



- (51) **International Patent Classification:**  
*G06F 12/08* (2006.01) *G06F 12/02* (2006.01)
- (21) **International Application Number:**  
PCT/US2013/071753
- (22) **International Filing Date:**  
25 November 2013 (25.11.2013)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**  
13/730,211 28 December 2012 (28.12.2012) US
- (71) **Applicant:** APPLE INC. [US/US]; 1 Infinite Loop, Cupertino, California 95014 (US).
- (72) **Inventor:** SOKOL, Joseph, Jr.; 1 Infinite Loop, M/S 37-1COS, Cupertino, California 95014 (US).
- (74) **Agents:** SCHELLER, James, C. et al.; Blakely, Sokoloff, Taylor & Zafinan LLP, 1279 Oakmead Parkway, Sunnyvale, California 94085 (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, LT, 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, 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).

[Continued on next page]

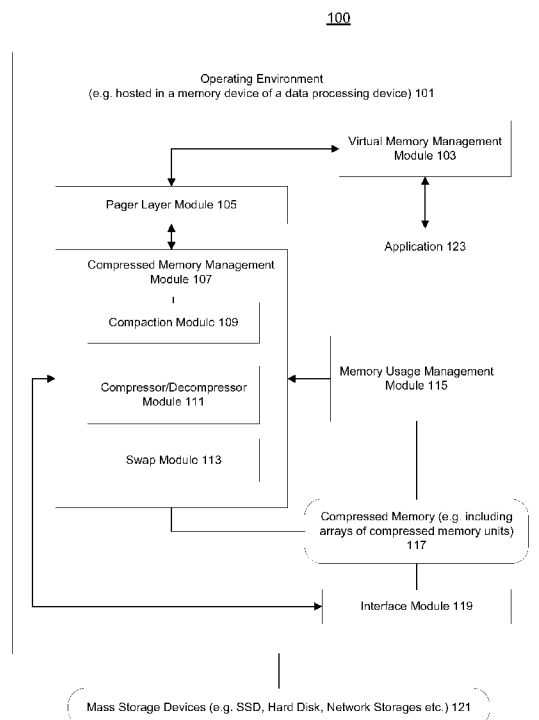
(54) **Title:** METHODS AND APPARATUS FOR COMPRESSED AND COMPACTED VIRTUAL MEMORY

Fig. 1

(57) **Abstract:** A method and an apparatus for a memory device including a dynamically updated portion of compressed memory for a virtual memory are described. The memory device can include an uncompressed portion of memory separate from the compressed portion of memory. The virtual memory may be capable of mapping a memory address to the compressed portion of memory. A memory region allocated in the uncompressed portion of memory can be compressed into the compressed portion of memory. As a result, the memory region can become available (e.g. after being compressed) for future allocation requested in the memory device. The compressed portion of memory may be updated to store the compressed memory region. The compressed memory region may be decompressed back to the uncompressed portion in the memory device in response to a request to access data in the compressed memory region.



---

**Published:**

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

— before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

## METHODS AND APPARATUS FOR COMPRESSED AND COMPACTED VIRTUAL MEMORY

### **FIELD OF INVENTION**

[0001] The present invention relates generally to memory systems. More particularly, this invention relates to compressed and/or compacted memory for virtual memory systems.

### **BACKGROUND**

[0002] Virtualization techniques have been developed to enhance memory management for memory space having a size exceeding actual physical capacity of a memory device. Typically, memory virtualization can be based on memory swapping using a storage device coupled to the memory device. As memory objects and/or IO (input/output) sizes are not deterministic, different sizes or pages of memory which are uncompressed may be swapped in/out between the storage and memory devices. As a result, IO throughput and/or latency for transferring stored memory data for the swap memory may negatively impact the performance of the virtual memory system.

[0003] Further, the performance of certain storage devices, such as spin based hard disks, may have strong dependency on data locality for memory access to reduce lengthy disk spins. However, data locality may not be guaranteed during runtime as clustering characteristics for memory access vary in different data processing tasks. Thus, memory swapping by paging in/out between the memory device and the storage device to support generic data processing tasks can further degrade the perceived performance of the memory system.

[0004] Although the trend of adopting larger and larger sizes of actual physical memory tends to alleviate performance cost in inter device memory swap, certain feature enhancements in data processing systems may be provided without memory size changes. For example, support for higher screen resolution may be imposed on existing devices based on substantially the same amount of memory size. However, a mere two times higher screen resolution may correspond to four times increase in window buffer sizes. As a result, without a compensating increase in physical memory size, the system performance for higher screen resolution may degrade and become visibly slower.

[0005] Therefore, traditional implementations of virtual memory systems are not capable of supporting performance requirements constrained by limited physical memory

sizes.

## **SUMMARY OF THE DESCRIPTION**

[0006] A compressed memory pool dynamically maintained in a memory device can provide an additional layer of support for a virtual memory based on the memory device coupled with a mass storage device. Virtual memory addresses may be paged or mapped into an uncompressed portion of the memory device, the compressed memory pool and/or the storage device. The compressed memory pool may grow and shrink within the memory device without a size limitation for the virtual memory.

[0007] In one embodiment, a memory page (e.g. a unit of area of uncompressed memory) may be compressed into a compressed memory unit (or compressed page) in the compressed memory. Multiple compressed memory units may be compacted together as a fixed sized segment in the compressed memory. The segment may be sized for efficient data transfer between the memory device and the mass storage device for swapping in/out segments of compressed memory units to support a virtual memory pager.

[0008] In another embodiment, a memory device may include a dynamically updated portion of compressed memory for a virtual memory. The memory device can include an uncompressed portion of memory separate from the compressed portion of memory. The virtual memory may be capable of mapping a memory address to the compressed portion of memory. A memory region allocated in the uncompressed portion of memory can be compressed into the compressed portion of memory. As a result, the memory region can become available (e.g. after being compressed) for future allocation requested in the memory device. The compressed portion of memory may be updated to store the compressed memory region. The compressed memory region may be decompressed back to the uncompressed portion in the memory device in response to a request to access data in the compressed memory region.

[0009] In another embodiment, a virtual memory based on a memory device is provided. The memory device can be dynamically partitioned into an uncompressed portion of memory and a compressed portion of memory, such as DRAM (volatile memory). The uncompressed portion of memory can store working data processed via a processor coupled to the memory device. One or more pages of the uncompressed portion of memory can be compressed into one or more varied sized compressed memory units in the compressed portion of memory. The compression can increase available memory space in the uncompressed portion of memory. The compressed memory units

can be decompressed back from the compressed portion of memory to the uncompressed portion of memory, for example, in response to a page fault for an access request to data in one of the compressed memory units. At least one of the varied sized compressed memory units can be swapped out from the compressed portion of memory to a mass storage device to increase available memory space in the uncompressed portion of memory.

[0010] In another embodiment, one or more memory pages of an uncompressed portion of a memory device can be compressed into one or more varied sized compressed memory units in a compressed portion of the memory device. A mass storage device (such as a magnetic hard drive (HDD) or a composite storage device, treated as one logical volume from a file system's perspective, that includes an HDD and flash memory) can be coupled with the memory device. The varied sized compressed memory units can be compacted into a segment in the compressed portion of the memory device. The one compressed memory unit can be decompressed from the compressed portion to the uncompressed portion of the memory device in response to an access request to one of the compressed memory units. The segment of the compressed memory units may be swapped out to the mass storage device to increase available memory space in the memory device.

[0011] Other features of the present invention will be apparent from the accompanying drawings and from the detailed description that follows.

### **BRIEF DESCRIPTION OF THE DRAWINGS**

[0012] The present invention is illustrated by way of example and not limitation in the figures of the accompanying drawings, in which like references indicate similar elements and in which:

[0013] Figure 1 is a block diagram illustrating one embodiment of a virtual memory system based on compressed and/or compacted memory;

[0014] Figure 2 is a flow diagram illustrating one embodiment of compressing and decompressing memory for a virtual memory;

[0015] Figure 3 is a flow diagram illustrating one embodiment of a process to swap compressed memory pages for a virtual memory;

[0016] Figure 4 is a flow diagram illustrating one embodiment of a process to swap a compacted segment of compressed memory for a virtual memory;

[0017] Figure 5 illustrates one example of a data processing system such as a computer

system, which may be used in conjunction with the embodiments described herein.

## **DETAILED DESCRIPTION**

[0018] Methods and apparatuses for compressed and compacted virtual memory are described herein. In the following description, numerous specific details are set forth to provide thorough explanation of embodiments of the present invention. It will be apparent, however, to one skilled in the art, that embodiments of the present invention may be practiced without these specific details. In other instances, well-known components, structures, and techniques have not been shown in detail in order not to obscure the understanding of this description.

[0019] Reference in the specification to “one embodiment” or “an embodiment” means that a particular feature, structure, or characteristic described in connection with the embodiment can be included in at least one embodiment of the invention. The appearances of the phrase “in one embodiment” in various places in the specification do not necessarily all refer to the same embodiment.

[0020] The processes depicted in the figures that follow, are performed by processing logic that comprises hardware (e.g., circuitry, dedicated logic, etc.), software (such as is run on a general-purpose computer system or a dedicated machine), or a combination of both. Although the processes are described below in terms of some sequential operations, it should be appreciated that some of the operations described may be performed in different order. Moreover, some operations may be performed in parallel rather than sequentially.

[0021] In one embodiment, a compression pool (or pool of compressed memory units) in a memory device may be managed dynamically using log based data structures (or log structures). Data can be put in the compression pool in a packed and/or compressed manner to shrink required memory space for storing the data. When to swap stored data between the compression pool and a mass storage device (or nonvolatile storage device) may be automatically determined in a self contained compressed memory management system, for example, to insulate a virtual memory system or other applicable memory usage module, from the underlying compression pool.

[0022] Memory data may be compressed to reduce swap file size, for example, for data transfer efficiency between a memory device and a swap device such as a mass storage device. Additionally, memory data may be compressed into a compression pool based on efficient memory compression schemes with compression/decompression cost (e.g. in

several micro seconds) significantly less (e.g. several orders of magnitudes) than conventional disk swap speed with a swap device. The compression pool may grow in size within the memory device when needed to reduce required swapping operations with the swap device.

**[0023]** For example, the compression pool can grow and work as large as the available physical memory size as long as there is enough memory space left (e.g. a small number of memory pages) to perform memory compression/decompression operations to shift data in/out of the compression pool. A working set may be maintained to fit in the uncompressed portion of memory with varied sizes as the compression pool grows/shrinks. If the working set (e.g. memory allocated for a set of live tasks for data operations of a processor) spills into the compression pool (or compressed portion of memory), memory pages may be pulled (e.g. based on a certain predetermined rate) from the compressed portion to substitute between the compressed and uncompressed portions (or layers) of memory to support the working set. Requests or requirements for additional memory space to allocate may be detected to indicate that the existing working set has spilled over. The working set may be extended via a LRU (least recently used) ordering strategy to select which memory pages to be compressed into the compression pool.

**[0024]** In one embodiment, multiple compression/decompression schemes or mechanisms may be dynamically selected during run time. Each scheme may be based on one or more compression/decompression algorithms. The selection may be made based on a hint or heuristic provided by the creator (e.g. an application or process initiating or requesting allocation) of the memory, or it could be determined by looking in each page of memory about to be compressed for a particular pattern of data. In one embodiment, one of the compression/decompression schemes may be designated as the default one applicable to a majority of allocated memories (e.g. without any hint provided by corresponding memory creators). This scheme would also be used as the default one if a pattern matching mechanism was in play, and no match was made. One candidate that might benefit (e.g. see more efficient compression) from a specialized compression/decompression scheme is the memory created for use by graphics operations (e.g. via a GPU).

**[0025]** In some embodiments, compaction operations may be performed in the compression pool to maintain log structures in the compressed memory segments. For example, compressed data units may be packed together within one compression memory segment (e.g. for minor compaction) and/or across multiple compression memory

segments (e.g. for major compaction). Compaction operations may be initiated on requests and/or automatically as compressed data are removed from or stored into the compressed memory. Log structures may allow efficient input/output operations (e.g. read/write operations) on the compression pool within a memory device and/or between coupling storage devices.

**[0026]** According to one embodiment, as compressed data is pulled out from the compression pool, holes (i.e. empty memory space previously occupied by the compressed data) may be created. As a result, memory usage of the compression pool may become less and less efficient if these holes are not removed. Minor compaction operations can be performed to move compressed data within a memory segment (e.g. a log structure) in the compression pool to reduce number of holes in the memory segment. In one embodiment, minor compaction operations may be performed to allow additional allocation of memory pages in the memory device. Minor compaction operations may cost insignificant processing resources to cause noticeable improvement in data processing latency.

**[0027]** In one embodiment, major compaction operations may be performed to move compressed data among different memory segments of the compression pool in addition to minor compaction operations, for example, when there is a need to swap compressed data out to a swap device. The number of memory segments required in the compression pool may be reduced as a result of the major compaction operations. Unused memory space in memory segments of the compression pool may be tightened up (or reduced), for example, to allow writing out large chunks of compressed memory pages to a swap device. As a result, available IO bandwidth (e.g. between the memory device and the swap device) can be maximized. Further, the amount of wear and tear against the swap device (such as a SSD device) can be minimized.

**[0028]** In one embodiment, compaction operations are performed based on a compaction model which determines when to perform minor compaction operations, major compaction operations or other applicable data movement operations to maintain efficient memory storage usage in the compression pool. Minor compaction operations may be performed to shuffle compressed data within one memory segment to increase or release available memory space. Major compaction operations may be performed to move compressed data among multiple memory segments, for example, to ensure readiness to swap out memory segments in full size from the compression pool to a swap device. A memory segment may be in full size if it is substantially without holes or with a memory



usage rate greater than, for example, a predetermined threshold (e.g. 99%).

[0029] Thus, log structured data (or model) can be maintained when performing IO operations writing to swap files (e.g. in a swap device) from the compression pool, for example, based on same sized chunk of data (e.g. a compression segment), such as 256KB, 512KB, 1024KB or other suitable size. In one embodiment, the chunk of data transferred during the IO operation for memory swapping may correspond to a collection of compression memory pages according to the IO interface with the swap device (e.g. matching an IO bus buffer size) to improve IO performance. Consequently, efficiency or through put in IO operations between the compression pool and the swap device may be improved based on the compression of the data transferred and the matching transfer data chunk size for the swap device.

[0030] In one embodiment, log structured data in the compression pool may include an ordered list of memory segments. Newly compressed data may be recorded in the head segment (or recording head) of the ordered segment list. Working set analysis for a virtual memory system may be performed to determine which portion of the compression pool to swap out based on the ordered segment lists in the compression pool. For example, the farther away a segment is located in the ordered list from the head segment, the more likely compressed data stored in the segment may no longer be of interest (or of immediate need) for the working set. As a result, the tail segment (or far end from the recording head) may be selected to be swapped out to a swap device. Compressed data may be brought in/out from the compression pool from/to the swap device based on compacted memory segments.

[0031] A need to perform memory swap operations may be identified, for example, to put an upper limit on the size of the compression pool. The size (or shape) of the compression pool may change dynamically based on runtime memory access patterns. In one embodiment, processing cost of compaction/compression/decompression operations vs. memory swap operations may be dynamically monitored (or estimated) based on a cost (e.g. power consumption cost) model to determine when to perform swap operations.

[0032] For example, the rate of compression, decompression and/or compaction operations may correspond to power (or processing resource consumption) to indicate health of the compression pool (or compression layer, compression portion) in the memory device. If this rate exceeds a certain threshold, swap operations may be performed or triggered to page out compressed data out of the compression pool to a swap device. In one embodiment, compression/decompression and/or compaction operations

may incur a fraction of processing cost compared with data IO transactions in swap operations. Power cost functions may be evaluated to allow the compression pool to grow before a need arises for paging out compressed data to a swap device. The cost function may depend on, for example, the shape of the compression pool, configured tolerance in processing latency caused by compaction/compression/decompression operations, or other applicable factors to maintain the working set in the memory device.

**[0033]** The compression pool in a memory device may allow power savings by reducing the required number of data paging (e.g. swap operations) to/from a swap device. For example, as data is compressed, less number of data paging or file paging is required. Further, efficiency or effectiveness in each data paging or swap operation can be improved with compacted segments carrying the compressed data via log based structures.

**[0034]** In some embodiments, lists of log structure based memory segments in the compression pool may be ordered corresponding to an age of compressed data stored in each segment. The age of the compressed data may indicate how early the data was compressed from the working set or accessed for the working set. Swap operations may write out memory segments which have been aged (e.g. the tail portions of the ordered segment list) out of the working set. As a result, compressed data selected to be swapped out of the compression pool to a swap device (e.g. disk) may be stored in a compacted memory segment which hasn't been used for a while. Swap operations may be performed to, for example, make room for or accommodate growth of the working set. In the meanwhile, periodic data write from the memory device to the swap device or disk device for maintenance purposes (e.g. file or data backup/synchronization) may be avoided or reduced with significant power saving.

**[0035]** The compression pool may provide mechanisms for speedy operations for system hibernation and/or waking up. For example, the compressed portion and uncompressed portion of a memory device may be separately backed up (or stored) in a non-volatile storage device (e.g. a disk) for a system to enter a hibernated (or sleep) state. The compressed portion (or compression pool) may be transferred (or pushed) to the non-volatile storage device via compacted memory segments. The uncompressed portion of the memory (e.g. including a current working set) may be compressed into the compression pool attached with special tags to be transferred to the non-volatile storage device as compacted memory segments. For example, applications may be suspended into the compression pool as compressed/compacted memory segments.

[0036] During system waking up, the specially tagged compressed memory segments corresponding to the uncompressed portion of the memory device may be retrieved (or pulled back) from the non-volatile storage device to populate the memory device with the previous working set (or hot working set). For example, suspended applications may be unsuspended or restored back to the uncompressed portion of the memory device from the specially tagged compressed memory segments. Instead of retrieving arbitrary memory pages (e.g. previously stored in either the compressed or uncompressed portions of the memory) associated with the application being restored, the system may be pre-heated (or pre-warmed for accelerated waking up) with valuable information previously stored in the uncompressed portion of the memory for the application before retrieving those memory pages previously stored in the compression pool.

[0037] Figure 1 is a block diagram illustrating one embodiment of a virtual memory system based on compressed and/or compacted memory. For example, system 100 may include operating environment 101 hosted in a memory device of a data processing device, such as a mobile device, a desktop computer, or a server, etc. The memory device may comprise, for example, a random access memory (RAM) for supporting data processing of at least one processor in the processing device. Compressed memory (or compression pool) 117 may be dynamically partitioned (or maintained) within the memory device. As a result, available memory space in the memory device for supporting operations of operating environment 101, which may be based on the uncompressed portion of the memory device, may be less than the total memory space provided by the memory device. In certain embodiments, multiple processors, such as CPUs, GPUs or other applicable processing units, can access compressed memory 117 in parallel.

[0038] In one embodiment, virtual memory management module 103 may virtualize various forms of computer data storage, such as RAM and mass storage devices 121, to allow a program or process, such as application 123, running in operating environment 101 to access memory addresses (or virtual memory addresses) as though there is only one kind of memory, namely virtual memory, which can behave like directly addressable read/write memory in RAM. For example, virtual memory management module 103 can maintain mapping between virtual memory addresses (or a virtual memory address) and actual physical data storage locations, such as physical memory pages within the RAM (or memory device) and storage references in mass storage devices 121.

[0039] Virtual memory management module 103 may move data into or out of the memory device (e.g. RAM) via pager layer module 105. For example, data (e.g.

anonymous data or dirty data not backed by a disk file in mass storage devices 121) may be swapped out of physical memory pages in the memory device to free up memory space in the physical memory pages in response to, for example, memory allocation requests from applications and/or other data processing requirements. Clean data (or memory data) which may be available from a disk file in mass storage devices 121 may be paged out (or purged) from the memory device to increase available memory space. When an application, such as application 123, terminates, memory space (e.g. stacks, heaps, etc.) allocated for the application may become available for future allocation.

**[0040]** In one embodiment, virtual memory management module 103 may send memory swapping in/out requests to pager layer module 105 to manage virtual memory addresses within uncompressed portion of the memory device. Pager layer module 105 may interface with mass storage devices 121 and/or compressed memory management module 107 to move memory pages to/from the uncompressed portion of the memory device. As a result, virtual memory management module 103 can map a virtual memory address outside of the uncompressed portion of the memory device without distinguishing (or knowledge of) whether the virtual memory address corresponds to a memory location within compressed memory 117 or mass storage devices 121. In other words, virtual memory management module 103 can be ignorant of the presence of compressed memory 117.

**[0041]** Pager layer module 105 may be based on multiple concurrent threads (or input/output threads) to perform memory paging operations. For example, pager layer module 105 may include a file pager to manage uncompressed memory pages storing clean data from disk files of mass storage devices 121. Pager layer module 105 may include memory pager (e.g. as one or more threads) interfacing with compressed memory management module 107 to move uncompressed memory pages to/from the uncompressed portion of the memory device via compressed memory 117.

**[0042]** In one embodiment, compressed memory 117 may include one or more log structures (e.g. based on an array or queue having a head and a tail) to store compressed memory units. Each log structure may include one or more segments with buffers allocated for storing compressed memory units. Each segment may be of a common size (e.g. 256 KB, 512KB, 1024KB, or other applicable sizes) with memory slots to store an ordered list of compressed memory units (or compressed pages). Each compressed unit maybe may be compressed from a memory page in the uncompressed portion of a memory device, e.g. via compressor/decompressor module 111 to the head of a log structure. In some embodiments, each log structure can be updated or operated via

separate threads concurrently and/or independently via compressed memory management module 107.

**[0043]** A compressed memory unit may be stored in a memory slot (or position) in one segment of a log structure. Compressor module 111 can record or store newly compressed memory page to a head slot of a log structure and update the next slot (to the current head slot) as the head slot of the log structure for future compression. Compressed memory management module 107 may create (or allocate from the uncompressed portion of a memory device) one or more new segments when the current segment (or head segment in the log structure) becomes full (or substantially full). A mapping between a memory page and its corresponding compressed memory unit in a log structure may be maintained for compressed memory 117, for example, based on identifiers identifying a log structure, a segment within the identified log structure and a memory slot within the identified segment storing the compressed memory unit.

**[0044]** In one embodiment, compressor/decompressor module 111 can manage log structures in compressed memory 117. Each log structure may be associated with a recording head which may point to a head slot in a current segment, e.g. head segment in an array of segments in the log structure. Multiple threads may be associated with compressor/decompressor module 111 to perform memory compress/decompress operations using separate log structures concurrently. Compressor/decompressor module 111 may perform compressing operations according to a request queue storing memory pages to be compressed, for example, as scheduled via pager layer module 105. A memory page when compressed may require less than one page of physical memory size to be stored in a memory slot within a segment in a log structure.

**[0045]** In one embodiment, as more memory pages are compressed into compressed memory 117, memory slots may be filled up following the head of corresponding log structure. Alternatively, some segments or slots in the segments in the corresponding log structure may become empty or available for allocation when the stored compressed memory units are released. For example, when an application or process terminates, objects or structures maintained for the application may be removed and physical memory pages or compressed memory units owned by the application may be released for future use (i.e. available for allocation). Alternatively, segments may become empty as a result of being swapped out, decompressed or compacted.

**[0046]** Compressor/decompressor module 111 may pull or decompress compressed memory units out of compressed memory 117 to the uncompressed portion of a memory

device, for example, in response to page faults encountered in virtual memory management module 103. Compressor/decompressor module 111 may perform compressing and decompressing operations concurrently on different segments of compressed memory 117, for example, via separate threads.

**[0047]** Compressed memory management module 107 can include swap module 113 to move segments of compressed memory units between compressed memory 117 and mass storage device 121 via interface module 119. For example, segments of compressed memory units may be swapped out of compressed memory 117 to make room for accommodating newly compressed memory pages, and/or to store a state of the memory device for future recovery during system hibernation process etc. Alternatively, a segment of compressed memory units may be swapped into compressed memory 117 because of page faults encountered in virtual memory management module 103 to access data in one or more of the segment of compressed memory units.

**[0048]** In one embodiment, one or more threads may be running for swap module 113. A swap thread may wake up according to a designated schedule or responding to requests. For example, pager layer module 105 may issue requests to resolve page faults. Memory usage management module 115 may send request to maintain enough available memory space in the memory device. According to one embodiment, a swap thread can wake up to examine the current state of the memory device to determine how many segments of compressed memory units to swap out of compressed memory 117. The swap thread can select a least used (e.g. aged or old) segment, compact it (e.g. via compaction module 109) if needed, and forward the selected segment to mass storage devices 121. Swapping operations via swap module 113 may be performed independent and/or concurrently with compression/decompression operations via compressor/decompressor module 111.

**[0049]** Compaction module 109 may perform compaction operations on compressed memory 117 to reduce memory fragmentation and/or to increase memory swapping efficiencies. For example, compaction module 109 can perform minor compaction operations to slide (or move) compressed memory units within a segment of a log structure without altering ordering relationship among the compressed memory units to consolidate empty memory slots left in between. Alternatively or optionally, compaction module 109 can perform major compaction operations to move compressed memory units across different segments to fill up and/or empty certain segments.

**[0050]** In one embodiment, swap module 113 may request compaction module 109 to compact segments (e.g. minimize or reduce empty slots) for better swapping

performance. For example, the interface between a memory device and a swap device may be based on a fixed size data buffer for data transfer. A tightly packed segment matching the data buffer size can maximize transfer efficiency (e.g. amount of data transferred per unit time or clock cycle).

[0051] Memory usage management module 115 may determine when to push out or swap out compressed memory units or pages from compressed memory 117 to mass storage devices 121. Memory usage management module can monitor usage status of the memory device to provide triggers for (or send request to) swap module 113 to initiate swapping operations. Usage status may include power usage, memory performance (e.g. read/write speed) , memory usage rate, ratio of compressed/uncompressed partitions in the memory device, size of compressed memory 117, size of working set of memory pages, or other applicable runtime measures, etc.

[0052] Figure 2 is a flow diagram illustrating one embodiment of compressing and decompressing memory for a virtual memory. For example, process 200 may be performed by some components of system 100 of Figure 1. At block 201, the processing logic of process 200 can dynamically update a compressed portion of memory in a memory device for a virtual memory supporting one or more processors. The memory device can include an uncompressed portion of memory separate from the compressed portion of memory. A memory address in the virtual memory may be mapped into the compressed portion or the uncompressed portion of the memory.

[0053] In one embodiment, the compressed portion of memory may be updated with additional memory allocated from the uncompressed portion of memory. A memory region allocated in the uncompressed portion of memory may be compressed into the compressed portion of memory to allow the memory region to be available for future allocation in the uncompressed portion of memory of the memory device. The updated compressed portion of memory may be capable of storing the compressed memory region. The size of available memory for allocation in the uncompressed portion of memory may increase as a result of the compression of the memory region.

[0054] According to one embodiment, the compressed memory region may be decompressed back to the uncompressed portion in the memory device, for example, in response to receiving a request to access data in the compressed memory region. The processing logic of process 200 may determine if the requested data is currently available in the uncompressed portion of memory. If the requested data is not currently available in the uncompressed portion of memory, the request may cause a page fault event to trigger

the update on the compressed portion of memory, e.g. to uncompress the compressed memory region.

**[0055]** The compressed portion of memory may be stored in multiple (or at least one) compressed memory units. Each compressed memory unit may correspond to a separate memory region compressed from the uncompressed portion of memory. Separate compressed memory units may not be stored consecutively (e.g. based on physical allocation within the memory device) in the compressed portion of memory. In one embodiment, the compressed portion of memory may be compacted to cause these separate compressed memory units to be stored consecutively in the compressed portion of memory.

**[0056]** A memory region may correspond to a memory unit of a fixed size, such as a page of memory, in the uncompressed portion of a memory device. A compressed memory unit may vary in size depending on, for example, amount or compression property of original data content compressed into the compressed memory unit. In one embodiment, a compressed memory region may correspond to a particular compressed memory unit in the compressed portion of memory.

**[0057]** In one embodiment, a particular compressed memory unit may be stored in between adjacent compressed memory units in the compressed portion of memory. When the particular compressed memory unit is removed from the compressed portion of memory, the adjacent memory units may be moved to be stored consecutively next to each other, e.g. via compaction operations. The particular compressed memory unit may be removed from the compressed portion of memory via, for example, decompression operation, swapping operations, or other applicable memory management operations.

**[0058]** In some embodiments, the compressed portion of memory may include multiple pools of memory slots. Each compressed memory unit may be stored in one or more consecutive (according to physical memory allocation) memory slots. Each pool may correspond to a separate partition of the compressed memory units. In one embodiment, each pool may be associated with a separate thread, e.g. belonging to a process running in a system using the memory. Compaction operations on the compressed portion of memory may be separately and asynchronously performed for each pool via its associated thread.

**[0059]** The memory device may be coupled with a separate storage device, such as a hard disk, a flash memory or other applicable non-volatile storage device. One or more of the compressed memory units may be swapped out to the separate storage device from the



compressed portion of memory (e.g. to update the compressed portion of memory). As a result, the size of the compressed portion of memory may be decreased to increase the size of the uncompressed portion of memory. Alternatively, the compressed memory units for compressed memory regions may be swapped to the compressed portion of memory from the separate storage device if data in the compressed memory units are not currently available in the compressed portion of memory when requested.

**[0060]** The processing logic of process 200 may swap compressed memory units between the storage device (e.g. a swap device) and the memory device via an interface having a data transfer buffer, such as a bus buffer with a certain bit size. In one embodiment, the processing logic of process 200 may compact the compressed portion of memory to allow the data transfer buffer to efficiently carry multiple compressed memory units together, for example, with substantially full capacity in each hardware data transaction.

**[0061]** In one embodiment, the memory region may correspond to a memory page in the memory device. The uncompressed portion of memory can include a working set of memory pages to support currently active data processing processes. The working set of memory pages may be compressed into a series of compressed memory units in the compressed portion of memory, for example, in response to a hibernation request to temporarily halt operations of the system. The series of compressed memory units can then be swapped out to a storage device coupled with the memory device for future use.

**[0062]** The processing logic of process 200 can select at least one compressed memory units to swap out of the compressed portion of memory. The selection may be based on usage properties of the compressed memory units. The usage properties may indicate when data is accessed for the compressed memory units. For example, the usage properties may include time stamps or other meta data associated with a compressed memory unit to indicate when the compressed memory unit was accessed, created, compressed, swapped in, or used etc.

**[0063]** In one embodiment, the processing logic of process 200 may swap out the compressed memory units which are least likely to be used or requested for the working set, such as via a LRU strategy. The processing logic of process 200 may heuristically determine likelihood of the compressed memory units to be accessed in the near future based on associated usage properties. In one embodiment, the compressed memory units may be stored linearly with an order in the compressed portion of memory. The order may correspond to the usage properties (e.g. access time stamps) associated with the

compressed memory units.

[0064] The compressed portion of memory may include one or more memory segments. Each memory segment can correspond to an array of one or more of the compressed memory units. The processing logic of process 200 can move or slide one or more compressed memory units within a particular memory segment for performing compaction operations to allow multiple compressed memory units to be stored in contiguous memory space within the particular memory segment.

[0065] In one embodiment, each memory segment can have a memory capacity and/or a usage rate. The usage rate can indicate how much of the memory capacity has been used to store a corresponding array of compressed memory units. The processing logic of process 200 can perform the compaction operations to move one or more of the compressed memory units across the memory segments to increase the usage rate of at least one of the memory segments. For example, a memory segment which is 60% full may become 99% full after the compaction operations.

[0066] Figure 3 is a flow diagram illustrating one embodiment of a process to swap compressed memory pages for a virtual memory. For example, process 300 may be performed by some components of system 100 of Figure 1. At block 301, the processing logic of process 300 can provide a virtual memory based on a memory device. The processing logic of process 300 may dynamically partition the memory device into an uncompressed portion (or layer) of memory and a compressed portion (or layer) of memory. A working set of memory space used to perform data processing tasks by one or more processors coupled to the memory device may be stored in the uncompressed portion of memory.

[0067] At block 303, the processing logic of process 300 can compress one or more memory pages of the uncompressed portion of memory into one or more varied sized compressed memory units in the compressed portion of memory. The compression can increase available memory space in the uncompressed portion of memory.

[0068] At block 305, the processing logic of process 300 can decompress a particular one of the compressed memory units from the compressed portion of memory to the uncompressed portion of memory, for example, in response to a page fault as a result of an access request to data in the particular compressed memory unit.

[0069] At block 307, the processing logic of process 300 may swap out at least one varied sized compressed memory unit from the compressed portion of memory to the mass storage device to increase available memory space in the uncompressed portion of

memory, for example, to accommodate additional allocation requirement for the working set in the uncompressed portion of memory.

[0070] Figure 4 is a flow diagram illustrating one embodiment of a process to swap a compacted segment of compressed memory for a virtual memory. For example, process 400 may be performed by some components of system 100 of Figure 1. At block 401, the processing logic of process 400 can compress one or more memory pages of an uncompressed portion of a memory device into one or more varied sized compressed memory units in a compressed portion of the memory device. The memory device may be coupled with a mass storage device.

[0071] At block 403, the processing logic of process 400 may compact the varied sized compressed memory units into a segment or other applicable log structure data in the compressed portion of the memory device.

[0072] In response to receiving an access request to compressed data in one of the compressed memory units, at block 405, the processing logic of process 400 can decompress one or more compressed memory units, for example, from the segment in the compressed portion to the uncompressed portion of the memory device.

[0073] At block 407, the processing logic of process 400 can swap out the segment of the compressed memory units to the mass storage device to increase available memory space in the memory device.

[0074] Figure 5 shows one example of a data processing system, such as a computer system, which may be used with one embodiment the present invention. For example, system 1 of Figure 1 may be implemented as a part of the system shown in Figure 5. Note that while Figure 5 illustrates various components of a computer system, it is not intended to represent any particular architecture or manner of interconnecting the components as such details are not germane to the present invention. It will also be appreciated that network computers and other data processing systems which have fewer components or perhaps more components may also be used with the present invention.

[0075] As shown in Figure 5, the computer system 500, which is a form of a data processing system, includes a bus 503 which is coupled to a microprocessor(s) 505 and a ROM (Read Only Memory) 507 and volatile RAM 509 and a non-volatile memory 511. The microprocessor 505 may retrieve the instructions from the memories 507, 509, 511 and execute the instructions to perform operations described above. The bus 503 interconnects these various components together and also interconnects these components 505, 507, 509, and 511 to a display controller and display device 513 and to peripheral

devices such as input/output (I/O) devices which may be mice, keyboards, modems, network interfaces, printers and other devices which are well known in the art. Typically, the input/output devices 515 are coupled to the system through input/output controllers 517. The volatile RAM (Random Access Memory) 509 is typically implemented as dynamic RAM (DRAM) which requires power continually in order to refresh or maintain the data in the memory.

[0076] The mass storage 511 is typically a magnetic hard drive or a magnetic optical drive or an optical drive or a DVD RAM or a flash memory or other types of memory systems which maintain data (e.g. large amounts of data) even after power is removed from the system. Typically, the mass storage 511 will also be a random access memory although this is not required. While Figure 5 shows that the mass storage 511 is a local device coupled directly to the rest of the components in the data processing system, it will be appreciated that the present invention may utilize a non-volatile memory which is remote from the system, such as a network storage device which is coupled to the data processing system through a network interface such as a modem or Ethernet interface or wireless networking interface. The bus 503 may include one or more buses connected to each other through various bridges, controllers and/or adapters as is well known in the art.

[0077] Portions of what was described above may be implemented with logic circuitry such as a dedicated logic circuit or with a microcontroller or other form of processing core that executes program code instructions. Thus processes taught by the discussion above may be performed with program code such as machine-executable instructions that cause a machine that executes these instructions to perform certain functions. In this context, a “machine” may be a machine that converts intermediate form (or “abstract”) instructions into processor specific instructions (e.g., an abstract execution environment such as a “virtual machine” (e.g., a Java Virtual Machine), an interpreter, a Common Language Runtime, a high-level language virtual machine, etc.), and/or, electronic circuitry disposed on a semiconductor chip (e.g., “logic circuitry” implemented with transistors) designed to execute instructions such as a general-purpose processor and/or a special-purpose processor. Processes taught by the discussion above may also be performed by (in the alternative to a machine or in combination with a machine) electronic circuitry designed to perform the processes (or a portion thereof) without the execution of program code.

[0078] An article of manufacture may be used to store program code. An article of manufacture that stores program code may be embodied as, but is not limited to, one or

more memories (e.g., one or more flash memories, random access memories (static, dynamic or other)), optical disks, CD-ROMs, DVD ROMs, EPROMs, EEPROMs, magnetic or optical cards or other type of machine-readable media suitable for storing electronic instructions. Program code may also be downloaded from a remote computer (e.g., a server) to a requesting computer (e.g., a client) by way of data signals embodied in a propagation medium (e.g., via a communication link (e.g., a network connection)).

[0079] The preceding detailed descriptions are presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the tools used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of operations leading to a desired result. The operations are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

[0080] It should be kept in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the above discussion, it is appreciated that throughout the description, discussions utilizing terms such as "processing" or "computing" or "calculating" or "determining" or "displaying" or the like, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

[0081] The present invention also relates to an apparatus for performing the operations described herein. This apparatus may be specially constructed for the required purpose, or it may comprise a general-purpose computer selectively activated or reconfigured by a computer program stored in the computer. Such a computer program may be stored in a computer readable storage medium, such as, but is not limited to, any type of disk including floppy disks, optical disks, CD-ROMs, and magnetic-optical disks, read-only memories (ROMs), RAMs, EPROMs, EEPROMs, magnetic or optical cards, or any type

of media suitable for storing electronic instructions, and each coupled to a computer system bus.

[0082] The processes and displays presented herein are not inherently related to any particular computer or other apparatus. Various general-purpose systems may be used with programs in accordance with the teachings herein, or it may prove convenient to construct a more specialized apparatus to perform the operations described. The required structure for a variety of these systems will be evident from the description below. In addition, the present invention is not described with reference to any particular programming language. It will be appreciated that a variety of programming languages may be used to implement the teachings of the invention as described herein.

[0083] The foregoing discussion merely describes some exemplary embodiments of the present invention. One skilled in the art will readily recognize from such discussion, the accompanying drawings and the claims that various modifications can be made without departing from the spirit and scope of the invention.

**CLAIMS**

What is claimed is:

1. A machine-readable non-transient storage medium having instructions stored therein, which when executed by a computer, cause the computer to perform a method comprising:

dynamically updating a compressed portion of memory in a memory device for a virtual memory, the memory device including an uncompressed portion of memory separate from the compressed portion of memory, the virtual memory capable of mapping a memory address to the compressed portion of memory;

compressing a memory region allocated in the uncompressed portion of memory into the compressed portion of memory, wherein the memory region becomes available for future allocation in the memory device, wherein the updated compressed portion is capable of storing the compressed memory region; and

in response to a request to access data in the compressed memory region, decompressing the compressed memory region back to the uncompressed portion in the memory device.

2. The medium of claim 1, wherein the request causes a page fault which indicates the data is not available in the uncompressed portion of memory.

3. The medium of claim 1, wherein the update comprises:

allocating additional memory from the uncompressed portion of memory for the compressed portion of memory, wherein size of available memory for allocation in the uncompressed portion of memory increases as a result of the compression of the memory region.

4. The medium of claim 1, wherein the compressed portion of memory stores a plurality of compressed memory units, each compressed memory unit corresponding to a separate memory region compressed from the uncompressed portion of memory, wherein at least two of the compressed memory units are not stored consecutively in the compressed portion of memory, further comprising:

compacting the compressed portion of memory, wherein the compacting causes the at least two of the compressed memory units to be stored consecutively in the compressed portion of memory.

5. The medium of claim 4, wherein the compressed memory region corresponds to a particular compressed memory unit in the compressed portion of memory, wherein the particular compressed memory unit was stored in between the at least two compressed memory units, wherein the decompression removes the particular compressed memory unit from the compressed portion of memory to allow the at least two of the compressed memory units to be stored consecutively.

6. The medium of claim 4, wherein the memory device is coupled with a separate storage device, wherein the update comprises:

swapping out one or more of the compressed memory units from the compressed portion of memory to the separate storage device, wherein the swapping out causes a decrease in size of the compressed portion of memory to increase size of the uncompressed portion of memory.

7. The medium of claim 6, further comprising:

determining if the data of the compressed memory region for the request is currently available in the compressed portion of memory, wherein the compressed memory region is decompressed if the data of the compressed memory region is currently available in the compressed portion of memory.

8. The medium of claim 7, further comprising:

swapping in the compressed memory region from the separate storage device to the compressed portion of memory if the data of the compressed memory region is not currently available in the compressed portion of memory.

9. The medium of claim 6, wherein a particular one of the one or more of the compressed memory units was stored in between the at least two compressed memory units, wherein the swapping out removes the particular compressed memory unit from the



compressed portion of memory to allow the at least two of the compressed memory units to be stored consecutively.

10. The medium of claim 6, wherein the swapping out comprises:  
selecting the one or more of the compressed memory units for the swapping and  
wherein the selection is based on usage properties associated with the  
compressed memory units.
11. The medium of claim 10, wherein the memory region corresponds to a memory  
unit of a fixed size in the memory device, wherein the compressed memory units are of  
varied sizes, and wherein the compressed memory units are stored linearly with an order  
in the compressed portion of memory.
12. The medium of claim 11, wherein the usage properties corresponds to the linear  
order associated with the compressed memory units.
13. The medium of claim 12, wherein usage properties indicate when data is accessed  
for the compressed memory units and wherein the one or more of the compressed  
memory units selected are least recently used according to the linear order for the  
selection.
14. The medium of claim 11, wherein the memory device is coupled with the separate  
storage device via an interface having a data transfer buffer, wherein the swapping is  
based on data transfer over the interface via the data transfer buffer, and wherein the  
compaction allows the data transfer buffer to carry a plurality of the compressed memory  
units together.
15. The medium of claim 11, wherein the compressed portion of memory includes one  
or more memory segments, each memory segment corresponds to an array of one or more  
of the compressed memory units, a particular array of one of the memory segments  
includes two or more compressed memory units, the compacting comprises:  
moving at least one of the compressed memory units within the particular  
memory segment such the two or more compressed memory units are  
stored in contiguous memory space within the particular memory segment.

16. The medium of claim 15, wherein each memory segment has a memory capacity, each memory segment having a usage rate indicating how much of the memory capacity has been used to store corresponding array of compressed memory units, the compacting further comprising:

moving one or more of the compressed memory units across the memory segments, wherein corresponding usage rate of at least one of the memory segments increases as a result the moving of the one or more of the compressed memory units.

17. The medium of claim 4, wherein the compressed portion of memory includes multiple pools of memory slots, each compressed memory unit stored in one or more consecutive ones of the memory slot, each pool corresponding to a separate partition of the compressed memory units, each pool is associated with a separate thread and wherein the compaction is separately and asynchronously performed for each pool via associated thread.

18. The medium of claim 6, wherein the memory region corresponds to a memory page in the memory device, wherein the uncompressed portion of memory includes a working set of memory pages to support currently active data processing processes, further comprising:

in response to a hibernation request, compressing the working set of memory pages to a series of compressed memory units in the compressed portion of memory; and

swapping out the series of compressed memory units to the storage device.

19. A machine-readable non-transient storage medium having instructions stored therein, which when executed by a computer, cause the computer to perform a method comprising:

providing a virtual memory based on a memory device, the memory device dynamically partitioned into an uncompressed portion of memory and a compressed portion of memory, the uncompressed portion of memory to store working data processed via a processor coupled to the memory device;

compressing one or more memory pages of the uncompressed portion of memory into one or more varied sized compressed memory units in the compressed portion of memory, wherein the compression increases available memory space in the uncompressed portion of memory;

in response to a page fault for an access request to data in one of the compressed memory units, decompressing the one compressed memory units from the compressed portion of memory to the uncompressed portion of memory;

and

swapping out at least one of the varied sized compressed memory units from the compressed portion of memory to the mass storage device, wherein the swapping out increases available memory space in the uncompressed portion of memory.

20. A machine-readable non-transient storage medium having instructions stored therein, which when executed by a computer, cause the computer to perform a method comprising:

compressing one or more memory pages of a uncompressed portion of a memory device into one or more varied sized compressed memory units in a compressed portion of the memory device, the memory device coupled with a mass storage device;

compacting the varied sized compressed memory units into a segment in the compressed portion of the memory device;

in response to access request to one of the compressed memory units,

decompressing the one compressed memory unit from the compressed portion to the uncompressed portion of the memory device; and

swapping out the segment of the compressed memory units to the mass storage device to increase available memory space in the memory device.

21. A computer implemented method comprising:

dynamically updating a compressed portion of memory in a memory device, the memory device including an uncompressed portion of memory capable of supporting data processing for at least one processor;

compressing a memory region allocated in the uncompressed portion of memory into the compressed portion of memory, wherein the memory region

becomes available for future allocation in the memory device, wherein the updated compressed portion is capable of storing the compressed memory region; and

in response to a request to access data in the compressed memory region, decompressing the compressed memory region back to the uncompressed portion in the memory device.

22. A computer system comprising:

a memory including executable instructions, the memory dynamically partitioned into a compressed portion and an uncompressed portion;

a mass storage device;

a processor coupled to the memory and the mass storage device, the processor to execute the executable instructions from the memory, the processor being configured to

compress a memory region allocated in the uncompressed portion of memory into the compressed portion of memory, wherein the memory region becomes available for future allocation in the memory device, wherein the updated compressed portion is capable of storing the compressed memory region;

in response to a request to access data in the compressed memory region, decompressing the compressed memory region back to the uncompressed portion in the memory device; and

in response to detecting a high memory usage in the memory device, swap out the compressed memory region to the storage device to increase available memory space in the memory device.

1/5

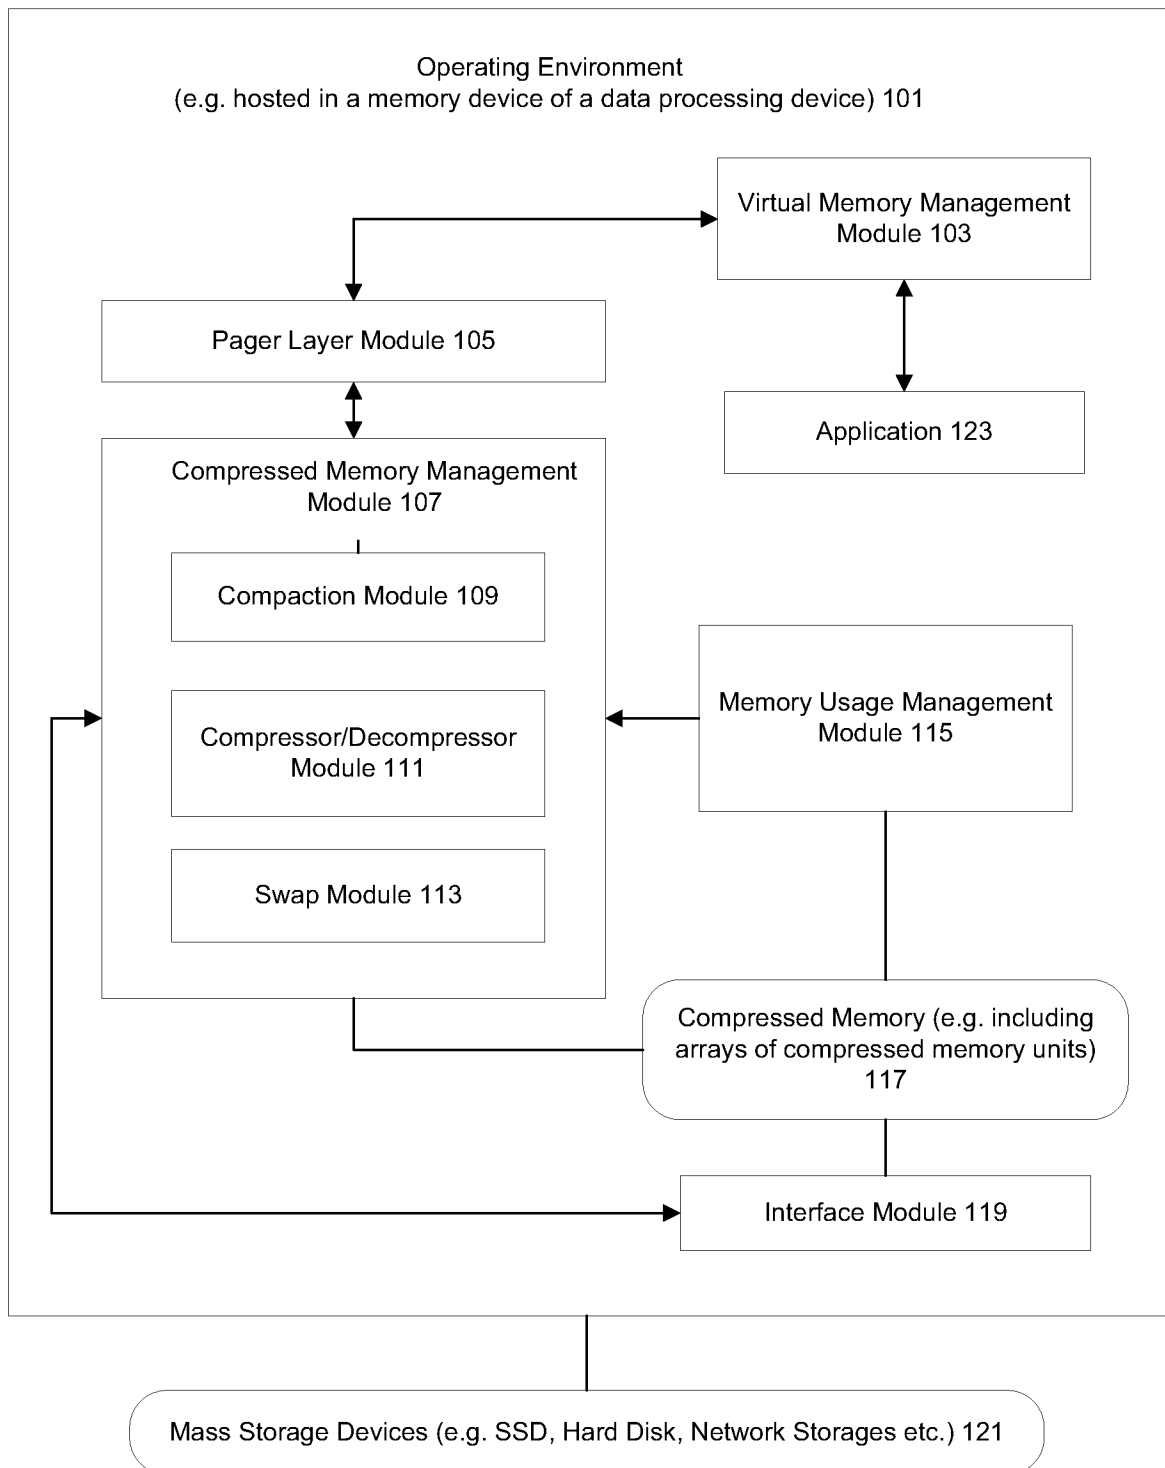
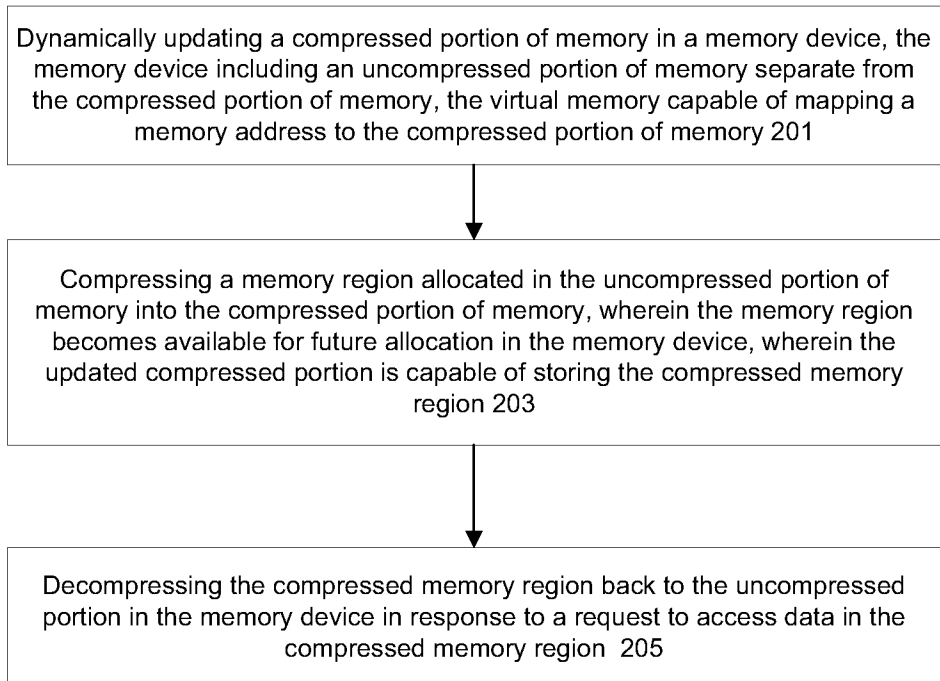
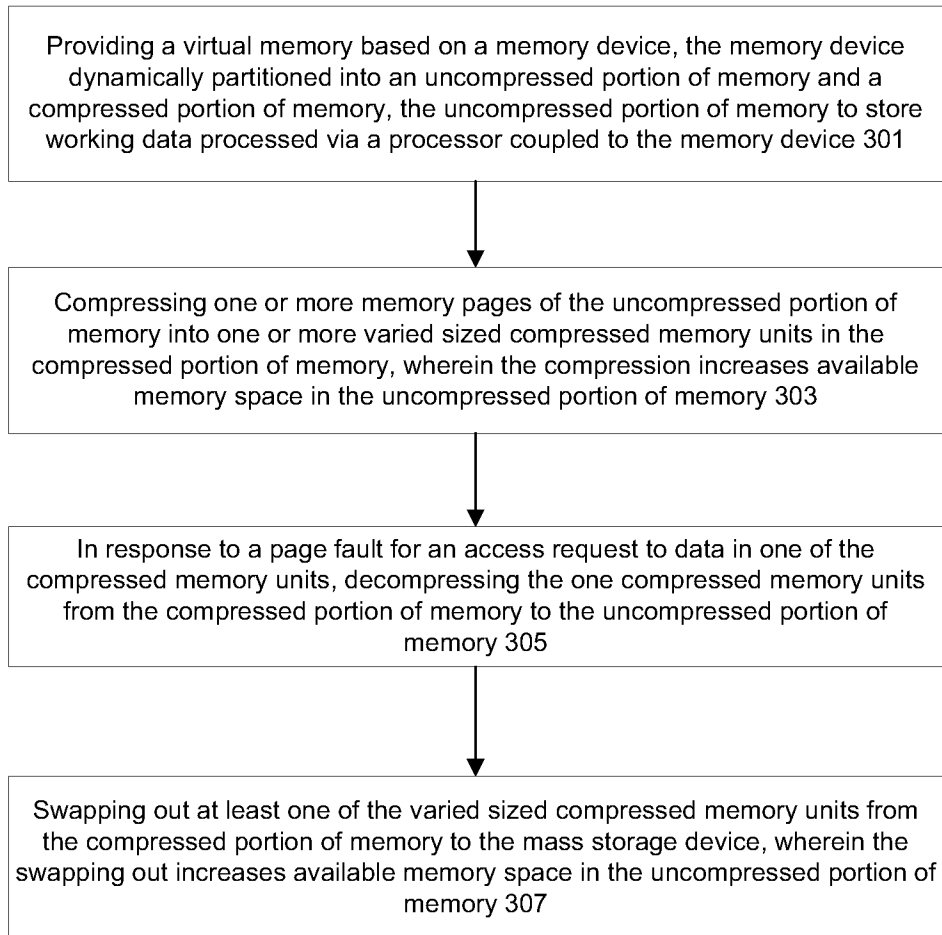
100

Fig. 1

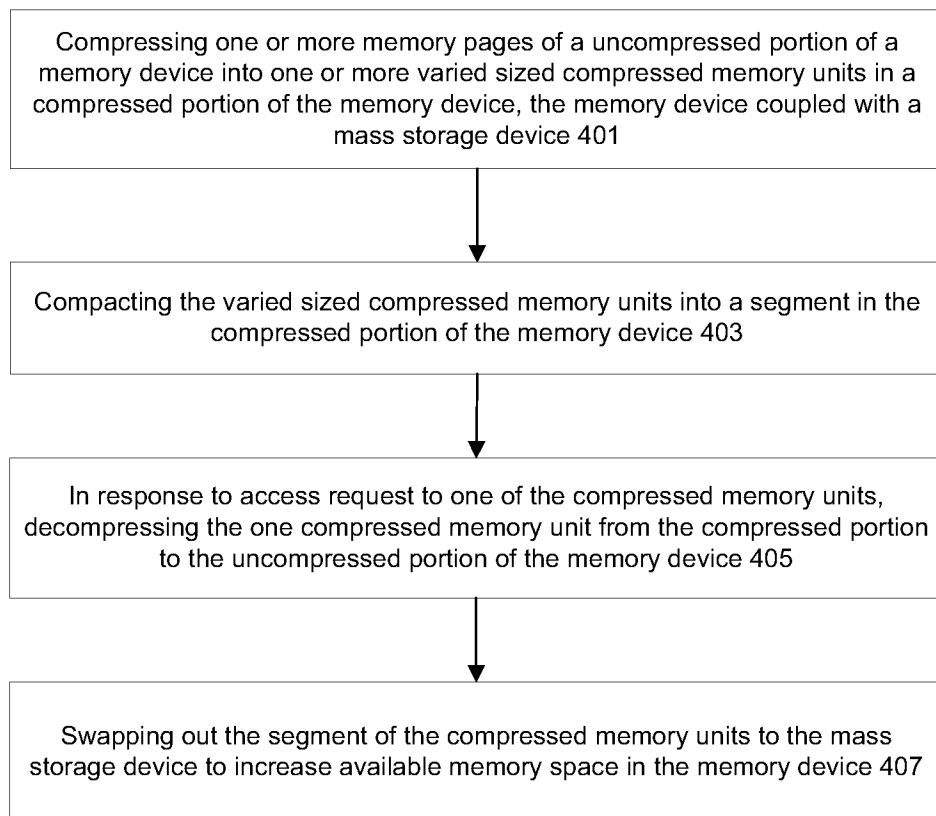
2/5

200**Fig. 2**

3/5

300**Fig. 3**

4/5

400**Fig. 4**



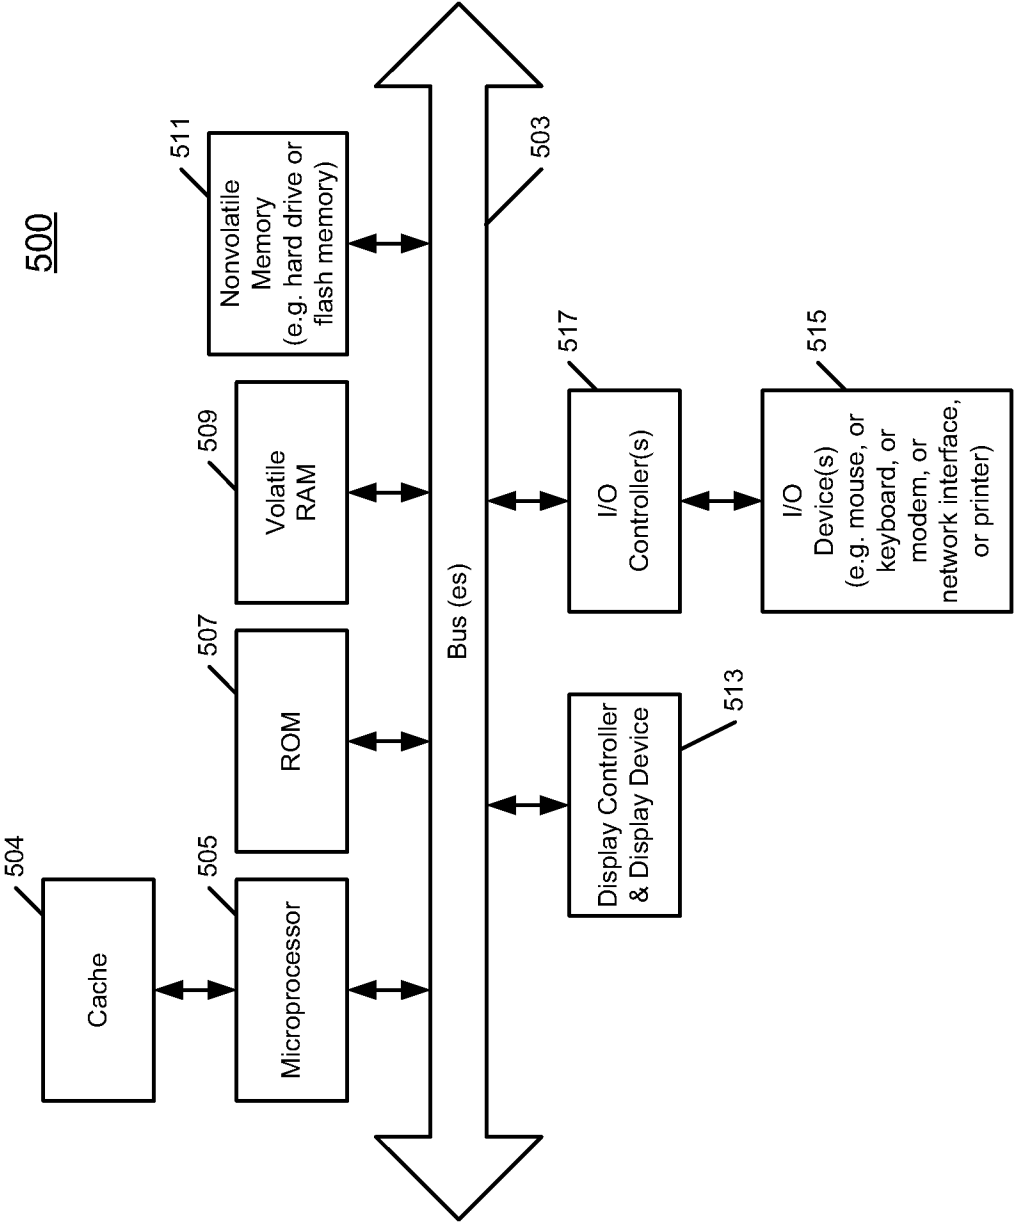


Fig. 5

## INTERNATIONAL SEARCH REPORT

International application No

PCT/US2013/071753

## A. CLASSIFICATION OF SUBJECT MATTER

INV. G06F12/08

ADD. G06F12/02

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 2009/112949 A1 (ERGAN CENK [US] ET AL) 30 April 2009 (2009-04-30)	1-11,15, 16,19, 21,22
Y	paragraphs [0002] - [0005], [0017] - [0026], [0039], [0064], [0065],	12,14, 17,18
A	[0070], [0075], [0076], [0085] - [0088]; figure 1	13
X	US 7 587 572 B1 (STENSTROM PER O [SE]) 8 September 2009 (2009-09-08) column 2, line 9 - column 3, line 32 column 11, lines 17-46	1-3,19, 21,22
Y	US 7 412 466 B1 (GARTHWAITE ALEXANDER T [US]) 12 August 2008 (2008-08-12)	12
A	column 2, line 43 - column 3, line 34	13
	-/--	



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

"&amp;" document member of the same patent family

Date of the actual completion of the international search

20 May 2014

Date of mailing of the international search report

11/06/2014

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

Nielsen, Ole

## INTERNATIONAL SEARCH REPORT

International application No

PCT/US2013/071753

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	CORTES T ET AL: "SWAP COMPRESSION: RESURRECTING OLD IDEAS", SOFTWARE PRACTICE & EXPERIENCE, WILEY & SONS, BOGNOR REGIS, GB, vol. 30, no. 5, 25 April 2000 (2000-04-25) , pages 567-587, XP000908375, ISSN: 0038-0644, DOI: 10.1002/(SICI)1097-024X(20000425)30:5<567: :AID-SPE312>3.0.CO;2-Z	20
Y	pages 569-570, page 574 - page 575, line 18	14
Y	----- US 7 533 228 B1 (GARTHWAITE ALEXANDER T [US]) 12 May 2009 (2009-05-12) column 3, line 66 - column 4, line 48	17
Y	----- US 2002/138669 A1 (KADATCH ANDREW V [US] ET AL) 26 September 2002 (2002-09-26) paragraphs [0048], [0049], [0052], [0053]; claim 16; figures 2,4a	18
X	----- US 2009/327621 A1 (KLIOT GABRIEL [IL] ET AL) 31 December 2009 (2009-12-31) paragraphs [0006], [0027], [0037] - [0039], [0041], [0071] - [0076], [0083]; figures 2,11,12-14	20
	-----	

# INTERNATIONAL SEARCH REPORT

International application No.  
PCT/US2013/071753

## Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:  
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☐ Claims Nos.:  
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
3. ☐ Claims Nos.:  
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

## Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

see additional sheet

1. ☒ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

### Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☒ No protest accompanied the payment of additional search fees.

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2013/071753

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2009112949 A1	30-04-2009	US 2009112949 A1	30-04-2009
		US 2011113432 A1	12-05-2011
US 7587572 B1	08-09-2009	NONE	
US 7412466 B1	12-08-2008	NONE	
US 7533228 B1	12-05-2009	NONE	
US 2002138669 A1	26-09-2002	US 2002138669 A1	26-09-2002
		US 2005097239 A1	05-05-2005
		US 2005108589 A1	19-05-2005
		US 2005149646 A1	07-07-2005
US 2009327621 A1	31-12-2009	NONE	

**FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210**

This International Searching Authority found multiple (groups of) inventions in this international application, as follows:

1. claims: 1-13, 15, 16, 19, 21, 22

Storing compressed memory units linearly with an order according to usage properties.

---

2. claim: 14

Swapping based on transfer via a data buffer such that the compaction allows the transfer of multiple compressed units together.

---

3. claim: 17

The compressed portion of memory including multiple pools of memory, each pool being associated with a separate thread, the compaction being performed separate for each pool via associated thread

---

4. claim: 18

In response to a hibernation request compressing working set in uncompressed portion and swapping out the compressed units to storage.

---

5. claim: 20

Swapping out a segment of compressed memory units to mass storage device.

---