



- (51) **International Patent Classification:**
G06F 12/08 (2006.01) *G06T 1/60* (2006.01)
- (21) **International Application Number:**
PCT/US2016/019274
- (22) **International Filing Date:**
24 February 2016 (24.02.2016)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
14/666,756 24 March 2015 (24.03.2015) US
- (71) **Applicant:** INTEL CORPORATION [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95054 (US).
- (72) **Inventors:** NILSSON, Jim K.; Svenska vägen 4, 226 39 Lund, M (SE). AKENINE-MÖLLER, Tomas G.; Rosenvägen 13, 22738 Lund, M (SE). CLARBERG, Franz P.; Filippavägen 4 D, 22241 Lund, M (SE).
- (74) **Agents:** TROP, Timothy N. et al.; Trop, Pruner & Hu, P.C., 1616 S. Voss Rd., Ste. 750, Houston, Texas 77057-2631 (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).

Declarations under Rule 4.17:

- as to the identity of the inventor (Rule 4.17(i))
- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

[Continued on next page]


(54) Title: COMPACTION FOR MEMORY HIERARCHIES

FIG. 1

(57) **Abstract:** In accordance with some embodiments, compaction, as contrasted with compression, is used to reduce the footprint of a near memory. In compaction, the density of data storage within a storage device is increased. In compression, the number of bits used to represent information is reduced. Thus you can have compression while still having sparse or non-contiguously arranged storage. As a result, compression may not always reduce the memory footprint. By compacting compressed data, the footprint of the information stored within the memory may be reduced. Compaction may reduce the need for far memory accesses in some cases.



Published:

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

COMPACTION FOR MEMORY HIERARCHIES

Background

[0001] Graphics and central processors may have a storage system with a large near memory (NM) cache, which is relatively fast, and this is in turn backed by far memory (FM), which is inexpensive, but also much slower. In such a hierarchical memory architecture, it is very important to get a hit rate in NM which is as high as possible, because on a miss, a request will be sent to FM, and it will take a relatively long time before the requested data is accessible. For example the near memory may be volatile memory (i.e. dynamic random access memory) mounted on a graphics/central processing unit of system-on-a-chip. The far memory may be non-volatile memory such as flash memory.

[0002] At the same time, the processor can compress some of its buffers to a small set of fixed sizes. The compressed cache lines are stored "sparse" (non-contiguously) for quick random access. Currently, since the NM cannot exploit this type of compression, the compressed data in the NM still has the same footprint as uncompressed data in NM. Thus, a substantial portion of NM, which is more expensive, is unused at any instance in time.

[0003] Current render cache data compression techniques, e.g., for color, do not reduce the memory footprint of render buffers. Those techniques only reduce the amount of data traffic to and from memory. Compaction is a way of reducing the memory footprint of already compressed data.

[0004] Other lossy data compression techniques can reduce footprint, but only at a potentially unbounded loss of quality. For example, a lossily compressed buffer that is read/modify/written repeatedly can deteriorate in quality. Data loss is in many cases not acceptable (e.g., due to API-compatibility, visual artifacts, non-image data).

[0005] Compaction is a way of reducing the memory footprint of already compressed data, without introducing any loss. Footprint reduction in NM enables storing more data, before triggering expensive far memory (FM) accesses.

Therefore, the apparent capacity of NM is increased, resulting in higher performance, less expensive system-on-a-chip, and reduced power consumption in some embodiments.

Brief Description Of The Drawings

[0006] Some embodiments are described with respect to the following figures:

Figure 1 is a schematic depiction of one embodiment;

Figure 2 is a flow chart for one embodiment of a compaction buffer;

Figure 3 is a flow chart for compaction according to one embodiment;

Figure 4 is a depiction of compaction according to one embodiment;

Figure 5 is a schematic depiction of compaction according to still another embodiment;

Figure 6 is a schematic depiction for another embodiment;

Figure 7 is an operational diagram for another embodiment;

Figure 8 is a flow chart for a compaction daemon according to one embodiment;

Figure 9 is a flow chart for an algorithm for selecting pages to compact;

Figure 10 is a block diagram of a processing system according to one embodiment;

Figure 11 is a block diagram of a processor according to one embodiment;

Figure 12 is a block diagram of a graphics processor according to one embodiment;

Figure 13 is a block diagram of a graphics processing engine according to one embodiment;

Figure 14 is a block diagram of another embodiment of a graphics processor;

Figure 15 is a depiction thread execution logic according to one embodiment;

Figure 16 is a block diagram of a graphics processor instruction format according to some embodiments;

Figure 17 is a block diagram of another embodiment of a graphics processor;

Figure 18A is a block diagram of a graphics processor command format according to some embodiments;

Figure 18B is a block diagram illustrating a graphics processor command sequence according to some embodiments;

Figure 19 is a depiction of an exemplary graphics software architecture according to some embodiments;

Figure 20 is a block diagram illustrating an IP core development system according to some embodiments; and

Figure 21 is a block diagram showing an exemplary system on chip integrated circuit according to some embodiments.

Detailed Description

[0007] In accordance with some embodiments, compaction, as contrasted with compression, is used to reduce the footprint of a near memory. In compaction, the density of data storage within a storage device is increased. In compression, the number of bits used to represent information is reduced. Thus you can have compression while still having sparse or non-contiguously arranged storage. As a result, compression may not always reduce the memory footprint. By compacting compressed data, the footprint of the information stored within the memory may be reduced. Compaction may reduce the need for far memory accesses in some cases.

[0008] In accordance with a first technique, data from a buffer such as a render buffer is compacted before it is entered into a near memory. This may be done using a buffered compaction technique.

[0009] In accordance with another embodiment, a compaction daemon may operate independently on the near memory to free up space. In accordance with

some embodiments, the compaction using the compaction daemon may be done as late as possible to reduce unnecessary memory traffic and logic upon premature compaction.

[0010] A daemon is a process that performs a specified operation at predetermined times or in response to certain events. It may be a computer program, firmware, or hardware. Generally a daemon is not called by a user.

[0011] Near and far memory may be coupled to a processor. The near memory is faster than the far memory and the near memory is backed by the far memory. A codec is coupled to a render cache. According to the first technique, a buffer between said codec and the near memory stores compressed data for the codec. The compacted and compressed data in the buffer, stored for more contiguous (or less sparse) addressing, has a smaller footprint than the same data when compressed but not compacted.

[0012] With compression, cache lines in near memory (NM) may be compressed down to a multiple of sub-blocks of a given size when evicted from the color cache, thereby reducing data traffic to lower level caches of a far memory (FM). For example, compression may be 2:N or 4:N. The cache lines may typically be 128B or 256B. Then the cache lines may be compressed to an integer multiple or number of 64B sub-blocks in one embodiment.

[0013] The near memory may be broken into pages of an allocated size called the near memory page size. A compaction page group is a group of contiguously addressed NM pages. A compaction page group size is a multiple of the page allocation size of the NM.

[0014] For example, with an NM page allocation size of 512 bytes, and a compaction page group size of 2,048 bytes, it is possible to coalesce compressed cache lines to fit into a fewer number of NM pages, in this case one, two, or three NM pages, instead of four, thereby reducing the required footprint in NM. If the amount of compressed cache lines is too low (i.e., few compressed lines, or low compression ratios), no compaction may be possible in some cases.

[0015] While examples are given in the context of graphics data from a render cache, embodiments are generally applicable to handling compressed data in memory hierarchies with two or more levels that use compressed random access data.

[0016] As shown in **Figure 1**, in an embodiment using a render cache, a central processing unit or graphics processing unit (GPU) 10 communicates with a render cache 12, such as a color cache, and a far memory 11. A render cache may be any cache used for storing compressed data for a graphics processor including a color cache. A compaction buffer (CB) 14 is provided after the codec 16 that compresses and decompresses pixel data. The compaction buffer gathers compressed cache lines (CLs) belonging to compaction page groups (e.g., 2kB for a NM page size of 512B). If a compaction page group is fully populated, it can be safely compacted and evicted from the compaction buffer to NM 18. If the data is sufficiently compacted it may require substantially less footprint in NM after compaction.

[0017] When compacted and compressed data is read from NM, modified, and written back, the data may occupy more or less data than before, giving rise to expansion and shrinkage, described in more detail below. Given a reasonably sized compaction buffer, many of these occurrences can be avoided.

[0018] A footprint reduction in NM enables storage of more and larger render targets and buffers in NM as opposed to in more distant parts of the memory system (i.e., far memory (FM)), thereby reducing power consumption and latency, and resulting in higher performance rendering in some embodiments. The techniques described herein may reduce the problem of limited NM capacity in some cases.

[0019] The compaction buffer (CB) holds cache lines evicted from a render (e.g. color) cache. Only cache lines that have been modified are evicted to the CB in some embodiments. These cache lines can be compressed or not. A compaction control surface 17 tracks and logs how and to what degree a particular cache line is compressed by the codec.

[0020] Compared to the near memory, the compaction buffer is fast like the near memory but much smaller. Generally the compaction buffer may be as fast as the

render cache that feeds compressed data to the compaction buffer. The compaction buffer may be located physically close to the render cache in some embodiments.

[0021] The items in the CB may be groups of contiguously addressed NM pages called compaction page groups (CPGs). At eviction from the render cache, the CB is inspected to see if it currently holds the corresponding CPG. The addressing logic to map a cache line to a CPG is, in one case, a right shift operation to strip least significant bits of the address, but the logic can be more elaborate, if needed.

[0022] A compaction control surface (CCS) in a control surface 15 keeps track of which NM pages are compacted. This CCS indicates whether a CPG is compacted. This surface may be a single bit per CPG in one embodiment and its location is arbitrary, i.e., it can be stored in a dedicated buffer on- or off-chip, or it can be a part of the general memory system.

[0023] A sequence 20 for implementing a compaction buffer, shown in **Figure 2**, may be implemented in software, firmware and/or hardware. In some embodiments the sequence may be implemented in software or firmware in the form of computer executed instructions stored in one or more non-transitory computer readable media such as magnetic, optical or semiconductor storage. For example the steps may be implemented by a graphics processing unit in some embodiments.

[0024] At eviction from the render cache (Figure 2, diamond 22) there are at least two possibilities:

1. If the CPG of the cache line is present in the CB (diamond 24, YES):
 - a. store the evicted line data in the CPG, along with its compression control bits that detail how and to what degree the cache line is compressed) (block 26).
2. If the CPG of the cache line is not present in the CB (diamond 24, NO):
 - a. if the CB is not full (diamond 28, YES), one free CPG may be allocated and the evicted line may be stored in it, along with the corresponding compression control bits (block 30).

b. If the CB is full (diamond 28, NO), a CPG needs to be evicted (block 32). The replacement policy used can be arbitrary, but in one embodiment a policy may be based on age (i.e., similar to a least-recently-used (LRU) replacement policy) combined with a metric of how populated the CPGs are, and possibly also on the amount of compactability. After the CPG is evicted, the empty CPG slot is used to store the evicted cache line as in 2a above.

[0025] A compaction sequence 34 shown in **Figure 3** may be implemented in software, firmware and/or hardware. In software and firmware embodiments it may be implemented by computer executed instructions stored in one or more non-transitory computer readable media such as magnetic, optical or semiconductor storage. For example the sequence may be implemented by the graphics processor in one embodiment.

[0026] Upon eviction from the CB, a suitable CPG is selected. If all data for the CPG is present in the CB, the CB inspects the compression control bits of the cache lines in the CPG (block 36) and calculates the amount of compaction possible (block 38). If compaction is achieved (diamond 40, YES), i.e., the number of NM pages to store the CPG is lower than storing them non-compacted, the bit in the CCS is set (block 42) and the cache lines are sent to near memory in a contiguous chunk (block 44). Any unused NM pages can be signaled as free (block 50). If compaction is not achieved (diamond 40, NO), the data is stored uncompacted in NM (block 46).

[0027] As an example, assume that a NM page size of 512 bytes and a CPG size of 1,024 bytes (i.e., two NM pages, enabling 2:1 compaction at most). Further, assume that the cache line size in the color buffer cache is 256 bytes with 4:N compression in place. At eviction of the CPG, the compression control bits are inspected, and if the total compression of all the cache lines is more than or equal to 2:1, compaction is possible. Any combination of compression ratios for individual cache lines (CLs), i.e., 4:1, 4:2, and 4:3 that results in a total compaction of 2:1, is possible. Other combinations of NM page sizes, CPG sizes, CL sizes, and compression ratios, are possible, that enable different amounts of compaction.

[0028] In any case, there can be unused parts in the last allocated NM page. For example, using 512 byte NM pages, if the compacted data of a 1,024 bytes, CPG occupies, say 576 bytes, then this algorithm allocates 2 NM pages since $1 \times 512 < 576 \leq 2 \times 512$.

[0029] A similar example is shown in **Figure 4**, where a 2kB CPG 51 is shown to the left, and it contains all CLs in it in compressed form. Hence, when such a CPG is marked for eviction from the CB, the data 51 is compacted and in this case two pages ($512B \times 2$) are needed for the compacted data. These can be allocated anywhere in this case, but most often they will be consecutive in near memory 18.

[0030] An example of expansion is shown in **Figure 5**. The data in one 2kB block 21 is read into the color cache and decompressed at 23, written to, and then compressed again at 25. At this point, the data occupies more space than before, and suddenly it does not fit into 2 pages as indicated at 27. Hence, a third page is allocated in NM at 29, and the data is shuffled into correct locations in these three pages.

[0031] When a CPG is evicted from the CB, it can happen that all data is not present for inspection, i.e., some data still resides in the render cache and some data can reside in NM. In any case, this data must be fetched into the CB to enable compaction of the CPG. Reading data from NM does require bandwidth, but may be kept low with a reasonably sized CB and efficient compression. The data belonging to the evicted CPG may be forcefully evicted from the render cache; however, it is unknown whether that data will be used again in a near future. One option is to only allow compaction of CPGs for which all CLs have been evicted from the render cache. If the rare case that none such exist, forceful eviction from the render cache may be used as a fallback.

[0032] Upon reading back data from the near memory, it is first decompressed and then placed directly in the cache, which may be a color/render cache, or other cache.

[0033] When a write to NM occurs as a result of an eviction of a dirty line from the LLC, and the page group is already in compacted form, the RCS can be consulted to decide whether expansion of the addressed page group is needed or not. If the line

evicted from the LLC is compressed in the same way as in the compacted representation in NM, no further action than writing the line is needed. An alternative is to let the daemon decompact the data in the page group, and then perform the write without consulting the RCS.

[0034] If the line evicted from LLC is not compressed to the same degree as previously, i.e., the compressed representation has either shrunk or expanded, the NM page group is first decompact before the writing of the line. This page group is then eligible for compaction at a future point in time.

[0035] If an eviction needs to be done in the compaction buffer, several replacement strategies are possible. To minimize data traffic then full CPGs may be prioritized first, followed by the oldest CPG in the compaction buffer. If no full CPGs exist, the CPG with the highest number of cache lines may be selected, since that would minimize the read traffic from NM. Other strategies are possible as well.

[0036] The compaction control surface (CCS) consisting of one bit per CPG requires relatively little space. For a 32 bits/pixel render surface of size 4096x2160, and using 2kB CPGs, $(4096 \times 2160 \times 4B) / 2048B \approx 17kB$ of CCS memory.

[0037] In the embodiments described above, data is compacted *before* it is entered in the NM using a buffered compaction approach. However, since data that is compacted and then read/modify/written may have to be expanded/shrunk, that approach advantageously uses a compaction buffer large enough to avoid most of these occurrences.

[0038] In accordance with the second technique, a compaction daemon operates independently on the near memory to free up space just-in-time or slightly ahead of time. Compacting as late as possible avoids wasting (on chip) memory traffic and logic on premature compaction.

[0039] As shown in **Figure 6**, a memory system includes buffer compression codec(s) 72 coupled to a cache hierarchy. The near memory (NM) 86 acts as a large cache after the last-level traditional cache (e.g., LLC) 78. Data written to NM is potentially compressed, but always in un-compacted form (i.e. stored sparse) in the same fashion as data is written to memory in current graphics processors.

[0040] A *compaction daemon* 90 is connected to the NM. The compaction daemon operates asynchronously on the NM to compact its data as the NM fills up. The NM is divided into memory pages of a certain size (e.g., 512B). The compaction daemon selects small ranges of memory pages that are sparse but unlikely to be written again, and compacts these into fewer memory page(s). This frees up empty pages in NM.

[0041] The compaction daemon can be an independent software (thread) that can be run at arbitrary points in time, as indicated by the current footprint level and/or budget, in turn as indicated by application program interface (API) events for example.

[0042] Several different policies for selecting pages for compaction are possible. The behavior may also benefit from software control (i.e. by the graphics driver). For example, when a render target is finished rendering and bound as a shader resource (texture), the relevant memory pages may be tagged as read-only by the driver. The compaction daemon can then operate on those pages to free up a potentially large number of pages.

[0043] The daemon effectively uses the *entire* NM as a compaction buffer. The data starts getting compacted when the NM begins to run out of space. For example when only a threshold number of open pages are available, the daemon may be triggered to start compacting in one embodiment. Hence, applications with a working set that natively fits within the NM may not incur extra on-chip traffic/processing for compaction in some embodiments.

[0044] Compaction of data can be done either before it enters the NM using a buffered approach, or *after* it has been stored in NM (but before eviction to FM), using a daemon compaction approach. Early compaction (i.e., before NM) incurs extra bandwidth/processing, even for applications with working sets that fit in NM. Additionally the control logic of the compression buffer (CB) may be more complex to handle the case when data to be evicted and compacted is not fully stored in the CB. For example, parts of a memory page(s) may reside in lower-level caches (which

must be forcefully evicted) and/or in NM (which must be read-back prior to compaction).

[0045] Referring to Figure 7, the L1 or render cache(s) 70 uses cache lines of a certain size (e.g., 256B), which are losslessly compressed by a codec 72 to one or more 64B sub-blocks, or stored uncompressed. A render control surface (RCS) 74 associated with each buffer encodes how/if a cache line is compressed. The L1 cache is coupled to central processing unit and/or graphics processing unit 68.

[0046] The codec(s) are connected to a hierarchy of higher-level cache(s), including an L2 cache 76, until a last-level cache 78. Data is stored in these caches in compressed form but un-compacted, i.e., sparse. Upon eviction of a modified cache line from the LLC cache, the data is written to near memory (NM) 86. Accesses to NM may be done through a near memory controller 80. The near memory is coupled to the far memory 92.

[0047] The near memory is organized into memory pages 84 of a certain size, e.g., 512B. These are in turn organized into larger page groups (e.g., 4 pages into 2KB). The compaction daemon 90 operates on page groups to compact them into fewer 512B pages, potentially freeing up one or more pages from each page group.

[0048] **Figure 7** shows an example of a compaction operation. Sparsely stored data in near memory 82a is compacted by compaction daemon 90 so that one page group 94 fits in one memory page 96. Then it is stored in the near memory as compacted at 82b. A page group may consist of pages with consecutive memory addresses, but which are stored non-consecutively in NM. A memory address translation table 91 (in Figure 7) in near memory controller 80 (Figure 7) keeps track of the address to page mapping.

[0049] A separate page group control surface (PGCS) 88 (Figure 7) encodes whether each page group is stored compacted or not. This may be done using as little as a single bit per page group. The near memory controller 80 (Figure 7) accesses the PGCS 88 to check whether a particular page is compacted or not.

[0050] Different possibilities exist, but one configuration lets the near memory (NM) operate as a large victim cache. Data that is evicted from a higher-level cache and

marked as modified is written to NM. Such data may be compressed, but is always un-compacted.

[0051] The compaction daemon may later compact the data to clear up memory pages in NM. Upon read access, the data is fetched (compacted or not) from NM and stored un-compacted in the last-level cache (LLC). The PGCS is inspected to find out if a certain page is compacted or not. Note that the page may remain in compacted form in NM.

[0052] When a write to NM occurs, it is generally the case that the daemon has been able to keep enough empty pages available. In the rare case it has not – either due to all data already being compacted, or due to the daemon not yet having had time to compact everything – an evict from NM to far memory (FM) may have to be done.

[0053] To further reduce the number of FM accesses, the memory controller may decide to wait for the daemon to finish more compaction operations, if such are underway.

[0054] The compaction daemon is connected via a port or bus (Figure 7, double arrows) to the near memory (NM), but the daemon may operate internally in its own clock/power domains. Since compaction occurs asynchronously and only when needed, the daemon may be idle for long periods of time. During such times, it is advantageous to power it down or lower its clock.

[0055] The daemon may either have its own data port 94 to NM, or share a data port with other units on the same interconnect (e.g., last-level traditional cache, or memory controller for far memory). In the latter cases, it may exploit unused read/write cycles on the interconnect to access the NM to perform compaction.

[0056] Pages that are being fetched by the daemon may still be simultaneously read-only accessed by other units. When compaction is complete and while the pages are being updated by the daemon (i.e., their content replaced by compacted content), they may be temporarily locked. Alternatively, if there are other empty pages in NM, the daemon may write the compacted data to those, and only when it is finished update the address translation table to point to the compacted pages. At

that point, the original un-compacted pages may be marked as empty. This strategy reduces system latencies.

[0057] The compaction daemon sequence 100, shown in **Figure 8**, may be implemented in software, firmware or hardware. In software and firmware embodiments it may be implemented by computer readable instructions stored in one or more non-transitory computer readable media such as magnetic, optical or semiconductor storage. In one embodiment it may be implemented by a graphics processing unit.

[0058] The sequence shown in **Figure 8** begins by determining whether compaction is needed at 102. If so, far memory accesses may be stayed in one embodiment as indicated in block 103. Then a range of sparse memory pages to compact may be selected as indicated in block 104. This range of pages is compacted into fewer pages as indicated in block 106. The address to page mapping is stored in an address translation table as indicated in block 108. Then a compaction indicator is stored in the PGCS as indicated in block 110 to indicate that the page group has been compacted. Finally far memory access may be unstayed as indicated in block 112 in some embodiments.

[0059] A sequence 120, shown in **Figure 9**, for compaction page selection may be implemented in software, firmware or hardware. In software and firmware embodiments it may implemented by computer readable instructions stored in one or more non-transitory computer readable media such as magnetic, optical or semiconductor storage. The sequence may be implemented for example in the graphics processing unit.

[0060] The compaction page selection sequence 120 in one embodiment begins by selecting page groups tagged as read only by a driver as indicated in block 122. Then the selected page groups are compacted as indicated in block 124.

[0061] A check at diamond 126 determines whether more compaction is needed. If so, select the oldest page groups for compaction as indicated in block 128. Those selected page groups are then compacted as indicated in block 130. Then a check at diamond 132 determines whether there are still additional compaction is desired.

If so, the RCS is consulted to find page groups that can be most compactly stored as indicated in block 134. Other techniques may also be used.

[0062] Since the compaction daemon operates asynchronously and independently in some embodiments, there are no complex interactions with other caches (e.g., forced evicts etc.). In comparison, a buffered compaction approach may have such intricacies. The daemon simply replaces groups of finished pages by compacted pages, at its own pace.

[0063] The compaction daemon selects page groups (e.g. 2KB blocks) for compaction based on an array of policies such as shown in Figure 10. For example, page groups that are tagged as read-only data by the driver are suitable for first selection, as such data may be compacted without risk of later being written to (and thus expanded).

[0064] Followed by this, page groups that are “old” by some metric can be chosen. For example, least-recently-written (LRW) is a useful metric, as such page groups are more likely to have a longer timespan before the next write (possibly never).

[0065] To select between different page groups, the compaction daemon can look at the associated render control surface (RCS) to see which page group can be most compactly stored (i.e., least number of pages after compaction).

[0066] Since the graphics driver and/or user application has additional high-level information, including the possibility to look-ahead in the command stream, hints may be sent to the compaction daemon to improve system performance.

[0067] A typical example is a render target that is later used as a texture, and hence goes from being read/write (R/W) memory to read-only. Another example is tiled rendering, where tiles that have been finished will not be written to in a long time (e.g., next frame). Such regions may be prioritized for compaction.

[0068] Discard commands available in the latest application program interfaces (APIs) (e.g. DX12) may also be used by the compaction daemon to asynchronously clear pages that are no longer needed, rather than compacting or evicting them to FM. Note this can be done similarly with a buffered compaction approach.

[0069] In practice, communicating information from the software stack to the compaction daemon may be done by a message passing mechanism directly to the daemon, which internally can keep a queue of page groups suitable for compaction, or by having additional control bits in the PGCS that can be written directly by the NM controller through commands sent to it.

[0070] The daemon performs less work than the buffered compaction (BC) approach. However, since BC performs compaction all the time, its compaction ratio can be better, but at a cost. For example, BC may consume more on-chip bandwidth since it performs more compaction operations, and each compaction can potentially trigger a read-back from NM.

[0071] In some embodiments both the compression buffer and the compaction daemon may be used. For example they may work in parallel.

[0072] **Figure 10** is a block diagram of a processing system 100, according to an embodiment. In various embodiments the system 100 includes one or more processors 102 and one or more graphics processors 108, and may be a single processor desktop system, a multiprocessor workstation system, or a server system having a large number of processors 102 or processor cores 107. In one embodiment, the system 100 is a processing platform incorporated within a system-on-a-chip (SoC) integrated circuit for use in mobile, handheld, or embedded devices.

[0073] An embodiment of system 100 can include, or be incorporated within a server-based gaming platform, a game console, including a game and media console, a mobile gaming console, a handheld game console, or an online game console. In some embodiments system 100 is a mobile phone, smart phone, tablet computing device or mobile Internet device. Data processing system 100 can also include, couple with, or be integrated within a wearable device, such as a smart watch wearable device, smart eyewear device, augmented reality device, or virtual reality device. In some embodiments, data processing system 100 is a television or set top box device having one or more processors 102 and a graphical interface generated by one or more graphics processors 108.

[0074] In some embodiments, the one or more processors 102 each include one or more processor cores 107 to process instructions which, when executed, perform operations for system and user software. In some embodiments, each of the one or more processor cores 107 is configured to process a specific instruction set 109. In some embodiments, instruction set 109 may facilitate Complex Instruction Set Computing (CISC), Reduced Instruction Set Computing (RISC), or computing via a Very Long Instruction Word (VLIW). Multiple processor cores 107 may each process a different instruction set 109, which may include instructions to facilitate the emulation of other instruction sets. Processor core 107 may also include other processing devices, such a Digital Signal Processor (DSP).

[0075] In some embodiments, the processor 102 includes cache memory 104. Depending on the architecture, the processor 102 can have a single internal cache or multiple levels of internal cache. In some embodiments, the cache memory is shared among various components of the processor 102. In some embodiments, the processor 102 also uses an external cache (e.g., a Level-3 (L3) cache or Last Level Cache (LLC)) (not shown), which may be shared among processor cores 107 using known cache coherency techniques. A register file 106 is additionally included in processor 102 which may include different types of registers for storing different types of data (e.g., integer registers, floating point registers, status registers, and an instruction pointer register). Some registers may be general-purpose registers, while other registers may be specific to the design of the processor 102.

[0076] In some embodiments, processor 102 is coupled to a processor bus 110 to transmit communication signals such as address, data, or control signals between processor 102 and other components in system 100. In one embodiment the system 100 uses an exemplary 'hub' system architecture, including a memory controller hub 116 and an Input Output (I/O) controller hub 130. A memory controller hub 116 facilitates communication between a memory device and other components of system 100, while an I/O Controller Hub (ICH) 130 provides connections to I/O devices via a local I/O bus. In one embodiment, the logic of the memory controller hub 116 is integrated within the processor.

[0077] Memory device 120 can be a dynamic random access memory (DRAM) device, a static random access memory (SRAM) device, flash memory device, phase-change memory device, or some other memory device having suitable performance to serve as process memory. In one embodiment the memory device 120 can operate as system memory for the system 100, to store data 122 and instructions 121 for use when the one or more processors 102 executes an application or process. Memory controller hub 116 also couples with an optional external graphics processor 112, which may communicate with the one or more graphics processors 108 in processors 102 to perform graphics and media operations.

[0078] In some embodiments, ICH 130 enables peripherals to connect to memory device 120 and processor 102 via a high-speed I/O bus. The I/O peripherals include, but are not limited to, an audio controller 146, a firmware interface 128, a wireless transceiver 126 (e.g., Wi-Fi, Bluetooth), a data storage device 124 (e.g., hard disk drive, flash memory, etc.), and a legacy I/O controller 140 for coupling legacy (e.g., Personal System 2 (PS/2)) devices to the system. One or more Universal Serial Bus (USB) controllers 142 connect input devices, such as keyboard and mouse 144 combinations. A network controller 134 may also couple to ICH 130. In some embodiments, a high-performance network controller (not shown) couples to processor bus 110. It will be appreciated that the system 100 shown is exemplary and not limiting, as other types of data processing systems that are differently configured may also be used. For example, the I/O controller hub 130 may be integrated within the one or more processor 102, or the memory controller hub 116 and I/O controller hub 130 may be integrated into a discreet external graphics processor, such as the external graphics processor 112.

[0079] **Figure 11** is a block diagram of an embodiment of a processor 200 having one or more processor cores 202A-202N, an integrated memory controller 214, and an integrated graphics processor 208. Those elements of **Figure 11** having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such. Processor 200 can include additional cores up to and including

additional core 202N represented by the dashed lined boxes. Each of processor cores 202A-202N includes one or more internal cache units 204A-204N. In some embodiments each processor core also has access to one or more shared cached units 206.

[0080] The internal cache units 204A-204N and shared cache units 206 represent a cache memory hierarchy within the processor 200. The cache memory hierarchy may include at least one level of instruction and data cache within each processor core and one or more levels of shared mid-level cache, such as a Level 2 (L2), Level 3 (L3), Level 4 (L4), or other levels of cache, where the highest level of cache before external memory is classified as the LLC. In some embodiments, cache coherency logic maintains coherency between the various cache units 206 and 204A-204N.

[0081] In some embodiments, processor 200 may also include a set of one or more bus controller units 216 and a system agent core 210. The one or more bus controller units 216 manage a set of peripheral buses, such as one or more Peripheral Component Interconnect buses (e.g., PCI, PCI Express). System agent core 210 provides management functionality for the various processor components. In some embodiments, system agent core 210 includes one or more integrated memory controllers 214 to manage access to various external memory devices (not shown).

[0082] In some embodiments, one or more of the processor cores 202A-202N include support for simultaneous multi-threading. In such embodiment, the system agent core 210 includes components for coordinating and operating cores 202A-202N during multi-threaded processing. System agent core 210 may additionally include a power control unit (PCU), which includes logic and components to regulate the power state of processor cores 202A-202N and graphics processor 208.

[0083] In some embodiments, processor 200 additionally includes graphics processor 208 to execute graphics processing operations. In some embodiments, the graphics processor 208 couples with the set of shared cache units 206, and the system agent core 210, including the one or more integrated memory controllers 214. In some embodiments, a display controller 211 is coupled with the graphics

processor 208 to drive graphics processor output to one or more coupled displays. In some embodiments, display controller 211 may be a separate module coupled with the graphics processor via at least one interconnect, or may be integrated within the graphics processor 208 or system agent core 210.

[0084] In some embodiments, a ring based interconnect unit 212 is used to couple the internal components of the processor 200. However, an alternative interconnect unit may be used, such as a point-to-point interconnect, a switched interconnect, or other techniques, including techniques well known in the art. In some embodiments, graphics processor 208 couples with the ring interconnect 212 via an I/O link 213.

[0085] The exemplary I/O link 213 represents at least one of multiple varieties of I/O interconnects, including an on package I/O interconnect which facilitates communication between various processor components and a high-performance embedded memory module 218, such as an eDRAM module. In some embodiments, each of the processor cores 202-202N and graphics processor 208 use embedded memory modules 218 as a shared Last Level Cache.

[0086] In some embodiments, processor cores 202A-202N are homogenous cores executing the same instruction set architecture. In another embodiment, processor cores 202A-202N are heterogeneous in terms of instruction set architecture (ISA), where one or more of processor cores 202A-N execute a first instruction set, while at least one of the other cores executes a subset of the first instruction set or a different instruction set. In one embodiment processor cores 202A-202N are heterogeneous in terms of microarchitecture, where one or more cores having a relatively higher power consumption couple with one or more power cores having a lower power consumption. Additionally, processor 200 can be implemented on one or more chips or as an SoC integrated circuit having the illustrated components, in addition to other components.

[0087] **Figure 12** is a block diagram of a graphics processor 300, which may be a discrete graphics processing unit, or may be a graphics processor integrated with a plurality of processing cores. In some embodiments, the graphics processor communicates via a memory mapped I/O interface to registers on the graphics

processor and with commands placed into the processor memory. In some embodiments, graphics processor 300 includes a memory interface 314 to access memory. Memory interface 314 can be an interface to local memory, one or more internal caches, one or more shared external caches, and/or to system memory.

[0088] In some embodiments, graphics processor 300 also includes a display controller 302 to drive display output data to a display device 320. Display controller 302 includes hardware for one or more overlay planes for the display and composition of multiple layers of video or user interface elements. In some embodiments, graphics processor 300 includes a video codec engine 306 to encode, decode, or transcode media to, from, or between one or more media encoding formats, including, but not limited to Moving Picture Experts Group (MPEG) formats such as MPEG-2, Advanced Video Coding (AVC) formats such as H.264/MPEG-4 AVC, as well as the Society of Motion Picture & Television Engineers (SMPTE) 421M/VC-1, and Joint Photographic Experts Group (JPEG) formats such as JPEG, and Motion JPEG (MJPEG) formats.

[0089] In some embodiments, graphics processor 300 includes a block image transfer (BLIT) engine 304 to perform two-dimensional (2D) rasterizer operations including, for example, bit-boundary block transfers. However, in one embodiment, 2D graphics operations are performed using one or more components of graphics processing engine (GPE) 310. In some embodiments, graphics processing engine 310 is a compute engine for performing graphics operations, including three-dimensional (3D) graphics operations and media operations.

[0090] In some embodiments, GPE 310 includes a 3D pipeline 312 for performing 3D operations, such as rendering three-dimensional images and scenes using processing functions that act upon 3D primitive shapes (e.g., rectangle, triangle, etc.). The 3D pipeline 312 includes programmable and fixed function elements that perform various tasks within the element and/or spawn execution threads to a 3D/Media sub-system 315. While 3D pipeline 312 can be used to perform media operations, an embodiment of GPE 310 also includes a media pipeline 316 that is specifically used to perform media operations, such as video post-processing and image enhancement.

[0091] In some embodiments, media pipeline 316 includes fixed function or programmable logic units to perform one or more specialized media operations, such as video decode acceleration, video de-interlacing, and video encode acceleration in place of, or on behalf of video codec engine 306. In some embodiments, media pipeline 316 additionally includes a thread spawning unit to spawn threads for execution on 3D/Media sub-system 315. The spawned threads perform computations for the media operations on one or more graphics execution units included in 3D/Media sub-system 315.

[0092] In some embodiments, 3D/Media subsystem 315 includes logic for executing threads spawned by 3D pipeline 312 and media pipeline 316. In one embodiment, the pipelines send thread execution requests to 3D/Media subsystem 315, which includes thread dispatch logic for arbitrating and dispatching the various requests to available thread execution resources. The execution resources include an array of graphics execution units to process the 3D and media threads. In some embodiments, 3D/Media subsystem 315 includes one or more internal caches for thread instructions and data. In some embodiments, the subsystem also includes shared memory, including registers and addressable memory, to share data between threads and to store output data.

[0093] **Figure 13** is a block diagram of a graphics processing engine 410 of a graphics processor in accordance with some embodiments. In one embodiment, the GPE 410 is a version of the GPE 310 shown in **Figure 12**. Elements of **Figure 13** having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such.

[0094] In some embodiments, GPE 410 couples with a command streamer 403, which provides a command stream to the GPE 3D and media pipelines 412, 416. In some embodiments, command streamer 403 is coupled to memory, which can be system memory, or one or more of internal cache memory and shared cache memory. In some embodiments, command streamer 403 receives commands from the memory and sends the commands to 3D pipeline 412 and/or media pipeline 416. The commands are directives fetched from a ring buffer, which stores commands for

the 3D and media pipelines 412, 416. In one embodiment, the ring buffer can additionally include batch command buffers storing batches of multiple commands. The 3D and media pipelines 412, 416 process the commands by performing operations via logic within the respective pipelines or by dispatching one or more execution threads to an execution unit array 414. In some embodiments, execution unit array 414 is scalable, such that the array includes a variable number of execution units based on the target power and performance level of GPE 410.

[0095] In some embodiments, a sampling engine 430 couples with memory (e.g., cache memory or system memory) and execution unit array 414. In some embodiments, sampling engine 430 provides a memory access mechanism for execution unit array 414 that allows execution array 414 to read graphics and media data from memory. In some embodiments, sampling engine 430 includes logic to perform specialized image sampling operations for media.

[0096] In some embodiments, the specialized media sampling logic in sampling engine 430 includes a de-noise/de-interlace module 432, a motion estimation module 434, and an image scaling and filtering module 436. In some embodiments, de-noise/de-interlace module 432 includes logic to perform one or more of a de-noise or a de-interlace algorithm on decoded video data. The de-interlace logic combines alternating fields of interlaced video content into a single frame of video. The de-noise logic reduces or removes data noise from video and image data. In some embodiments, the de-noise logic and de-interlace logic are motion adaptive and use spatial or temporal filtering based on the amount of motion detected in the video data. In some embodiments, the de-noise/de-interlace module 432 includes dedicated motion detection logic (e.g., within the motion estimation engine 434).

[0097] In some embodiments, motion estimation engine 434 provides hardware acceleration for video operations by performing video acceleration functions such as motion vector estimation and prediction on video data. The motion estimation engine determines motion vectors that describe the transformation of image data between successive video frames. In some embodiments, a graphics processor media codec uses video motion estimation engine 434 to perform operations on video at the macro-block level that may otherwise be too computationally intensive to

perform with a general-purpose processor. In some embodiments, motion estimation engine 434 is generally available to graphics processor components to assist with video decode and processing functions that are sensitive or adaptive to the direction or magnitude of the motion within video data.

[0098] In some embodiments, image scaling and filtering module 436 performs image-processing operations to enhance the visual quality of generated images and video. In some embodiments, scaling and filtering module 436 processes image and video data during the sampling operation before providing the data to execution unit array 414.

[0099] In some embodiments, the GPE 410 includes a data port 444, which provides an additional mechanism for graphics subsystems to access memory. In some embodiments, data port 444 facilitates memory access for operations including render target writes, constant buffer reads, scratch memory space reads/writes, and media surface accesses. In some embodiments, data port 444 includes cache memory space to cache accesses to memory. The cache memory can be a single data cache or separated into multiple caches for the multiple subsystems that access memory via the data port (e.g., a render buffer cache, a constant buffer cache, etc.). In some embodiments, threads executing on an execution unit in execution unit array 414 communicate with the data port by exchanging messages via a data distribution interconnect that couples each of the sub-systems of GPE 410.

[0100] **Figure 14** is a block diagram of another embodiment of a graphics processor 500. Elements of **Figure 14** having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such.

[0101] In some embodiments, graphics processor 500 includes a ring interconnect 502, a pipeline front-end 504, a media engine 537, and graphics cores 580A-580N. In some embodiments, ring interconnect 502 couples the graphics processor to other processing units, including other graphics processors or one or more general-purpose processor cores. In some embodiments, the graphics processor is one of many processors integrated within a multi-core processing system.

[0102] In some embodiments, graphics processor 500 receives batches of commands via ring interconnect 502. The incoming commands are interpreted by a command streamer 503 in the pipeline front-end 504. In some embodiments, graphics processor 500 includes scalable execution logic to perform 3D geometry processing and media processing via the graphics core(s) 580A-580N. For 3D geometry processing commands, command streamer 503 supplies commands to geometry pipeline 536. For at least some media processing commands, command streamer 503 supplies the commands to a video front end 534, which couples with a media engine 537. In some embodiments, media engine 537 includes a Video Quality Engine (VQE) 530 for video and image post-processing and a multi-format encode/decode (MFX) 533 engine to provide hardware-accelerated media data encode and decode. In some embodiments, geometry pipeline 536 and media engine 537 each generate execution threads for the thread execution resources provided by at least one graphics core 580A.

[0103] In some embodiments, graphics processor 500 includes scalable thread execution resources featuring modular cores 580A-580N (sometimes referred to as core slices), each having multiple sub-cores 550A-550N, 560A-560N (sometimes referred to as core sub-slices). In some embodiments, graphics processor 500 can have any number of graphics cores 580A through 580N. In some embodiments, graphics processor 500 includes a graphics core 580A having at least a first sub-core 550A and a second core sub-core 560A. In other embodiments, the graphics processor is a low power processor with a single sub-core (e.g., 550A). In some embodiments, graphics processor 500 includes multiple graphics cores 580A-580N, each including a set of first sub-cores 550A-550N and a set of second sub-cores 560A-560N. Each sub-core in the set of first sub-cores 550A-550N includes at least a first set of execution units 552A-552N and media/texture samplers 554A-554N. Each sub-core in the set of second sub-cores 560A-560N includes at least a second set of execution units 562A-562N and samplers 564A-564N. In some embodiments, each sub-core 550A-550N, 560A-560N shares a set of shared resources 570A-570N. In some embodiments, the shared resources include shared cache memory and pixel operation logic. Other shared resources may also be included in the various embodiments of the graphics processor.

[0104] **Figure 15** illustrates thread execution logic 600 including an array of processing elements employed in some embodiments of a GPE. Elements of **Figure 15** having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such.

[0105] In some embodiments, thread execution logic 600 includes a pixel shader 602, a thread dispatcher 604, instruction cache 606, a scalable execution unit array including a plurality of execution units 608A-608N, a sampler 610, a data cache 612, and a data port 614. In one embodiment the included components are interconnected via an interconnect fabric that links to each of the components. In some embodiments, thread execution logic 600 includes one or more connections to memory, such as system memory or cache memory, through one or more of instruction cache 606, data port 614, sampler 610, and execution unit array 608A-608N. In some embodiments, each execution unit (e.g. 608A) is an individual vector processor capable of executing multiple simultaneous threads and processing multiple data elements in parallel for each thread. In some embodiments, execution unit array 608A-608N includes any number individual execution units.

[0106] In some embodiments, execution unit array 608A-608N is primarily used to execute “shader” programs. In some embodiments, the execution units in array 608A-608N execute an instruction set that includes native support for many standard 3D graphics shader instructions, such that shader programs from graphics libraries (e.g., Direct 3D and OpenGL) are executed with a minimal translation. The execution units support vertex and geometry processing (e.g., vertex programs, geometry programs, vertex shaders), pixel processing (e.g., pixel shaders, fragment shaders) and general-purpose processing (e.g., compute and media shaders).

[0107] Each execution unit in execution unit array 608A-608N operates on arrays of data elements. The number of data elements is the “execution size,” or the number of channels for the instruction. An execution channel is a logical unit of execution for data element access, masking, and flow control within instructions. The number of channels may be independent of the number of physical Arithmetic Logic Units (ALUs) or Floating Point Units (FPUs) for a particular graphics processor. In some

embodiments, execution units 608A-608N support integer and floating-point data types.

[0108] The execution unit instruction set includes single instruction multiple data (SIMD) instructions. The various data elements can be stored as a packed data type in a register and the execution unit will process the various elements based on the data size of the elements. For example, when operating on a 256-bit wide vector, the 256 bits of the vector are stored in a register and the execution unit operates on the vector as four separate 64-bit packed data elements (Quad-Word (QW) size data elements), eight separate 32-bit packed data elements (Double Word (DW) size data elements), sixteen separate 16-bit packed data elements (Word (W) size data elements), or thirty-two separate 8-bit data elements (byte (B) size data elements). However, different vector widths and register sizes are possible.

[0109] One or more internal instruction caches (e.g., 606) are included in the thread execution logic 600 to cache thread instructions for the execution units. In some embodiments, one or more data caches (e.g., 612) are included to cache thread data during thread execution. In some embodiments, sampler 610 is included to provide texture sampling for 3D operations and media sampling for media operations. In some embodiments, sampler 610 includes specialized texture or media sampling functionality to process texture or media data during the sampling process before providing the sampled data to an execution unit.

[0110] During execution, the graphics and media pipelines send thread initiation requests to thread execution logic 600 via thread spawning and dispatch logic. In some embodiments, thread execution logic 600 includes a local thread dispatcher 604 that arbitrates thread initiation requests from the graphics and media pipelines and instantiates the requested threads on one or more execution units 608A-608N. For example, the geometry pipeline (e.g., 536 of **Figure 14**) dispatches vertex processing, tessellation, or geometry processing threads to thread execution logic 600 (**Figure 15**). In some embodiments, thread dispatcher 604 can also process runtime thread spawning requests from the executing shader programs.

[0111] Once a group of geometric objects has been processed and rasterized into pixel data, pixel shader 602 is invoked to further compute output information and cause results to be written to output surfaces (e.g., color buffers, depth buffers, stencil buffers, etc.). In some embodiments, pixel shader 602 calculates the values of the various vertex attributes that are to be interpolated across the rasterized object. In some embodiments, pixel shader 602 then executes an application programming interface (API)-supplied pixel shader program. To execute the pixel shader program, pixel shader 602 dispatches threads to an execution unit (e.g., 608A) via thread dispatcher 604. In some embodiments, pixel shader 602 uses texture sampling logic in sampler 610 to access texture data in texture maps stored in memory. Arithmetic operations on the texture data and the input geometry data compute pixel color data for each geometric fragment, or discards one or more pixels from further processing.

[0112] In some embodiments, the data port 614 provides a memory access mechanism for the thread execution logic 600 output processed data to memory for processing on a graphics processor output pipeline. In some embodiments, the data port 614 includes or couples to one or more cache memories (e.g., data cache 612) to cache data for memory access via the data port.

[0113] **Figure 16** is a block diagram illustrating a graphics processor instruction formats 700 according to some embodiments. In one or more embodiment, the graphics processor execution units support an instruction set having instructions in multiple formats. The solid lined boxes illustrate the components that are generally included in an execution unit instruction, while the dashed lines include components that are optional or that are only included in a sub-set of the instructions. In some embodiments, instruction format 700 described and illustrated are macro-instructions, in that they are instructions supplied to the execution unit, as opposed to micro-operations resulting from instruction decode once the instruction is processed.

[0114] In some embodiments, the graphics processor execution units natively support instructions in a 128-bit format 710. A 64-bit compacted instruction format 730 is available for some instructions based on the selected instruction, instruction

options, and number of operands. The native 128-bit format 710 provides access to all instruction options, while some options and operations are restricted in the 64-bit format 730. The native instructions available in the 64-bit format 730 vary by embodiment. In some embodiments, the instruction is compacted in part using a set of index values in an index field 713. The execution unit hardware references a set of compaction tables based on the index values and uses the compaction table outputs to reconstruct a native instruction in the 128-bit format 710.

[0115] For each format, instruction opcode 712 defines the operation that the execution unit is to perform. The execution units execute each instruction in parallel across the multiple data elements of each operand. For example, in response to an add instruction the execution unit performs a simultaneous add operation across each color channel representing a texture element or picture element. By default, the execution unit performs each instruction across all data channels of the operands. In some embodiments, instruction control field 714 enables control over certain execution options, such as channels selection (e.g., predication) and data channel order (e.g., swizzle). For 128-bit instructions 710 an exec-size field 716 limits the number of data channels that will be executed in parallel. In some embodiments, exec-size field 716 is not available for use in the 64-bit compact instruction format 730.

[0116] Some execution unit instructions have up to three operands including two source operands, src0 722, src1 722, and one destination 718. In some embodiments, the execution units support dual destination instructions, where one of the destinations is implied. Data manipulation instructions can have a third source operand (e.g., SRC2 724), where the instruction opcode 712 determines the number of source operands. An instruction's last source operand can be an immediate (e.g., hard-coded) value passed with the instruction.

[0117] In some embodiments, the 128-bit instruction format 710 includes an access/address mode information 726 specifying, for example, whether direct register addressing mode or indirect register addressing mode is used. When direct register addressing mode is used, the register address of one or more operands is directly provided by bits in the instruction 710.

[0118] In some embodiments, the 128-bit instruction format 710 includes an access/address mode field 726, which specifies an address mode and/or an access mode for the instruction. In one embodiment the access mode to define a data access alignment for the instruction. Some embodiments support access modes including a 16-byte aligned access mode and a 1-byte aligned access mode, where the byte alignment of the access mode determines the access alignment of the instruction operands. For example, when in a first mode, the instruction 710 may use byte-aligned addressing for source and destination operands and when in a second mode, the instruction 710 may use 16-byte-aligned addressing for all source and destination operands.

[0119] In one embodiment, the address mode portion of the access/address mode field 726 determines whether the instruction is to use direct or indirect addressing. When direct register addressing mode is used bits in the instruction 710 directly provide the register address of one or more operands. When indirect register addressing mode is used, the register address of one or more operands may be computed based on an address register value and an address immediate field in the instruction.

[0120] In some embodiments instructions are grouped based on opcode 712 bit-fields to simplify Opcode decode 740. For an 8-bit opcode, bits 4, 5, and 6 allow the execution unit to determine the type of opcode. The precise opcode grouping shown is merely an example. In some embodiments, a move and logic opcode group 742 includes data movement and logic instructions (e.g., move (mov), compare (cmp)). In some embodiments, move and logic group 742 shares the five most significant bits (MSB), where move (mov) instructions are in the form of 0000xxxxb and logic instructions are in the form of 0001xxxxb. A flow control instruction group 744 (e.g., call, jump (jmp)) includes instructions in the form of 0010xxxxb (e.g., 0x20). A miscellaneous instruction group 746 includes a mix of instructions, including synchronization instructions (e.g., wait, send) in the form of 0011xxxxb (e.g., 0x30). A parallel math instruction group 748 includes component-wise arithmetic instructions (e.g., add, multiply (mul)) in the form of 0100xxxxb (e.g., 0x40). The parallel math group 748 performs the arithmetic operations in parallel across data

channels. The vector math group 750 includes arithmetic instructions (e.g., dp4) in the form of 0101xxxxb (e.g., 0x50). The vector math group performs arithmetic such as dot product calculations on vector operands.

[0121] **Figure 17** is a block diagram of another embodiment of a graphics processor 800. Elements of **Figure 17** having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such.

[0122] In some embodiments, graphics processor 800 includes a graphics pipeline 820, a media pipeline 830, a display engine 840, thread execution logic 850, and a render output pipeline 870. In some embodiments, graphics processor 800 is a graphics processor within a multi-core processing system that includes one or more general purpose processing cores. The graphics processor is controlled by register writes to one or more control registers (not shown) or via commands issued to graphics processor 800 via a ring interconnect 802. In some embodiments, ring interconnect 802 couples graphics processor 800 to other processing components, such as other graphics processors or general-purpose processors. Commands from ring interconnect 802 are interpreted by a command streamer 803, which supplies instructions to individual components of graphics pipeline 820 or media pipeline 830.

[0123] In some embodiments, command streamer 803 directs the operation of a vertex fetcher 805 that reads vertex data from memory and executes vertex-processing commands provided by command streamer 803. In some embodiments, vertex fetcher 805 provides vertex data to a vertex shader 807, which performs coordinate space transformation and lighting operations to each vertex. In some embodiments, vertex fetcher 805 and vertex shader 807 execute vertex-processing instructions by dispatching execution threads to execution units 852A, 852B via a thread dispatcher 831.

[0124] In some embodiments, execution units 852A, 852B are an array of vector processors having an instruction set for performing graphics and media operations. In some embodiments, execution units 852A, 852B have an attached L1 cache 851 that is specific for each array or shared between the arrays. The cache can be

configured as a data cache, an instruction cache, or a single cache that is partitioned to contain data and instructions in different partitions.

[0125] In some embodiments, graphics pipeline 820 includes tessellation components to perform hardware-accelerated tessellation of 3D objects. In some embodiments, a programmable hull shader 811 configures the tessellation operations. A programmable domain shader 817 provides back-end evaluation of tessellation output. A tessellator 813 operates at the direction of hull shader 811 and contains special purpose logic to generate a set of detailed geometric objects based on a coarse geometric model that is provided as input to graphics pipeline 820. In some embodiments, if tessellation is not used, tessellation components 811, 813, 817 can be bypassed.

[0126] In some embodiments, complete geometric objects can be processed by a geometry shader 819 via one or more threads dispatched to execution units 852A, 852B, or can proceed directly to the clipper 829. In some embodiments, the geometry shader operates on entire geometric objects, rather than vertices or patches of vertices as in previous stages of the graphics pipeline. If the tessellation is disabled the geometry shader 819 receives input from the vertex shader 807. In some embodiments, geometry shader 819 is programmable by a geometry shader program to perform geometry tessellation if the tessellation units are disabled.

[0127] Before rasterization, a clipper 829 processes vertex data. The clipper 829 may be a fixed function clipper or a programmable clipper having clipping and geometry shader functions. In some embodiments, a rasterizer/depth 873 in the render output pipeline 870 dispatches pixel shaders to convert the geometric objects into their per pixel representations. In some embodiments, pixel shader logic is included in thread execution logic 850. In some embodiments, an application can bypass the rasterizer 873 and access un-rasterized vertex data via a stream out unit 823.

[0128] The graphics processor 800 has an interconnect bus, interconnect fabric, or some other interconnect mechanism that allows data and message passing amongst the major components of the processor. In some embodiments, execution units

852A, 852B and associated cache(s) 851, texture and media sampler 854, and texture/sampler cache 858 interconnect via a data port 856 to perform memory access and communicate with render output pipeline components of the processor. In some embodiments, sampler 854, caches 851, 858 and execution units 852A, 852B each have separate memory access paths.

[0129] In some embodiments, render output pipeline 870 contains a rasterizer and depth test component 873 that converts vertex-based objects into an associated pixel-based representation. In some embodiments, the rasterizer logic includes a windower/masker unit to perform fixed function triangle and line rasterization. An associated render cache 878 and depth cache 879 are also available in some embodiments. A pixel operations component 877 performs pixel-based operations on the data, though in some instances, pixel operations associated with 2D operations (e.g. bit block image transfers with blending) are performed by the 2D engine 841, or substituted at display time by the display controller 843 using overlay display planes. In some embodiments, a shared L3 cache 875 is available to all graphics components, allowing the sharing of data without the use of main system memory.

[0130] In some embodiments, graphics processor media pipeline 830 includes a media engine 837 and a video front end 834. In some embodiments, video front end 834 receives pipeline commands from the command streamer 803. In some embodiments, media pipeline 830 includes a separate command streamer. In some embodiments, video front-end 834 processes media commands before sending the command to the media engine 837. In some embodiments, media engine 337 includes thread spawning functionality to spawn threads for dispatch to thread execution logic 850 via thread dispatcher 831.

[0131] In some embodiments, graphics processor 800 includes a display engine 840. In some embodiments, display engine 840 is external to processor 800 and couples with the graphics processor via the ring interconnect 802, or some other interconnect bus or fabric. In some embodiments, display engine 840 includes a 2D engine 841 and a display controller 843. In some embodiments, display engine 840 contains special purpose logic capable of operating independently of the 3D pipeline.

In some embodiments, display controller 843 couples with a display device (not shown), which may be a system integrated display device, as in a laptop computer, or an external display device attached via a display device connector.

[0132] In some embodiments, graphics pipeline 820 and media pipeline 830 are configurable to perform operations based on multiple graphics and media programming interfaces and are not specific to any one application programming interface (API). In some embodiments, driver software for the graphics processor translates API calls that are specific to a particular graphics or media library into commands that can be processed by the graphics processor. In some embodiments, support is provided for the Open Graphics Library (OpenGL) and Open Computing Language (OpenCL) from the Khronos Group, the Direct3D library from the Microsoft Corporation, or support may be provided to both OpenGL and D3D. Support may also be provided for the Open Source Computer Vision Library (OpenCV). A future API with a compatible 3D pipeline would also be supported if a mapping can be made from the pipeline of the future API to the pipeline of the graphics processor.

[0133] **Figure 18A** is a block diagram illustrating a graphics processor command format 900 according to some embodiments. **Figure 18B** is a block diagram illustrating a graphics processor command sequence 910 according to an embodiment. The solid lined boxes in **Figure 18A** illustrate the components that are generally included in a graphics command while the dashed lines include components that are optional or that are only included in a sub-set of the graphics commands. The exemplary graphics processor command format 900 of **Figure 18A** includes data fields to identify a target client 902 of the command, a command operation code (opcode) 904, and the relevant data 906 for the command. A sub-opcode 905 and a command size 908 are also included in some commands.

[0134] In some embodiments, client 902 specifies the client unit of the graphics device that processes the command data. In some embodiments, a graphics processor command parser examines the client field of each command to condition the further processing of the command and route the command data to the appropriate client unit. In some embodiments, the graphics processor client units

include a memory interface unit, a render unit, a 2D unit, a 3D unit, and a media unit. Each client unit has a corresponding processing pipeline that processes the commands. Once the command is received by the client unit, the client unit reads the opcode 904 and, if present, sub-opcode 905 to determine the operation to perform. The client unit performs the command using information in data field 906. For some commands an explicit command size 908 is expected to specify the size of the command. In some embodiments, the command parser automatically determines the size of at least some of the commands based on the command opcode. In some embodiments commands are aligned via multiples of a double word.

[0135] The flow diagram in **Figure 18B** shows an exemplary graphics processor command sequence 910. In some embodiments, software or firmware of a data processing system that features an embodiment of a graphics processor uses a version of the command sequence shown to set up, execute, and terminate a set of graphics operations. A sample command sequence is shown and described for purposes of example only as embodiments are not limited to these specific commands or to this command sequence. Moreover, the commands may be issued as batch of commands in a command sequence, such that the graphics processor will process the sequence of commands in at least partially concurrence.

[0136] In some embodiments, the graphics processor command sequence 910 may begin with a pipeline flush command 912 to cause any active graphics pipeline to complete the currently pending commands for the pipeline. In some embodiments, the 3D pipeline 922 and the media pipeline 924 do not operate concurrently. The pipeline flush is performed to cause the active graphics pipeline to complete any pending commands. In response to a pipeline flush, the command parser for the graphics processor will pause command processing until the active drawing engines complete pending operations and the relevant read caches are invalidated. Optionally, any data in the render cache that is marked 'dirty' can be flushed to memory. In some embodiments, pipeline flush command 912 can be used for pipeline synchronization or before placing the graphics processor into a low power state.

[0137] In some embodiments, a pipeline select command 913 is used when a command sequence requires the graphics processor to explicitly switch between pipelines. In some embodiments, a pipeline select command 913 is required only once within an execution context before issuing pipeline commands unless the context is to issue commands for both pipelines. In some embodiments, a pipeline flush command is 912 is required immediately before a pipeline switch via the pipeline select command 913.

[0138] In some embodiments, a pipeline control command 914 configures a graphics pipeline for operation and is used to program the 3D pipeline 922 and the media pipeline 924. In some embodiments, pipeline control command 914 configures the pipeline state for the active pipeline. In one embodiment, the pipeline control command 914 is used for pipeline synchronization and to clear data from one or more cache memories within the active pipeline before processing a batch of commands.

[0139] In some embodiments, return buffer state commands 916 are used to configure a set of return buffers for the respective pipelines to write data. Some pipeline operations require the allocation, selection, or configuration of one or more return buffers into which the operations write intermediate data during processing. In some embodiments, the graphics processor also uses one or more return buffers to store output data and to perform cross thread communication. In some embodiments, the return buffer state 916 includes selecting the size and number of return buffers to use for a set of pipeline operations.

[0140] The remaining commands in the command sequence differ based on the active pipeline for operations. Based on a pipeline determination 920, the command sequence is tailored to the 3D pipeline 922 beginning with the 3D pipeline state 930, or the media pipeline 924 beginning at the media pipeline state 940.

[0141] The commands for the 3D pipeline state 930 include 3D state setting commands for vertex buffer state, vertex element state, constant color state, depth buffer state, and other state variables that are to be configured before 3D primitive commands are processed. The values of these commands are determined at least

in part based the particular 3D API in use. In some embodiments, 3D pipeline state 930 commands are also able to selectively disable or bypass certain pipeline elements if those elements will not be used.

[0142] In some embodiments, 3D primitive 932 command is used to submit 3D primitives to be processed by the 3D pipeline. Commands and associated parameters that are passed to the graphics processor via the 3D primitive 932 command are forwarded to the vertex fetch function in the graphics pipeline. The vertex fetch function uses the 3D primitive 932 command data to generate vertex data structures. The vertex data structures are stored in one or more return buffers. In some embodiments, 3D primitive 932 command is used to perform vertex operations on 3D primitives via vertex shaders. To process vertex shaders, 3D pipeline 922 dispatches shader execution threads to graphics processor execution units.

[0143] In some embodiments, 3D pipeline 922 is triggered via an execute 934 command or event. In some embodiments, a register write triggers command execution. In some embodiments execution is triggered via a 'go' or 'kick' command in the command sequence. In one embodiment command execution is triggered using a pipeline synchronization command to flush the command sequence through the graphics pipeline. The 3D pipeline will perform geometry processing for the 3D primitives. Once operations are complete, the resulting geometric objects are rasterized and the pixel engine colors the resulting pixels. Additional commands to control pixel shading and pixel back end operations may also be included for those operations.

[0144] In some embodiments, the graphics processor command sequence 910 follows the media pipeline 924 path when performing media operations. In general, the specific use and manner of programming for the media pipeline 924 depends on the media or compute operations to be performed. Specific media decode operations may be offloaded to the media pipeline during media decode. In some embodiments, the media pipeline can also be bypassed and media decode can be performed in whole or in part using resources provided by one or more general purpose processing cores. In one embodiment, the media pipeline also includes

elements for general-purpose graphics processor unit (GPGPU) operations, where the graphics processor is used to perform SIMD vector operations using computational shader programs that are not explicitly related to the rendering of graphics primitives.

[0145] In some embodiments, media pipeline 924 is configured in a similar manner as the 3D pipeline 922. A set of media pipeline state commands 940 are dispatched or placed into in a command queue before the media object commands 942. In some embodiments, media pipeline state commands 940 include data to configure the media pipeline elements that will be used to process the media objects. This includes data to configure the video decode and video encode logic within the media pipeline, such as encode or decode format. In some embodiments, media pipeline state commands 940 also support the use one or more pointers to “indirect” state elements that contain a batch of state settings.

[0146] In some embodiments, media object commands 942 supply pointers to media objects for processing by the media pipeline. The media objects include memory buffers containing video data to be processed. In some embodiments, all media pipeline states must be valid before issuing a media object command 942. Once the pipeline state is configured and media object commands 942 are queued, the media pipeline 924 is triggered via an execute command 944 or an equivalent execute event (e.g., register write). Output from media pipeline 924 may then be post processed by operations provided by the 3D pipeline 922 or the media pipeline 924. In some embodiments, GPGPU operations are configured and executed in a similar manner as media operations.

[0147] **Figure 19** illustrates exemplary graphics software architecture for a data processing system 1000 according to some embodiments. In some embodiments, software architecture includes a 3D graphics application 1010, an operating system 1020, and at least one processor 1030. In some embodiments, processor 1030 includes a graphics processor 1032 and one or more general-purpose processor core(s) 1034. The graphics application 1010 and operating system 1020 each execute in the system memory 1050 of the data processing system.

[0148] In some embodiments, 3D graphics application 1010 contains one or more shader programs including shader instructions 1012. The shader language instructions may be in a high-level shader language, such as the High Level Shader Language (HLSL) or the OpenGL Shader Language (GLSL). The application also includes executable instructions 1014 in a machine language suitable for execution by the general-purpose processor core 1034. The application also includes graphics objects 1016 defined by vertex data.

[0149] In some embodiments, operating system 1020 is a Microsoft® Windows® operating system from the Microsoft Corporation, a proprietary UNIX-like operating system, or an open source UNIX-like operating system using a variant of the Linux kernel. When the Direct3D API is in use, the operating system 1020 uses a front-end shader compiler 1024 to compile any shader instructions 1012 in HLSL into a lower-level shader language. The compilation may be a just-in-time (JIT) compilation or the application can perform shader pre-compilation. In some embodiments, high-level shaders are compiled into low-level shaders during the compilation of the 3D graphics application 1010.

[0150] In some embodiments, user mode graphics driver 1026 contains a back-end shader compiler 1027 to convert the shader instructions 1012 into a hardware specific representation. When the OpenGL API is in use, shader instructions 1012 in the GLSL high-level language are passed to a user mode graphics driver 1026 for compilation. In some embodiments, user mode graphics driver 1026 uses operating system kernel mode functions 1028 to communicate with a kernel mode graphics driver 1029. In some embodiments, kernel mode graphics driver 1029 communicates with graphics processor 1032 to dispatch commands and instructions.

[0151] One or more aspects of at least one embodiment may be implemented by representative code stored on a machine-readable medium which represents and/or defines logic within an integrated circuit such as a processor. For example, the machine-readable medium may include instructions which represent various logic within the processor. When read by a machine, the instructions may cause the machine to fabricate the logic to perform the techniques described herein. Such representations, known as "IP cores," are reusable units of logic for an integrated

circuit that may be stored on a tangible, machine-readable medium as a hardware model that describes the structure of the integrated circuit. The hardware model may be supplied to various customers or manufacturing facilities, which load the hardware model on fabrication machines that manufacture the integrated circuit. The integrated circuit may be fabricated such that the circuit performs operations described in association with any of the embodiments described herein.

[0152] **Figure 20** is a block diagram illustrating an IP core development system 1100 that may be used to manufacture an integrated circuit to perform operations according to an embodiment. The IP core development system 1100 may be used to generate modular, re-usable designs that can be incorporated into a larger design or used to construct an entire integrated circuit (e.g., an SOC integrated circuit). A design facility 1130 can generate a software simulation 1110 of an IP core design in a high level programming language (e.g., C/C++). The software simulation 1110 can be used to design, test, and verify the behavior of the IP core. A register transfer level (RTL) design can then be created or synthesized from the simulation model 1100. The RTL design 1115 is an abstraction of the behavior of the integrated circuit that models the flow of digital signals between hardware registers, including the associated logic performed using the modeled digital signals. In addition to an RTL design 1115, lower-level designs at the logic level or transistor level may also be created, designed, or synthesized. Thus, the particular details of the initial design and simulation may vary.

[0153] The RTL design 1115 or equivalent may be further synthesized by the design facility into a hardware model 1120, which may be in a hardware description language (HDL), or some other representation of physical design data. The HDL may be further simulated or tested to verify the IP core design. The IP core design can be stored for delivery to a 3rd party fabrication facility 1165 using non-volatile memory 1140 (e.g., hard disk, flash memory, or any non-volatile storage medium). Alternatively, the IP core design may be transmitted (e.g., via the Internet) over a wired connection 1150 or wireless connection 1160. The fabrication facility 1165 may then fabricate an integrated circuit that is based at least in part on the IP core

design. The fabricated integrated circuit can be configured to perform operations in accordance with at least one embodiment described herein.

[0154] **Figure 21** is a block diagram illustrating an exemplary system on a chip integrated circuit 1200 that may be fabricated using one or more IP cores, according to an embodiment. The exemplary integrated circuit includes one or more application processors 1205 (e.g., CPUs), at least one graphics processor 1210, and may additionally include an image processor 1215 and/or a video processor 1220, any of which may be a modular IP core from the same or multiple different design facilities. The integrated circuit includes peripheral or bus logic including a USB controller 1225, UART controller 1230, an SPI/SDIO controller 1235, and an I²S/I²C controller 1240. Additionally, the integrated circuit can include a display device 1245 coupled to one or more of a high-definition multimedia interface (HDMI) controller 1250 and a mobile industry processor interface (MIPI) display interface 1255. Storage may be provided by a flash memory subsystem 1260 including flash memory and a flash memory controller. Memory interface may be provided via a memory controller 1265 for access to SDRAM or SRAM memory devices. Some integrated circuits additionally include an embedded security engine 1270.

[0155] Additionally, other logic and circuits may be included in the processor of integrated circuit 1200, including additional graphics processors/cores, peripheral interface controllers, or general purpose processor cores.

[0156] The following clauses and/or examples pertain to further embodiments:

One example embodiment may be a method comprising providing a near and far memory coupled to a processor, wherein said near memory is faster than said far memory and said near memory is backed by said far memory, compressing data to be stored in said near memory using a codec, compacting said compressed data, and storing said compacted data so that said compressed and compacted data's footprint is less than that of compressed data before compaction. The method may also include compacting before storing compressed data in the near memory. The method may also include compacting a plurality of blocks of cache lines together as an addressable group. The method may also include providing a buffer between

said codec and said near memory, said buffer to store compressed data for said codec, such that said compressed data in said buffer has a smaller footprint than the compressed data from the codec. The method may also include determining whether a plurality of cache lines can be stored as a contiguous block and if so storing said block contiguously in near memory. The method may also include storing said block as an integer multiple of a cache line size in the near memory. The method may also include providing an indication that said block as stored is compacted. The method may also include reading data from said near memory by determining from said indication whether said block is compacted and if so decompressing said block and storing said decompressed chunk in said buffer as an integer multiple of a cache line size. The method may also include compacting said data in the near memory. The method may also include compacting in response to an indication that free memory is needed. The method may also include compacting a selected range of memory pages. The method may also include storing a compacted range of page addresses in an address translation table. The method may also include storing an indication that the range of pages have been compacted. The method may also include compacting using an independent software that can be run at arbitrary points in time. The method may also include storing both compacted and uncompact compressed data in near memory. The method may also include compressing then decompressing on writing to the near memory and reading data by decompressing then decompressing.

[0157] Another example embodiment may be one or more non-transitory computer readable media storing instructions executed by a processor to perform a sequence comprising providing a near and far memory coupled to a processor, wherein said near memory is faster than said far memory and said near memory is backed by said far memory, compressing data to be stored in said near memory using a codec, compacting said compressed data, and storing said compacted data so that said compressed and compacted data's footprint is less than that of compressed data before compaction. The media may include said sequence including compacting before storing compressed data in the near memory. The media may include said sequence including compacting a plurality of blocks of cache lines together as an addressable group. The media may include said sequence including providing a

buffer between said codec and said near memory, said buffer to store compressed data for said codec, such that said compressed data in said buffer has a smaller footprint than the compressed data from the codec. The media may include determining whether a plurality of cache lines can be stored as a contiguous block and if so storing said block contiguously in near memory. The media may include said sequence including storing said block as an integer multiple of a cache line size in the near memory. The media may include said sequence including providing an indication that said block as stored is compacted. The media may include said sequence including reading data from said near memory by determining from said indication whether said block is compacted and if so decompacting said block and storing said decompacted chunk in said buffer as an integer multiple of a cache line size.

[0158] In another example embodiment may be an apparatus comprising a processor, a near and far memory coupled to the processor, wherein said near memory is faster than said far memory and said near memory is backed by said far memory, a codec to compress data to be stored in said near memory, and said processor to compact said compressed data and store said compacted data so that said compressed and compacted data's footprint is less than that of compressed data before compaction. The apparatus may include said codec to compact before storing compressed data in the near memory. The apparatus may include said codec to compact a plurality of blocks of cache lines together as an addressable group. The apparatus may include a buffer between said codec and said near memory, said buffer to store compressed data for said codec, such that said compressed data in said buffer has a smaller footprint than the compressed data from the codec. The apparatus may include said processor to determine whether a plurality of cache lines can be stored as a contiguous block and if so storing said block contiguously in near memory. The apparatus may include said processor to store said block as an integer multiple of a cache line size in the near memory.

[0159] The graphics processing techniques described herein may be implemented in various hardware architectures. For example, graphics functionality may be integrated within a chipset. Alternatively, a discrete graphics processor may be

used. As still another embodiment, the graphics functions may be implemented by a general purpose processor, including a multicore processor.

[0160] References throughout this specification to “one embodiment” or “an embodiment” mean that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one implementation encompassed within the present disclosure. Thus, appearances of the phrase “one embodiment” or “in an embodiment” are not necessarily referring to the same embodiment. Furthermore, the particular features, structures, or characteristics may be instituted in other suitable forms other than the particular embodiment illustrated and all such forms may be encompassed within the claims of the present application.

[0161] While a limited number of embodiments have been described, those skilled in the art will appreciate numerous modifications and variations therefrom. It is intended that the appended claims cover all such modifications and variations as fall within the true spirit and scope of this disclosure.

What is claimed is:

- 1 1. A method comprising:
2 providing a near and far memory coupled to a processor, wherein said near
3 memory is faster than said far memory and said near memory is backed by said far
4 memory;
5 compressing data to be stored in said near memory using a codec;
6 compacting said compressed data; and
7 storing said compacted data so that said compressed and compacted data's
8 footprint is less than that of compressed data before compaction.
- 1 2. The method of claim 1 including compacting before storing compressed data
2 in the near memory.
- 1 3. The method of claim 1 including compacting a plurality of blocks of cache
2 lines together as an addressable group.
- 1 4. The method of claim 1 including providing a buffer between said codec and
2 said near memory, said buffer to store compressed data for said codec, such that
3 said compressed data in said buffer has a smaller footprint than the compressed
4 data from the codec.
- 1 5. The method of claim 4 including determining whether a plurality of cache lines
2 can be stored as a contiguous block and if so storing said block contiguously in near
3 memory.
- 1 6. The method of claim 5 including storing said block as an integer multiple of a
2 cache line size in the near memory.
- 1 7. The method of claim 6 including providing an indication that said block as
2 stored is compacted.

- 1 8. The method of claim 7 including reading data from said near memory by
2 determining from said indication whether said block is compacted and if so
3 decompacting said block and storing said decompacted chunk in said buffer as an
4 integer multiple of a cache line size.
- 1 9. The method of claim 1 including compacting said data in the near memory.
- 1 10. The method of claim 9 including compacting in response to an indication that
2 free memory is needed.
- 1 11. The method of claim 10 including compacting a selected range of memory
2 pages.
- 1 12. The method of claim 10 including storing a compacted range of page
2 addresses in an address translation table.
- 1 13. The method of claim 9 including storing an indication that the range of pages
2 have been compacted.
- 1 14. The method of claim 9 including compacting using an independent software
2 that can be run at arbitrary points in time.
- 1 15. The method of claim 1 including storing both compacted and uncompactd
2 compressed data in near memory.
- 1 16. The method of claim 1 including compressing then decompacting on writing to
2 the near memory and reading data by decompacting then decompressing.
- 1 17. One or more non-transitory computer readable media storing instructions
2 executed by a processor to perform a sequence comprising:

3 providing a near and far memory coupled to a processor, wherein said near
4 memory is faster than said far memory and said near memory is backed by said far
5 memory;

6 compressing data to be stored in said near memory using a codec;

7 compacting said compressed data; and

8 storing said compacted data so that said compressed and compacted data's
9 footprint is less than that of compressed data before compaction.

1 18. The media of claim 17, said sequence including compacting before storing
2 compressed data in the near memory.

1 19. The media of claim 17, said sequence including compacting a plurality of
2 blocks of cache lines together as an addressable group.

1 20. The media of claim 17, said sequence including providing a buffer between
2 said codec and said near memory, said buffer to store compressed data for said
3 codec, such that said compressed data in said buffer has a smaller footprint than the
4 compressed data from the codec.

1 21. The media of claim 20, said sequence including determining whether a
2 plurality of cache lines can be stored as a contiguous block and if so storing said
3 block contiguously in near memory.

1 22. The media of claim 21, said sequence including storing said block as an
2 integer multiple of a cache line size in the near memory.

1 23. The media of claim 22, said sequence including providing an indication that
2 said block as stored is compacted.

1 24. The media of claim 23, said sequence including reading data from said near
2 memory by determining from said indication whether said block is compacted and if

3 so decompacting said block and storing said decompacted chunk in said buffer as an
4 integer multiple of a cache line size.

1 25. An apparatus comprising:
2 a processor;
3 a near and far memory coupled to the processor, wherein said near memory
4 is faster than said far memory and said near memory is backed by said far memory;
5 a codec to compress data to be stored in said near memory; and
6 said processor to compact said compressed data and store said compacted
7 data so that said compressed and compacted data's footprint is less than that of
8 compressed data before compaction.

1 / 18




FIG. 1




FIG. 2

2 / 18




FIG. 3




FIG. 4

3 / 18




FIG. 5

4 / 18




FIG. 6

5 / 18




FIG. 7




FIG. 8

6 / 18




FIG. 9

7 / 18




FIG. 10




FIG. 11




FIG. 12




FIG. 13




FIG. 14





FIG. 15

GRAPHICS CORE INSTRUCTION
FORMATS

700



64-BIT COMPACT
INSTRUCTION 730



OPCODE DECODE
740




FIG. 16

14 / 18




FIG. 17

15 / 18


GRAPHICS PROCESSOR COMMAND FORMAT
900

FIG. 18A


GRAPHICS PROCESSOR COMMAND SEQUENCE
910

FIG. 18B




FIG. 19




FIG. 20




FIG. 21

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2016/019274**A. CLASSIFICATION OF SUBJECT MATTER****G06F 12/08(2006.01)i, G06T 1/60(2006.01)i**

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 12/08; G06F 12/02; G06F 17/30; H04N 7/26; G06F 12/00; G06T 1/60

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: volatile memory, non volatile memory, compression, compaction, codec, buffer

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2012-0290798 A1 (CHENGKUO HUANG et al.) 15 November 2012 See paragraphs [0018], [0023], [0036], [0039]–[0043], [0061]; claims 1, 5, 12; and figures 1, 6.	1, 3, 16–17, 19, 25
Y		2, 9–15, 18
A		4–8, 20–24
Y	US 2009-0327621 A1 (GABRIEL KLIOT et al.) 31 December 2009 See paragraphs [0037]–[0038], [0064], [0066]; claim 2; and figures 4, 8A, 11.	2, 9–15, 18
Y	US 2013-0346378 A1 (DIMITRIOS TSIROGIANNIS et al.) 26 December 2013 See paragraph [0047]; and figure 4.	10–12
A	WO 2014-105323 A1 (APPLE INC.) 03 July 2014 See paragraphs [0070]–[0073]; claim 1; and figure 4.	1–25
A	US 2013-0077673 A1 (ROHIT PURI) 28 March 2013 See paragraphs [0032]–[0033]; claim 1; and figure 3.	1–25



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

28 June 2016 (28.06.2016)

Date of mailing of the international search report

29 June 2016 (29.06.2016)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

CHIN, Sang Bum

Telephone No. +82-42-481-8398



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2016/019274

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2012-0290798 A1	15/11/2012	CN 103620564 A US 8949513 B2 WO 2012-154980 A1	05/03/2014 03/02/2015 15/11/2012
US 2009-0327621 A1	31/12/2009	None	
US 2013-0346378 A1	26/12/2013	CN 104412241 A EP 2864890 A1 US 9053003 B2 WO 2013-192020 A1	11/03/2015 29/04/2015 09/06/2015 27/12/2013
WO 2014-105323 A1	03/07/2014	CN 105027093 A DE 112013006298 T5 KR 10-2015-0090211 A TW 201432456 A TW I494764 B US 2014-0189281 A1	04/11/2015 17/09/2015 05/08/2015 16/08/2014 01/08/2015 03/07/2014
US 2013-0077673 A1	28/03/2013	None	