirectly in communication with through one or more components, which may or may not be shown or described herein.

[0026] In general, the controller 110 transmits command sequences for read, program, and erase operations to the memory die 120, and performs other functions, as will be discussed below. The controller 110 can be implemented in any suitable manner. For example, the controller 110 can take the form of processing circuitry, a microprocessor or processor, and a computer-readable medium that stores computer-readable program code (e.g., firmware) executable by the (micro)processor, logic gates, switches, an application specific integrated circuit (ASIC), a programmable logic controller, and an embedded microcontroller, for example. The controller 110 can be configured with hardware and/or firmware to perform the various functions described below and shown in the flow diagrams. Also, some of the components shown as being internal to the controller 110 can also be stored external to the controller 110, and other components can be used.

[0027] In some embodiments, the memory system 100 comprises a single memory die 120, while, in other embodiments, the memory system 100 comprises at least one additional memory die 125. As used herein, the term “die” refers to the collection of non-volatile memory cells, and associated circuitry for managing the physical operation of those non-volatile memory cells, that are formed on a single semiconductor substrate. A memory die may include any suitable non-volatile storage medium, including NAND flash memory cells and/or NOR flash memory cells. The memory cells can take the form of solid-state (e.g., flash) memory cells and can be one-time programmable, few-time programmable, or many-time programmable. The memory cells can also be single-level cells (SLC), multiple-level cells (MLC), triple-level cells (TLC), or use other memory cell level technologies, now known or later developed. Also, the memory cells can be fabricated in a two-dimensional or three-dimensional fashion.

[0028] As shown in Figure 1, in one embodiment, the memory system's controller 110 comprises a host interface (or host interface module (HIM)) 115, a memory (or backend) interface 125, interrupt circuitry 130, a clock 135, and self-learning circuitry 140. The controller 110 can contain other components as well, such as, but not limited to, random access memory (RAM), read only memory (ROM) (e.g., to store system boot code), an error correction controller (ECC) engine to encode/decode data written to and

read from the memory die 120, and a crypto-engine to provide encryption and/or decryption operations.

[0029] The host interface 115 can take any suitable form, such as, but are not limited to, a SATA, SATA Express, SAS, Fibre Channel, USB, PCIe, and NVMe interface. The memory interface 125 can also take any suitable form, such as, but not limited to, a flash interface (e.g., Toggle Mode 200, 400, or 800). The interrupt circuitry 130, which takes the form of a programmable hardware block in one embodiment, is configured to receive interrupts from multiple interrupt sources on every change associated with data of a serviced command. For example, the interrupts can be host interrupts that are received via the host interface 115 and/or interrupts received from the memory 120 via the memory interface 125. In one embodiment, the host interface 115 receives commands from the host and provides a transaction start signal as an interrupt to the interrupt circuitry 130 in response to receiving every new command. Similarly, the memory interface 125 receives a transaction end (end-of-operation (EOP)) signal from the memory 120 with every completion of a transaction in the memory 120.

[0030] In this embodiment, when the interrupt circuitry 130 receives a transaction end signal from the memory interface 125, the interrupt circuitry 130 activates the clock 135 (which can be a hardware component that serves as a real time or non-real time clock). The interrupt circuitry 130 stops the clock 135 in response to receiving the transaction start signal from the host interface 115. In this way, the clock 135 measures the amount of time between when the memory 120 completes an operation to when the controller 110 receives a command to perform another operation in the memory 120. This measured time is then provided to the self-learning circuitry 140, which uses this measured time to adjust an auto-sleep threshold (i.e., a time period after which the controller 110 enters an auto-sleep mode), which will be discussed in more detail below.

[0031] As used herein, the term "circuitry" can refer to an instruction processor, such as a Central Processing Unit (CPU), microcontroller, or a microprocessor; or an Application Specific Integrated Circuit (ASIC), Programmable Logic Device (PLD), or Field Programmable Gate Array (FPGA); or a collection of discrete logic or other circuit components, including analog circuit components, digital circuit components or both; or any combination thereof. Circuitry may include discrete interconnected hardware components or may be combined on a single integrated circuit die, distributed among

multiple integrated circuit dies, or implemented in a Multiple Chip Module (MCM) of multiple integrated circuit dies in a common package, as examples.

[0032] Accordingly, “circuitry” may store or access instructions for execution or may implement its functionality in hardware alone. The instructions may be stored in a tangible storage medium that is other than a transitory signal, such as a flash memory, a Random Access Memory (RAM), a Read Only Memory (ROM), an Erasable Programmable Read Only Memory (EPROM); or on a magnetic or optical disc, such as a Compact Disc Read Only Memory (CDROM), Hard Disk Drive (HDD), or other magnetic or optical disk; or in or on another machine-readable medium. A product, such as a computer program product, may include a storage medium and instructions stored in or on the medium, and the instructions when executed by the circuitry in a device may cause the device to implement any of the processing described above or illustrated in the drawings.

[0033] The implementations may be distributed. For instance, circuitry may include multiple distinct system components, such as multiple processors and memories, and may span multiple distributed processing systems. Parameters, databases, and other data structures may be separately stored and managed, may be incorporated into a single memory or database, may be logically and physically organized in many different ways, and may be implemented in many different ways. Example implementations include linked lists, program variables, hash tables, arrays, records (e.g., database records), objects, and implicit storage mechanisms. Instructions may form parts (e.g., subroutines or other code sections) of a single program, may form multiple separate programs, may be distributed across multiple memories and processors, and may be implemented in many different ways. Example implementations include stand-alone programs and part of a library, such as a shared library like a Dynamic Link Library (DLL). The library, for example, may contain shared data and one or more shared programs that include instructions that perform any of the processing described above or illustrated in the drawings, when executed by the circuitry. Accordingly, “circuitry” may be implemented in many different ways and in many different combinations of hardware and software.

[0034] The memory system 100 can take any suitable form. For example, the memory system 100 can be a removable mass storage device, such as, but not limited to, a

handheld, removable memory device, such as a memory card (e.g., a Secure Digital (SD) card, a micro Secure Digital (micro-SD) card, or a MultiMedia Card (MMC)), or a universal serial bus (USB) device. The memory system 100 can also take the form of an embedded mass storage device, such as an eSD/eMMC embedded flash drive or an embedded or removable solid state drive (SSD) found in, for example, portable computing devices, such as laptop computers, and tablet computers.

[0035] As shown in Figure 2A, in some embodiments, the memory system 100 is removably connected to a host 50 (e.g., where the host 50 and memory system 100 can each have mating physical connectors that allow the memory system 100 to be removably connected to the host 50). In other embodiments (shown in Figure 2B), the memory system 100 is embedded in a host 50 (Figure 2B).

[0036] A host can take any suitable form. A host 50 typically has a controller and an interface for communicating with the memory system 100. In some implementations, the host 50 has a crypto-engine operative to provide encryption and/or decryption operations for communication with the memory system 100 (e.g., for mutual authentication, key exchange, and establishment of a secure channel). The host 50 can contain other components (e.g., a display device, a speaker, a headphone jack, a video output connection, etc.). Examples of hosts 50 include, but are not limited to, a mobile phone, a digital media player, a game device, a personal digital assistant (PDA), a personal computer (PC), a kiosk, a set-top box, a TV system, a book reader, or any combination thereof.

[0037] As mentioned in the background section above, some memory systems have an auto-sleep function to reduce the current and power consumed by the memory system when not in use. Typically, the controller of such a memory system is programmed with a predetermined time threshold, and if the memory has been idle for longer than the predetermined time threshold (e.g., due to inactivity of a host), the controller places the memory system in an auto-sleep mode. Before the memory system is put to sleep, the controller typically performs various tasks to prepare for sleep mode, such as saving user and management data from volatile memory to non-volatile memory, powering off hardware components (e.g., hardware cores and volatile memories inside the controller) and latching logic or hardware cores. When the controller senses a new command from

the host that requires a memory operation, the controller typically performs various tasks to exit from sleep mode, such as copying user and management data from non-volatile memory to volatile memory, powering on hardware components, and unlatching logic or hardware cores. The time it takes for the controller to exit auto-sleep mode will be referred to herein as "t1," and the time it takes for the controller to enter auto-sleep mode will be referred to herein as "t2."

[0038] The controller starts the auto-sleep process after measuring a time T of no memory activity, e.g., due to inactivity of a host. However, host processors vary in their behavior during idle (no-activity) time based on the type, brand, and model of the host processor. Also, even the same host processor can change its behavior (e.g., bus speed, bus bits width, single data rate vs. double data rate, etc.) depending on the mode in which it operates and the main task it is performing (e.g., sequential vs. random activity, write vs. read vs. erase, size of data chunk, etc.). This results in idle time and frequency that varies from one device controller to the other. As such, using a constant (fixed) time T against which to measure idle time of the memory can lead to non-ideal behavior, as a given time T may be too short in some situations and too long in others. If T is fixed at a time that is too short, the controller will enter auto-sleep too frequently, incurring the time penalty of t1 and t2 too often, thereby degrading the speed and performance of the memory system. This can occur, for example, when the host is frequently stopping its activity after  $T + \Delta t$  time, where  $\Delta t \ll T$ , and then going back to activity right afterwards. Likewise, if T is fixed at a time that is too long, the controller will not enter auto-sleep frequently enough, which means that the memory system is consuming more current and power than it needs to.

[0039] To address this situation, the following embodiments can be used to provide an adaptive auto-sleep mode. Instead of using a constant time T against which to measure the idle time of the memory 120, an adjustable time period can be used. More specifically, the time T in which the controller 110 enters auto-sleep mode is adaptive and dynamically self-adjusted in an optimal manner based on the actual operation of the controller 110. This allows for optimizing operation of the controller 110 in sleep mode in a dynamic manner, thereby improving overall system performance and achieving less power consumption.

[0040] Turning again to the drawings, Figure 3 is a flow chart 300 that illustrates this embodiment. As shown in Figure 3, the controller 110 measures the idle time of the memory (act 310). Next, the controller 110 determines if the measured idle is less than or greater than an auto-sleep threshold (act 320). In response to determining that the measured idle time is less than the auto-sleep threshold, the controller 110 decreases the auto-sleep threshold (act 330). In response to determining that the measured idle time is greater than the auto-sleep threshold, the controller 110 increases the auto-sleep threshold (act 340). In this way, the controller 110 can reduce the time period when the measured amount of time is less than the auto-sleep threshold (time period T), or less than T plus t1 and t2. That is, if the memory 120 is idle less than the time period T (or less than T plus t1 and t2), the controller 110 can reduce the time period T, which would cause the memory system 100 to go to sleep more frequently. Conversely, the controller 110 can increase the time period when the measured amount of time is more than the time period T (or more than T plus t1 and t2). Of course, other algorithms can be used, e.g., depending on the need of the memory system 100 for performance and power optimization.

[0041] Measuring the idle time and dynamically self-adjusting the time period T after which the controller 110 enters an auto-sleep mode can be performed in any suitable way. For example, in one embodiment, idle time is measured as the amount of time between when the memory 120 completes an operation and when the controller 110 receives a command to perform another operation in the memory 120.

[0042] In one exemplary implementation, the components in Figure 1 are used, and Figure 4 is a flowchart 400 that illustrates the use of those components. As shown in Figure 1, interrupt circuitry 130 is configured to receive a host interrupt via the host interface 115 and an interrupt from the memory 120 via the memory interface 125. In one embodiment, the host interface 115 receives commands from the host and provides a transaction start signal as an interrupt to the interrupt circuitry 130 in response to receiving every new command. (An interrupt from the host 50 can also prompt the controller 110 to exit sleep mode.) Similarly, the memory interface 125 receives a transaction end (end-of-operation (EOP)) signal from the memory 120 with every completion of a transaction in the memory 120.

[0043] As shown in Figure 4, in this embodiment, when the interrupt circuitry 130 receives a transaction end signal from the memory interface 125 (act 410), the interrupt circuitry 130 activates the clock 135 (act 420), which can be a hardware component that serves as a real time or non-real time clock. The clock measures time (act 430), and the controller 110 determines if a transaction start signal is received from a host (act 440). If the transaction start signal is received, the interrupt circuitry 130 stops the clock 135 (act 450). In this way, the clock 135 measures the amount of time between when the memory 120 completes an operation to when the controller 110 receives a command to perform another operation in the memory 120.

[0044] This measured time is then provided to the self-learning circuitry 140 (act 460), which uses this measured time to adjust an auto-sleep threshold (i.e., a time period  $T$  after which the controller 110 enters an auto-sleep mode based on the measured amount of time) (act 470). The self-learning circuitry 140 can make these adjustments after the completion of every transaction (associated with one or more commands) or after a threshold (e.g., after  $N$  transactions). Also, the controller 110 can periodically or occasionally initiate a self-adjusting process associated with the operation of the controller 110 for adjusting (increasing or decreasing) the time period  $T$  based on input from the self-learning circuitry 140. Accordingly, self-adjusting may be performed by the controller 110 in a dynamic manner. Also, the self-learning circuitry 140 can be used only when the memory system 100 is in a certain mode. In this way, the adaptive changing of the auto-sleep threshold can be performed only some of the time.

[0045] Figure 5 is a graph that illustrates this embodiment. Figure 5 is a graph of current consumed by the memory device 100 versus time. As shown in Figure 5, when the memory device 100 is in sleep mode (during time  $T_s$ ), the memory device 100 is consuming a relatively-low level of current (in this example, 6.0 mA). When the memory device 100 receives a command from a host (e.g., when the controller 110 receives a transaction start signal), the memory device 100 spends  $t_1$  time coming out of sleep mode and  $T_{act}$  time performing the command. When the memory device 100 is out of auto-sleep mode, it consumes a relatively-higher level of current (in this example, 106 mA). The controller 110 then waits for  $T$  time to pass before putting the memory device 100 back into auto-sleep mode, which requires  $t_2$  time.

[0046] In this embodiment, the self-learning circuitry 140 in the controller 110 monitors the actual time that passes since the controller 110 completed the last operation for an existing command and until receiving an indication for a new request. The time since the controller 110 completes the operation for an existing command until it receives an indication from the host for a new request is represented as  $T_d = T_s + t_2 + T_{act} + t_1$  and is monitored by the self-learning circuitry 140 on a regular basis and updated accordingly. For example, the self-learning circuitry 140 can measure the time frame ( $T_d$ ) since the controller 110 finishes its activity and until it gets an indication from the host for the next activity. If this time  $T_d$  is much smaller (say by a threshold  $T_a$ , which can be configured by the controller 110) than the time  $T + t_1 + t_2$  for  $N$  number of transactions ( $N$  can be a defined number of transactions and can be associated with one or more commands), then  $T$  can be adjusted (reduced) as much as possible, provided that, in this particular embodiment,  $T + t_1 + t_2 > T_d$ . Similarly, if  $T_d > T + T_1 + T_2$ , then the time  $T$  can be adjusted (increased) such that  $T + t_1 + t_2 > T_d$ .

[0047] There are several alternatives that can be used with these embodiments. For example, the self-learning mechanism discussed above for adjusting an auto-sleep threshold can instead or additionally be used for adjusting a time period after which the controller performs a background operation. A background operation is in contrast to a foreground operation, which is an operation that is preferably performed to completion without being interrupted and typically has a predetermined completion time. A foreground command is typically issued by the host controller and sent to the memory system 100 for execution. Examples of a foreground command include, but are not limited to, a read command, a write command, and an erase command.

[0048] In contrast, a background operation refers to an operation that can be preempted/interrupted by a foreground command and continued at a later time. Unlike foreground commands which typically have a predetermined completion time, background commands can have an uncertain completion time, depending on how much they are interrupted by foreground commands. A background command can be performed whenever possible when no host foreground commands are currently being executed and after a period of idle time. Examples of background operations include, but are not limited to, data scrambling, column replacement, handling write aborts and/or

program failures (via safe zones), read scrubbing, wear leveling, bad block and/or spare block management, error detection code (EDC) functionality, status functionality, encryption functionality, error recovery, and address mapping (e.g., mapping of logical to physical blocks).

[0049] In the same manner as described above, if  $T_d$  (the time frame since the controller 110 finishes its activity and until it gets an indication from the host for the next activity) is smaller than the time  $T_b+t_1+t_2$  for  $N$  number of transactions, then  $T_b$  (the time after which a background operation can be performed) can be adjusted (reduced) as much as possible, preferably provided that  $T_b+t_1+t_2 > T_d$ . Accordingly, if  $T_d > T_b+T_1+T_2$ , then the time  $T$  can be adjusted (increased) such that  $T+t_1+t_2 > T_d$ .

[0050] The flow chart 600 in Figure 6 illustrates this embodiment. As shown in the flow chart 600 in Figure 6, in performing adaptive background operations, the controller 110 measures idle time of the memory 110 (act 610) and determines if the measured idle time is less than or greater than a background operation threshold time after which the controller 110 performs a background operation (act 620). If the measured idle time is less than the threshold time, the controller 110 decreases the background operation threshold time (act 630). However, if the measured idle time is greater than the threshold time, the controller 110 increases the background operation threshold time (act 640).

[0051] Finally, as mentioned above, any suitable type of memory can be used. Semiconductor memory devices include volatile memory devices, such as dynamic random access memory ("DRAM") or static random access memory ("SRAM") devices, non-volatile memory devices, such as resistive random access memory ("ReRAM"), electrically erasable programmable read only memory ("EEPROM"), flash memory (which can also be considered a subset of EEPROM), ferroelectric random access memory ("FRAM"), and magnetoresistive random access memory ("MRAM"), and other semiconductor elements capable of storing information. Each type of memory device may have different configurations. For example, flash memory devices may be configured in a NAND or a NOR configuration.

[0052] The memory devices can be formed from passive and/or active elements, in any combinations. By way of non-limiting example, passive semiconductor memory elements include ReRAM device elements, which in some embodiments include a

resistivity switching storage element, such as an anti-fuse, phase change material, etc., and optionally a steering element, such as a diode, etc. Further by way of non-limiting example, active semiconductor memory elements include EEPROM and flash memory device elements, which in some embodiments include elements containing a charge storage region, such as a floating gate, conductive nanoparticles, or a charge storage dielectric material.

[0053] Multiple memory elements may be configured so that they are connected in series or so that each element is individually accessible. By way of non-limiting example, flash memory devices in a NAND configuration (NAND memory) typically contain memory elements connected in series. A NAND memory array may be configured so that the array is composed of multiple strings of memory in which a string is composed of multiple memory elements sharing a single bit line and accessed as a group. Alternatively, memory elements may be configured so that each element is individually accessible, e.g., a NOR memory array. NAND and NOR memory configurations are exemplary, and memory elements may be otherwise configured.

[0054] The semiconductor memory elements located within and/or over a substrate may be arranged in two or three dimensions, such as a two dimensional memory structure or a three dimensional memory structure.

[0055] In a two dimensional memory structure, the semiconductor memory elements are arranged in a single plane or a single memory device level. Typically, in a two dimensional memory structure, memory elements are arranged in a plane (e.g., in an x-z direction plane) which extends substantially parallel to a major surface of a substrate that supports the memory elements. The substrate may be a wafer over or in which the layer of the memory elements are formed or it may be a carrier substrate which is attached to the memory elements after they are formed. As a non-limiting example, the substrate may include a semiconductor such as silicon.

[0056] The memory elements may be arranged in the single memory device level in an ordered array, such as in a plurality of rows and/or columns. However, the memory elements may be arrayed in non-regular or non-orthogonal configurations. The memory elements may each have two or more electrodes or contact lines, such as bit lines and word lines.

[0057] A three dimensional memory array is arranged so that memory elements occupy multiple planes or multiple memory device levels, thereby forming a structure in three dimensions (i.e., in the x, y and z directions, where the y direction is substantially perpendicular and the x and z directions are substantially parallel to the major surface of the substrate).

[0058] As a non-limiting example, a three dimensional memory structure may be vertically arranged as a stack of multiple two dimensional memory device levels. As another non-limiting example, a three dimensional memory array may be arranged as multiple vertical columns (e.g., columns extending substantially perpendicular to the major surface of the substrate, i.e., in the y direction) with each column having multiple memory elements in each column. The columns may be arranged in a two dimensional configuration, e.g., in an x-z plane, resulting in a three dimensional arrangement of memory elements with elements on multiple vertically stacked memory planes. Other configurations of memory elements in three dimensions can also constitute a three dimensional memory array.

[0059] By way of non-limiting example, in a three dimensional NAND memory array, the memory elements may be coupled together to form a NAND string within a single horizontal (e.g., x-z) memory device levels. Alternatively, the memory elements may be coupled together to form a vertical NAND string that traverses across multiple horizontal memory device levels. Other three dimensional configurations can be envisioned wherein some NAND strings contain memory elements in a single memory level while other strings contain memory elements which span through multiple memory levels. Three dimensional memory arrays may also be designed in a NOR configuration and in a ReRAM configuration.

[0060] Typically, in a monolithic three dimensional memory array, one or more memory device levels are formed above a single substrate. Optionally, the monolithic three dimensional memory array may also have one or more memory layers at least partially within the single substrate. As a non-limiting example, the substrate may include a semiconductor such as silicon. In a monolithic three dimensional array, the layers constituting each memory device level of the array are typically formed on the layers of the underlying memory device levels of the array. However, layers of adjacent

memory device levels of a monolithic three dimensional memory array may be shared or have intervening layers between memory device levels.

[0061] Then again, two dimensional arrays may be formed separately and then packaged together to form a non-monolithic memory device having multiple layers of memory. For example, non-monolithic stacked memories can be constructed by forming memory levels on separate substrates and then stacking the memory levels atop each other. The substrates may be thinned or removed from the memory device levels before stacking, but as the memory device levels are initially formed over separate substrates, the resulting memory arrays are not monolithic three dimensional memory arrays. Further, multiple two dimensional memory arrays or three dimensional memory arrays (monolithic or non-monolithic) may be formed on separate chips and then packaged together to form a stacked-chip memory device.

[0062] Associated circuitry is typically required for operation of the memory elements and for communication with the memory elements. As non-limiting examples, memory devices may have circuitry used for controlling and driving memory elements to accomplish functions such as programming and reading. This associated circuitry may be on the same substrate as the memory elements and/or on a separate substrate. For example, a controller for memory read-write operations may be located on a separate controller chip and/or on the same substrate as the memory elements.

[0063] One of skill in the art will recognize that this invention is not limited to the two dimensional and three dimensional exemplary structures described but cover all relevant memory structures within the spirit and scope of the invention as described herein and as understood by one of skill in the art.

[0064] It is intended that the foregoing detailed description be understood as an illustration of selected forms that the invention can take and not as a definition of the invention. It is only the following claims, including all equivalents, that are intended to define the scope of the claimed invention. Finally, it should be noted that any aspect of any of the preferred embodiments described herein can be used alone or in combination with one another.

What is claimed is:

1. A method for performing adaptive auto-sleep in a memory system, the method comprising:  
performing the following in a controller of a memory system having a memory:  
measuring an amount of time between when the memory completes an operation and when the controller receives a command to perform another operation in the memory; and  
adjusting a time period after which the controller enters an auto-sleep mode based on the measured amount of time.
2. The method of Claim 1, wherein the time period is adjusted based on the measured amount of time and on at least one additional amount of time for at least one additional operation.
3. The method of Claim 1, wherein the time period is reduced when the measured amount of time is less than the time period.
4. The method of Claim 1, wherein the time period is reduced when the measured amount of time is less than a sum of the time period and a time it takes the controller to enter and exit the auto-sleep mode.
5. The method of Claim 1, wherein the time period is increased when the measured amount of time is greater than the time period.

6. The method of Claim 1, wherein the time period is increased when the measured amount of time is greater than a sum of the time period and a time it takes the controller to enter and exit the auto-sleep mode.
7. The method of Claim 1, wherein the controller measures the amount of time by measuring time between when the controller receives an end signal via a memory interface and when the controller received a start signal via a host interface.
8. The method of Claim 1, wherein the memory is a three-dimensional memory.
9. The method of Claim 1, wherein the memory system is embedded in a host.
10. The method of Claim 1, wherein the memory system is removably connected to a host.
11. A memory system comprising:
  - a memory; and
  - a controller in communication with the memory, wherein the controller is configured to:
    - measure idle time of the memory; and
    - determine if the measured idle time of the memory is less than or greater than an auto-sleep threshold;

in response to determining that the measured idle time is less than the auto-sleep threshold, decrease the auto-sleep threshold; and

in response to determining that the measured idle time is greater than the auto-sleep threshold, increase the auto-sleep threshold.

12. The memory system of Claim 11, wherein the controller is configured to measure idle time by measuring an amount of time between when the controller completes an operation in the memory and when the controller receives a command to perform another operation in the memory.

13. The memory system of Claim 12, wherein the controller is configured to measure the amount of time by measuring time between when the controller receives an end signal via a memory interface and when the controller received a start signal via a host interface.

14. The memory system of Claim 11, wherein the controller is configured to use a plurality of measured idle times to determine if the auto-sleep threshold should be increased or decreased.

15. The memory system of Claim 11, wherein the auto-sleep threshold includes a time it takes the controller to enter and exit the auto-sleep mode.

16. The memory system of Claim 11, wherein the controller is configured to measure idle time using interrupt circuitry and a clock, and wherein the controller is further

configured to determine if the measured idle time of the controller is less than or greater than the auto-sleep threshold using self-learning circuitry.

17. The memory system of Claim 1, wherein the memory is a three-dimensional memory.

18. The memory system of Claim 1, wherein the memory system is embedded in a host.

19. The memory system of Claim 1, wherein the memory system is removably connected to a host.

20. A method for performing adaptive background operations in a memory system, the method comprising:

performing the following in a controller of a memory system having a memory:

measuring an amount of time between when the controller completes an operation in the memory and when the controller receives a command to perform another operation in the memory; and

adjusting a time period after which the controller performs a background operation based on the measured amount of time.

21. The method of Claim 20, wherein the time period is adjusted based on the measured amount of time and on at least one additional amount of time for at least one additional operation.
22. The method of Claim 20, wherein the time period is reduced when the measured amount of time is less than the time period.
23. The method of Claim 20, wherein the time period is increased when the measured amount of time is greater than the time period.
24. The method of Claim 20, wherein the controller measures the amount of time by measuring time between when the controller receives an end signal via a memory interface and when the controller received a start signal via a host interface.
25. The method of Claim 20, wherein the memory is a three-dimensional memory.
26. The method of Claim 20, wherein the memory system is embedded in a host.
27. The method of Claim 20, wherein the memory system is removably connected to a host.
28. A memory system comprising:  
a memory; and

a controller in communication with the memory, wherein the controller is configured to:

measure idle time of the memory; and

determine if the measured idle time is less than or greater than a background operation threshold time after which the controller performs a background operation;

in response to determining that the measured idle time is less than the background operation threshold time, decrease the background operation threshold time; and

in response to determining that the measured idle time is greater than the threshold time, increase the background operation threshold time.

29. The memory system of Claim 28, wherein the controller is configured to measure idle time by measuring an amount of time between when the controller completes an operation in the memory and when the controller receives a command to perform another operation in the memory.

30. The memory system of Claim 29, wherein the controller is configured to measure the amount of time by measuring time between when the controller receives an end signal via a memory interface and when the controller received a start signal via a host interface.

31. The memory system of Claim 28, wherein the controller is configured to use a plurality of measured idle times to determine if the threshold should be increased or decreased.

32. The memory system of Claim 28, wherein the controller is configured to measure idle time using interrupt circuitry and a clock, and wherein the controller is further configured to determine if the measured idle time of the controller is less than or greater than the threshold using self-learning circuitry.

33. The memory system of Claim 28, wherein the memory is a three-dimensional memory.

34. The memory system of Claim 28, wherein the memory system is embedded in a host.

35. The memory system of Claim 28, wherein the memory system is removably connected to a host.




Figure 1

2/6




Figure 2A




Figure 2B

3/6




Figure 3

4/6




Figure 4




Figure 5

6/6




Figure 6

## INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2016/036354

A. CLASSIFICATION OF SUBJECT MATTER  
INV. G06F1/32 G06F13/16  
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)  
G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, WPI Data

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages                                                                                                                                       | Relevant to claim No. |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| X         | US 2013/290758 A1 (QUICK CHRISTOPHER [US] ET AL) 31 October 2013 (2013-10-31)<br>paragraphs [0045], [0048] - [0050],<br>[0058], [0059]<br>paragraph [0087]; figure 5<br>paragraph [0108]; figure 12<br>-----             | 1-35                  |
| X         | US 2005/198542 A1 (FREKER DAVID [US] ET AL) 8 September 2005 (2005-09-08)<br>paragraphs [0013], [0014]<br>paragraph [0017]; figure 1<br>paragraph [0033] - paragraph [0045];<br>figures 3-5<br>paragraph [0051]<br>----- | 1-19<br>20-35         |
| A         | -----<br>-/-                                                                                                                                                                                                             |                       |



Further documents are listed in the continuation of Box C.



See patent family annex.

\* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

30 August 2016

Date of mailing of the international search report

05/09/2016

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2  
NL - 2280 HV Rijswijk  
Tel. (+31-70) 340-2040,  
Fax: (+31-70) 340-3016

Authorized officer

Pascual Vallés, E

## INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2016/036354

| C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT |                                                                                                                                                                                                                                                                                                              |                       |
|------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| Category*                                            | Citation of document, with indication, where appropriate, of the relevant passages                                                                                                                                                                                                                           | Relevant to claim No. |
| A                                                    | US 2012/159209 A1 (STEMEN PATRICK L [US]<br>ET AL) 21 June 2012 (2012-06-21)<br>abstract<br>paragraph [0014] - paragraph [0021];<br>figure 1<br>-----                                                                                                                                                        | 20-35                 |
| A                                                    | US 2012/151238 A1 (YANG JUNG WOONG [KR])<br>14 June 2012 (2012-06-14)<br>paragraph [0006]<br>paragraph [0028] - paragraph [0032];<br>figures 2,3<br>-----                                                                                                                                                    | 1-35                  |
| A                                                    | WO 2010/080172 A1 (RAMBUS INC [US]; WARE<br>FREDERICK A [US]; PALMER ROBERT E [US];<br>POULTON) 15 July 2010 (2010-07-15)<br>paragraph [0053] - paragraph [0054];<br>figures 1A-1B<br>paragraph [0119] - paragraph [0121];<br>figure 12A<br>paragraph [0153] - paragraph [0160];<br>figures 18A-18D<br>----- | 1-35                  |

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2016/036354

| Patent document<br>cited in search report | Publication<br>date | Patent family<br>member(s) | Publication<br>date                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------------------------|---------------------|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| US 2013290758                             | A1                  | 31-10-2013                 | NONE                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| US 2005198542                             | A1                  | 08-09-2005                 | NONE                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| US 2012159209                             | A1                  | 21-06-2012                 | CN 102567110 A 11-07-2012<br>US 2012159209 A1 21-06-2012<br>US 2015286503 A1 08-10-2015                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| US 2012151238                             | A1                  | 14-06-2012                 | KR 20120063828 A 18-06-2012<br>US 2012151238 A1 14-06-2012                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| WO 2010080172                             | A1                  | 15-07-2010                 | CN 102257569 A 23-11-2011<br>CN 102257571 A 23-11-2011<br>CN 102257572 A 23-11-2011<br>EP 2351037 A1 03-08-2011<br>EP 2353058 A1 10-08-2011<br>EP 2356657 A1 17-08-2011<br>EP 2370975 A1 05-10-2011<br>JP 5627603 B2 19-11-2014<br>JP 5855726 B2 09-02-2016<br>JP 2012515376 A 05-07-2012<br>JP 2012515377 A 05-07-2012<br>JP 2012515408 A 05-07-2012<br>JP 2015038802 A 26-02-2015<br>KR 20110113736 A 18-10-2011<br>KR 20110113737 A 18-10-2011<br>KR 20110113738 A 18-10-2011<br>US 2011235459 A1 29-09-2011<br>US 2011235763 A1 29-09-2011<br>US 2011235764 A1 29-09-2011<br>US 2011239030 A1 29-09-2011<br>US 2011239031 A1 29-09-2011<br>US 2015074437 A1 12-03-2015<br>US 2015227188 A1 13-08-2015<br>US 2016147281 A1 26-05-2016<br>WO 2010080172 A1 15-07-2010<br>WO 2010080173 A1 15-07-2010<br>WO 2010080174 A1 15-07-2010<br>WO 2010080175 A1 15-07-2010<br>WO 2010080176 A1 15-07-2010 |