



- (51) International Patent Classification:
G06T 15/00 (2006.01) *G06T 1/60* (2006.01)
- (21) International Application Number:
PCT/US2016/018575
- (22) International Filing Date:
19 February 2016 (19.02.2016)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
14/662,603 19 March 2015 (19.03.2015) US
- (71) Applicant: INTEL CORPORATION [US/US]; 2200
Mission College Boulevard, Santa Clara, California 95054
(US).

- (72) Inventors: APODACA, Michael; 1563 Ballou Circle, Folsom, California 95630 (US). PIAZZA, Thomas; 9 White Pine Lane, Setauket, New York 11733 (US). PODDAR, Bimal; 4710 Freno Way, EL Dorado Hills, California 95762 (US).
- (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.

[Continued on next page]

(54) Title: DYNAMICALLY MANAGING MEMORY FOOTPRINT FOR TILE BASED RENDERING




FIG. 2

(57) Abstract: The introduction of an "out-of-memory" marker in the sorted tile geometry sequence for a tile may aid in handling out-of-memory frames. This marker allows hardware to continue rendering using the original data stream instead of the sorted data stream. This enables use of the original data stream allows the system to continue rendering without requiring any driver intervention. During the visibility generation/sorting phase, the number of memory pages required for storing the data for a rendering pass is continuously tracked. This tracking includes tracking the pages that are required even if the hardware had not run out-of-memory. This information can be monitored by a graphics driver and the driver can provide more memory pages for the system to work at full efficiency.



(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))*

Published:

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

DYNAMICALLY MANAGING MEMORY FOOTPRINT FOR TILE BASED RENDERING

Background

[0001] Tile based rendering is a hardware rendering technique where the three-dimensional (3D) geometry for a “rendering pass” is first sorted into the “bins” which correspond to “tiles” in the frame buffer and then each tile is rendered. Tile based rendering hardware derives its efficiency by rendering the full geometry of a render-pass for each tile and keeping the rendering buffers, associated with tiles, in hardware cache. An end-user visible application frame consists of a sequence of “rendering passes”.

[0002] Since the tile based renderers derive efficiency from the rendering the full render-pass, they need to buffer up the geometry information for the full render-pass. If the scene is very complex, the allocated buffer storage may not be enough. A graphics driver can try to compensate for this by allocating larger buffers, but then this triggers another problem. In order to maintain the responsiveness of the overall system, it is desirable to minimize the memory footprint of all the running applications. It does not make sense to allocate the same amount of memory for geometry buffers for rendering a widget as it does for rendering a complex 3D scene.

[0003] Due to the above trade off, there is no good solution. If too little memory is allocated, the graphics hardware can run out-of-memory and then would have to signal to driver components to add more memory which can result in an end user visible stall. If too much memory is allocated per application, then overall system responsiveness is adversely affected.

Brief Description Of The Drawings

[0004] Some embodiments are described with respect to the following figures:

Figure 1 is a flow chart for one embodiment;

Figure 2 is a flow chart for another embodiment;

Figure 3 is a block diagram of a processing system according to one embodiment;

Figure 4 is a block diagram of a processor according to one embodiment;

Figure 5 is a block diagram of a graphics processor according to one embodiment;

Figure 6 is a block diagram of a graphics processing engine according to one embodiment;

Figure 7 is a block diagram of another embodiment of a graphics processor;

Figure 8 is a depiction thread execution logic according to one embodiment;

Figure 9 is a block diagram of a graphics processor instruction format according to some embodiments;

Figure 10 is a block diagram of another embodiment of a graphics processor;

Figure 11A is a block diagram of a graphics processor command format according to some embodiments;

Figure 11B is a block diagram illustrating a graphics processor command sequence according to some embodiments;

Figure 12 is a depiction of an exemplary graphics software architecture according to some embodiments;

Figure 13 is a block diagram illustrating an IP core development system according to some embodiments; and

Figure 14 is a block diagram showing an exemplary system on chip integrated circuit according to some embodiments.

Detailed Description

[0005] The two conflicting conditions can be balanced by rendering the frames with “out-of-memory” somewhat more slowly without any stalls and providing an opportunity for software to recover and increase the memory footprint when increased application footprint is detected. Additionally, information is provided on how much additional memory would have been required.

[0006] “Visibility recording” driven “tile based rendering” includes two passes. In the first pass, the incoming application stream is processed and “visibility” information for each primitive (e.g. triangle”) is generated per-tile. In the second pass, the incoming application stream is re-processed for each tile together with the recorded “visibility” information. “Visibility” information records which triangles have survived back-face culling together with information about whether a primitive would be visible in a tile including visibility due to rasterization rules.

[0007] Traditional tile based renderers read input object position and output window positions during the first phase and then re-read the window position. Input object position typically indicates coordinates of object vertices in some local coordinate system where each object may have its own origin. Output window position indicates the positions of objects relative to the screen space and all objects are now described with respect to a common origin. In contrast to traditional tile based rendering architectures, the visibility based tile rendering algorithm reads the input object position, writes out a visibility bit per-primitive per tile during the first phase to indicate whether a particular primitive is visible within a particular tile and reads the visibility during the rendering phase. If the primitive is visible, it re-reads the input object position during the rendering phase. This leads to reading the position twice but writing only 1-bit of visibility per-primitive instead of a lot of intermediate information.

[0008] Visibility information is stored in memory pages. The number of pages required is a function of the application complexity. As was mentioned previously, a graphics driver tries to minimize the footprint of this allocation.

[0009] A high level flow of tiled rendering, using visibility recording, is shown in **Figure 1**. The sequence shown in Figure 1 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 optical, magnetic or semiconductor storages.

[0010] The sequence begins when an application submits a frame for rendering (stage 10). The frame commands in the application are recorded as indicated in stage 12. Then the flow splits.

[0011] At step 14, for each tile, primitives that are not visible are culled and whether a primitive is visible in the tile is recorded. Then the tile data is sent to the renderer indicated at 16. The tile is rendered using the recorded visibility information. A check at diamond 18 determines whether all tiles have been rendered. If not, a count is incremented at 20 and the remaining tiles are rendered at 16. When all the tiles have been rendered, the recorded data buffers are reset at block 22.

[0012] While the system described above works robustly when enough memory is available, one of the key considerations is how to handle the condition when no more memory is available for recording the visibility information. "Introduction" of an "out-of-memory" token in the sorted tile visibility information helps with handling out-of-memory frames without requiring additional memory. The visibility generation stage, since it does not have any more memory, records an "out-of-memory" marker as a terminating marker and all subsequent visibility information for that tile is discarded. As a result the visibility generation stage continues to function without requiring any driver intervention.

[0013] While the "out-of-memory" token handled the processing of the input stream, since the visibility information is not complete, the rendering systems need to produce the correct image. When the tile is rendered, the "out-of-memory" token signals to the rendering sub-system to continue rendering the tile using the original data stream without the help of the "visibility data". Any geometry following the "out-of-memory" token for a tile is considered visible in that tile. It may be slower since more geometry, even things not intersecting the tile, would be processed per-tile instead of being discarded right away, but is ultimately discarded since it does not intersect the tile. Overall, the rendering results are correct and comply with the application program interface (API) requirements since each tile is rendered one at a time.

[0014] During the visibility generation/sorting phase, the number of memory pages required for storing the data for a rendering pass is continuously tracked. This tracking includes tracking the amount of required memory (e.g. pages) even if the hardware had not run out-of-memory. This information can be monitored by a graphics driver (as will be explained later) and the driver can provide more memory pages for the system to work at full efficiency.

[0015] The graphics driver starts off with a conservative number of pages for recording the visibility information. But this “conservative” or even a “generous” value may be too small when the application frame is very complex. In such cases, the visibility information recording is terminated with an “out-of-memory” marker or token.

[0016] When a tile is rendered, it starts out having visibility for the incoming geometry and that part is processed regularly. However, if during the rendering phase for a particular tile, the “out-of-memory” marker is encountered, all the remaining geometry is processed as if everything is visible to the tile i.e. without using culling information. This requires more work since culled primitives are also processed but it produces the right results since the rendering to a tile is constrained by a scissor region which precludes rendering outside the tile.

[0017] In other words, if a primitive is allowed, because of the out-of-memory condition, to be processed initially even though it is not visible, the primitive will not be drawn because it is outside the pertinent tile’s scissor region, so a low level scissor operation will ultimately exclude the primitive that passed because of the out-of-memory condition. It should be noted that such a scissoring operation is required anyways for processing a primitive that is partly inside the tile and partly outside the tile.

[0018] If the driver’s conservative guess for memory allocation is based on some measure of incoming vertex count, a good chunk of the render-pass is drawn at high efficiency using the visibility information. Only the last sub-section of the render-pass, where the “out-of-memory” marker was generated, may experience performance degradation.

[0019] It is not necessary that all tiles have the “out-of-memory” condition. For example, if a tile contains a very complex object with lots of primitives, then that tile may be the only tile with the out-of-memory condition and all the remaining tiles may all render normally. The out-of-memory condition is processed automatically in hardware without blocking the rendering sub-system by generating an interrupt and requiring an interrupt to be serviced before resumption. Additionally, there is no guarantee that a graphics driver is always able to allocate additional memory. But even when additional memory cannot be allocated, the out-of-memory scheme described here will continue to function and allow the scene to be rendered.

[0020] While rendering with some performance loss may be fine for a frame or two, continuing to render in an “out-of-memory” mode may result in overall reduced performance. In order to address the “recovery” from the out-of-memory condition, a counter in the hardware continuously records the number of memory pages allocated for the “visibility information”. This counter not only records the number of pages used, it also records “what would be required” when out-of-memory condition is reached.

[0021] Since the system has reach out-of-memory condition, no new pages can be allocated but the count is maintained as if pages were allocated. In the out-of-memory condition, the visibility information is generated as it normally would be, and page counts are incremented. However the visibility information for a tile is discarded because of the out-of-memory condition. Processing of the graphics primitive commands is continued without requiring an agent, such as a software driver, to resolve the out-of-memory condition before processing the remaining commands, for example by allocating more memory pages to store the tile’s visibility information.

[0022] An application frame may include several render-passes. One render-pass may be for generation of shadow information and another may be for generation of deferred lighting parameters and so on. The tiled rendering system is not only consuming pages for recording but is also freeing up pages as rendering is completed for each render-pass.

[0023] The out-of-memory marker described previously is per-render-pass in one embodiment. If render-pass 3 generated an out-of-memory marker, this does not mean render-pass 4 will generate out-of-memory marker. Each render-pass has its own memory requirements, since the processing of render-pass 3 will free up memory pages that can be used for render-pass 4. Given that each render-pass has its own memory requirements, a graphics driver can, for example, do the following:

Submit request to copy hardware page counter register to memory location 1

Render-pass1

Submit request to copy hardware page counter register to memory location 2

Render-pass2

Submit request to copy hardware page counter register to memory location 3

Render-pass3

[0024] As the counter values are stored in memory, a graphics driver can continuously monitor the values and take actions. Consider as an example below the “counter values” reported for different render-passes.

	Counter Value	Pages required for Render-pass	Pages required for 2 consecutive Render-passes
....			
End of RenderPass 2	200	...	
End of Render-pass 3	250	50	
End of Render-pass 4	350	100	150
End of Render-pass 5	475	125	225

[0025] The driver can subtract two adjacent values to determine how many pages are required for each render-pass. As an example render-pass 5 requires 125 memory pages and if only 75 pages are provided, the graphics driver can add more pages for recording visibility. The driver also has the flexibility to compute the number of pages required for two consecutive “render-passes” ahead when

generating visibility information. For example, since Render-pass 4 and 5 require the maximum number of pages, if the driver wants to allow for visibility information to be two render-passes ahead, it would need to provide 225 memory pages.

[0026] Overall, with the out-of-memory marker and a counter that continuously counts the number of pages required for visibility recording together with the proposed usage, the solution reduces or eliminates the memory allocation dilemma/stalls prevalent in the existing solutions in some embodiments.

[0027] **Figure 2** shows a sequence for managing footprint during tile-based rendering according to one embodiment. The sequence may be implemented in software, firmware and/or hardware. In software and firmware embodiments it may be implemented using computer executed instructions stored in one or more non-transitory computer readable media such as magnetic, optical or semiconductor storages. In one embodiment it may be implemented by a graphics processing unit.

[0028] The sequence begins by processing visibility information for tiles as indicated in block 30. In a first pass, visibility of different primitives with any given tile may be analyzed for recording and used in a subsequent rendering pass.

[0029] A check at diamond 32 determines whether there is sufficient memory available to store the visibility information for a particular tile. If not, the tile is marked as indicated in block 36. Then the tile is processed as if the next set of primitives is entirely visible (block 38). This is done in some embodiments without recording any visibility information. This may mean that extra processing is done but the tile will be correctly processed and processed without being delayed waiting for additional memory.

[0030] However, if the check at diamond 32 determines that sufficient memory is available then the visibility information is recorded for use in a subsequent pass as indicated in block 34. As indicated in block 40, the amount of available memory is assessed and additional memory may be allocated asynchronously. For example the amount of remaining available memory may be determined together with an estimate of the amount of memory that may be needed. In this way additional memory can be allocated to deal with the out-of-memory condition as it exists in any

particular rendering pass. As a result, the out-of-memory condition is reduced in duration and improving performance in some embodiments.

[0031] While the method described here refers to tiled rendering, this method can be extended to other situations. As an example, if there is a single tile in the system, it as if visibility information is generated for the full render-pass and then rendering is performed using the visibility information. Since it is rendering to the full framebuffer, while it may not be called tiled rendering, the out-of-memory and asynchronous method of providing more memory can also be used for such conditions.

[0032] **Figure 3** 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 on 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.

[0033] 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.

[0034] 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).

[0035] 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.

[0036] 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.

[0037] 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.

[0038] 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 discrete external graphics processor, such as the external graphics processor 112.

[0039] **Figure 4** 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 4** 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.

[0040] 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.

[0041] 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).

[0042] 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.

[0043] 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.

[0044] 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.

[0045] 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.

[0046] 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.

[0047] **Figure 5** 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.

[0048] 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.

[0049] 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.

[0050] 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.

[0051] 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.

[0052] 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.

[0053] **Figure 6** 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 5**. Elements of **Figure 6** 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.

[0054] 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.

[0055] 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.

[0056] 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).

[0057] 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.

[0058] 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.

[0059] 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.

[0060] **Figure 7** is a block diagram of another embodiment of a graphics processor 500. Elements of **Figure 7** 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.

[0061] 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.

[0062] 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.

[0063] 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.

[0064] **Figure 8** illustrates thread execution logic 600 including an array of processing elements employed in some embodiments of a GPE. Elements of **Figure 8** 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.

[0065] 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.

[0066] 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).

[0067] 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.

[0068] 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.

[0069] 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.

[0070] 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 7**) dispatches vertex processing, tessellation, or geometry processing threads to thread execution logic 600 (**Figure 8**). In some embodiments, thread dispatcher 604 can also process runtime thread spawning requests from the executing shader programs.

[0071] 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.

[0072] 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.

[0073] **Figure 9** 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.

[0074] 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.

[0075] 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.

[0076] 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.

[0077] 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.

[0078] 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.

[0079] 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.

[0080] 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.

[0081] **Figure 10** is a block diagram of another embodiment of a graphics processor 800. Elements of **Figure 10** 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.

[0082] 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.

[0083] 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.

[0084] 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.

[0085] 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.

[0086] 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.

[0087] 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.

[0088] 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.

[0089] 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.

[0090] 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.

[0091] 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.

[0092] 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.

[0093] **Figure 11A** is a block diagram illustrating a graphics processor command format 900 according to some embodiments. **Figure 11B** is a block diagram illustrating a graphics processor command sequence 910 according to an embodiment. The solid lined boxes in **Figure 11A** 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 11A** 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.

[0094] 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.

[0095] The flow diagram in **Figure 11B** 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.

[0096] 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.

[0097] 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 912 is required immediately before a pipeline switch via the pipeline select command 913.

[0098] 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.

[0099] 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.

[0100] 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.

[0101] 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.

[0102] 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.

[0103] 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.

[0104] 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.

[0105] 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.

[0106] 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.

[0107] **Figure 12** 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.

[0108] 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.

[0109] 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.

[0110] 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.

[0111] 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.

[0112] **Figure 13** 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.

[0113] 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.

[0114] **Figure 14** 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.

[0115] 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.

[0116] The following clauses and/or examples pertain to further embodiments:

One example embodiment may be a method comprising during tile rendering, determining in a hardware processor, whether sufficient storage is available in a memory that stores visibility information to store additional visibility information and if an out-of-memory condition exists because sufficient storage is not available, continuing to render without storing visibility information during the out-of-memory condition before requiring an agent to resolve the out-of-memory condition. The method may also include resolving an out-of-memory condition, arising in one rendering pass, in a subsequent rendering pass. The method may also include recording, for each tile, memory that would have been needed absent the out-of-memory condition. The method may also include performing tile based rendering, while rendering, detecting an out-of-memory condition in a first rendering pass and marking said condition with a marker, and detecting the marker in another rendering

pass. The method may also include reducing the rendering speed in response to detecting the marker. The method may also include monitoring, during visibility generation, an amount of memory needed to store visibility data for a rendering pass and allocating more memory in response to detecting a need for more memory. The method may also include determining an amount of available memory for storing visibility information. The method may also include proceeding to render without culling primitives for a rendering pass with an out-of-memory condition. The method may also include in response to detection of an out-of-memory condition, rendering as if all primitives in a tile are visible. The method may also include rendering without storing visibility information only during a rendering pass in which an out-of-memory condition was detected, while continuing to render while storing visibility information in another rendering pass.

[0117] Another example embodiment may be one or more non-transitory computer readable media storing instructions executed by a processor to perform a sequence comprising during tile rendering, determining whether sufficient storage is available in a memory that stores visibility information to store additional visibility information, and if an out-of-memory condition exists because sufficient storage is not available, continuing to render without storing visibility information during the out-of-memory condition before requiring an agent to resolve the out-of-memory condition. The media may include said sequence including resolving an out-of-memory condition, arising in one rendering pass, in a subsequent rendering pass. The media may include said sequence including recording, for each tile, memory that would have been needed absent the out-of-memory condition. The media may include performing tile based rendering, while rendering, detecting an out-of-memory condition in a first rendering pass and marking said condition with a marker, and detecting the marker in another rendering pass. The media may include said sequence including reducing the rendering speed in response to detecting the marker. The media may include said sequence further including monitoring, during visibility generation, an amount of memory needed to store visibility data for a rendering pass and allocating more memory in response to detecting a need for more memory. The media may include said sequence including determining an amount of available memory for storing visibility information. The media may include

said sequence including proceeding to render without culling primitives for a rendering pass with an out-of-memory condition. The media may include said sequence including in response to detection of an out-of-memory condition, rendering as if all primitives in a tile are visible. The media may include said sequence including proceeding with rendering without storing visibility information only during a rendering pass in which an out-of-memory condition was detected, while continuing to render while storing visibility information in another rendering pass.

[0118] In another example embodiment it may be an apparatus comprising a hardware processor to determine, during tile rendering, whether sufficient storage is available in a memory that stores visibility information to store additional visibility information, and if an out-of-memory condition exists because sufficient storage is not available, continue to render without storing visibility information during the out-of-memory condition before requiring an agent to resolve the out-of-memory condition, and a storage coupled to said processor. The apparatus may include said processor to resolve an out-of-memory condition, arising in one rendering pass, in a subsequent rendering pass. The apparatus may include said processor to record, for each tile, memory that would have been needed absent the out-of-memory condition. The apparatus may include said processor to perform tile based rendering, while rendering, detect an out-of-memory condition in a first rendering pass and mark said condition with a marker and detect the marker in another rendering pass. The apparatus may include said processor to reduce the rendering speed in response to detecting the marker. The apparatus may include said processor to monitor, during visibility generation, an amount of memory needed to store visibility data for a rendering pass and allocate more memory in response to detecting a need for more memory. The apparatus may include said processor to determine an amount of available memory for storing visibility information. The apparatus may include said processor to proceed to render without culling primitives for a rendering pass with an out-of-memory condition. The apparatus may include said processor in response to detection of an out-of-memory condition, render as if all primitives in a tile are visible. The apparatus may include said processor to proceed with rendering without storing visibility information only during a rendering

pass in which an out-of-memory condition was detected, while continuing to render while storing visibility information in another rendering pass.

[0119] 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.

[0120] 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.

[0121] 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 during tile rendering, determining in a hardware processor, whether sufficient
3 storage is available in a memory that stores visibility information to store additional
4 visibility information; and
5 if an out of memory condition exists because sufficient storage is not
6 available, continuing to render without storing visibility information during the out of
7 memory condition before requiring an agent to resolve the out of memory condition.
- 1 2. The method of claim 1 including resolving an out of memory condition, arising
2 in one rendering pass, in a subsequent rendering pass.
- 1 3. The method of claim 1 including recording, for each tile, memory that would
2 have been needed absent the out of memory condition.
- 1 4. The method of claim 3 including:
2 performing tile based rendering;
3 while rendering, detecting an out-of-memory condition in a first rendering pass
4 and marking said condition with a marker; and
5 detecting the marker in another rendering pass.
- 1 5. The method of claim 4 including reducing the rendering speed in response to
2 detecting the marker.
- 1 6. The method of claim 1 further including monitoring, during visibility generation,
2 an amount of memory needed to store visibility data for a rendering pass and
3 allocating more memory in response to detecting a need for more memory.
- 1 7. The method of claim 1 including determining an amount of available memory
2 for storing visibility information.

- 1 8. The method of claim 1 including proceeding to render without culling
2 primitives for a rendering pass with an out-of-memory condition.
- 1 9. The method of claim 1 including in response to detection of an out-of-memory
2 condition, rendering as if all primitives in a tile are visible.
- 1 10. The method of claim 1 including proceeding with rendering without storing
2 visibility information only during a rendering pass in which an out-of-memory
3 condition was detected, while continuing to render while storing visibility information
4 in another rendering pass.
- 1 11. One or more non-transitory computer readable media storing instructions
2 executed by a processor to perform a sequence comprising:
3 during tile rendering, determining whether sufficient storage is available in a
4 memory that stores visibility information to store additional visibility information; and
5 if sufficient storage is not available, continuing to render without storing
6 visibility information during the out of memory condition before requiring an agent to
7 resolve the out of memory condition.
1
- 1 12. The media of claim 11, said sequence including resolving an out of memory
2 condition, arising in one rendering pass, in a subsequent rendering pass.
1
- 1 13. The media of claim 11, said sequence including recording, for each tile,
2 memory that would have been needed absent the out of memory condition.
- 1 14. The media of claim 13, said sequence including:
2 performing tile based rendering;
3 while rendering, detecting an out-of-memory condition in a first rendering pass
4 and marking said condition with a marker; and
5 detecting the marker in another rendering pass.

1 15. The media of claim 14, said sequence including reducing the rendering speed
2 in response to detecting the marker.

1 16. The media of claim 11, said sequence further including monitoring, during
2 visibility generation, an amount of memory needed to store visibility data for a
3 rendering pass and allocating more memory in response to detecting a need for
4 more memory.

1 17. The media of claim 11, said sequence including determining an amount of
2 available memory for storing visibility information.

1 18. The media of claim 11, said sequence including proceeding to render without
2 culling primitives for a rendering pass with an out-of-memory condition.

1 19. The media of claim 11, said sequence including in response to detection of an
2 out-of-memory condition, rendering as if all primitives in a tile are visible.

1 20. The media of claim 11, said sequence including proceeding with rendering
2 without storing visibility information only during a rendering pass in which an out-of-
3 memory condition was detected, while continuing to render while storing visibility
4 information in another rendering pass.

1 21. An apparatus comprising:
2 a hardware processor to determine, during tile rendering, whether sufficient
3 storage is available in a memory that stores visibility information to store additional
4 visibility information, and if sufficient storage is not available, continue to render
5 without storing visibility information during the out of memory condition before
6 requiring an agent to resolve the out of memory condition; and
7 a storage coupled to said processor.

- 1 22. The apparatus of claim 21, said processor to resolve an out of memory
2 condition, arising in one rendering pass, in a subsequent rendering pass.
- 1 23. The apparatus of claim 21, said processor to record, for each tile, memory
2 that would have been needed absent the out of memory condition.
- 1 24. The apparatus of claim 23, said processor to perform tile based rendering,
2 while rendering, detect an out-of-memory condition in a first rendering pass and
3 mark said condition with a marker and detect the marker in another rendering pass.
- 1 25. The apparatus of claim 24, said processor to reduce the rendering speed in
2 response to detecting the marker.




FIG. 1

2 / 14




FIG. 2

3 / 14




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8

GRAPHICS CORE INSTRUCTION
FORMATS

700



64-BIT COMPACT



OPCODE DECODE

740




FIG. 9

10 / 14




FIG. 10

11 / 14


GRAPHICS PROCESSOR COMMAND FORMAT
900

FIG. 11A


GRAPHICS PROCESSOR COMMAND SEQUENCE
910

FIG. 11B




FIG. 12




FIG. 13




FIG. 14

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2016/018575**A. CLASSIFICATION OF SUBJECT MATTER****G06T 15/00(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)

G06T 15/00; G06F 12/02; G06T 17/00; G06T 17/20; G06T 1/00; 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: tile-based rendering, visibility information, memory, marker

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2011-0227936 A1 (JONATHAN REDSHAW) 22 September 2011 See paragraphs [0011], [0030]–[0063]; and figure 1.	1–25
Y	US 2007-0229521 A1 (WEI LI et al.) 04 October 2007 See paragraphs [0017], [0031], [0060]–[0062]; and figures 1–3.	1–25
A	US 2012-0206455 A1 (DAVID ROBERT SHREINER) 16 August 2012 See paragraphs [0054]–[0062]; and figures 5–7.	1–25
A	US 2014-0267259 A1 (QUALCOMM INCORPORATED) 18 September 2014 See paragraphs [0252]–[0286]; and figures 12–16.	1–25
A	US 2013-0241938 A1 (ANDREW GRUBER et al.) 19 September 2013 See paragraphs [0048]–[0053]; 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

21 June 2016 (21.06.2016)

Date of mailing of the international search report

21 June 2016 (21.06.2016)

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

LEE, Dong Yun

Telephone No. +82-42-481-8734



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2016/018575

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2011-0227936 A1	22/09/2011	CN 102804225 A EP 2548171 A1 EP 2548171 B1 GB 2478660 A GB 2478660 B US 2015-0109314 A1 US 8836710 B2 WO 2011-114114 A1	28/11/2012 23/01/2013 08/10/2014 14/09/2011 13/06/2012 23/04/2015 16/09/2014 22/09/2011
US 2007-0229521 A1	04/10/2007	DE 102007016179 A1 US 7911474 B2	25/10/2007 22/03/2011
US 2012-0206455 A1	16/08/2012	CN 102708579 A GB 2488196 A JP 2012-168951 A US 8339409 B2	03/10/2012 22/08/2012 06/09/2012 25/12/2012
US 2014-0267259 A1	18/09/2014	None	
US 2013-0241938 A1	19/09/2013	CN 104169974 A EP 2826024 A1 JP 2015-514246 A KR 10-2014-0139553 A WO 2013-138061 A1	26/11/2014 21/01/2015 18/05/2015 05/12/2014 19/09/2013