



(51) International Patent Classification:

*G06T 1/20* (2006.01) *G06T 7/40* (2006.01)  
*G06T 3/40* (2006.01)

(21) International Application Number:

PCT/US2016/036184

(22) International Filing Date:

7 June 2016 (07.06.2016)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

14/751,513 26 June 2015 (26.06.2015) US

(71) Applicant: **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95054 (US).

(72) Inventor: **SEILER, Larry**; 198 Linden Street, Boylston, Massachusetts 01505 (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.

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

Published:

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

(54) Title: FILTERING MULTI-SAMPLE SURFACES




FIG. 5

(57) Abstract: In accordance with some embodiments, multi-sampling may be used together with texture filtering and particularly texture filtering that generally uses rectangular grids of samples. This is accomplished by performing the texture filtering before doing the resolve, while conventionally the resolve is done and then the texture filtering is done. In addition, each sample is filtered as if it were the only sample.

## FILTERING MULTI-SAMPLE SURFACES

### Background

[0001] This relates generally to graphics processing and particularly to a shading technique call multi-sampling.

[0002] Multi-sampled anti-aliasing (multi-sampling or MSAA) is a technique that allows higher image quality by storing multiple color and depth sample values at each pixel. This permits higher precision for object edges, since each sample position is separately tested for whether it is inside or outside the primitive. The multi-sample image is “resolved” to a single sample per pixel when rendering is finished. A typical resolve function simply averages the sample values within each pixel.

[0003] Typically, when using MSAA, a single color is computed for all of the samples in a pixel that are within the primitive. As a result, interior pixels of a primitive have the same value at each sample. Samples have different values only in pixels that contain a triangle edge or vertex. This makes the rendering time required for MSAA rendering the same as for single-sampling for interior pixels.

[0004] Super-sampled anti-aliasing (super-sampling or SSAA) differs from MSAA in that it uses a higher precision pixel grid for rendering, which is then resolved to the desired pixel resolution. An MSAA and an SSAA image may have the same number of samples, but the SSAA image requires dramatically more processing, since a separate color must be computed for each sample.

[0005] Another difference between MSAA and SSAA is that SSAA typically uses an ordered grid of samples. That is, just as the pixels are defined at positions on a rectangular grid, the samples are also defined at positions on a rectangular grid. MSAA typically uses scattered sample positions within each pixel. E.g., 4x SSAA (4 samples per pixel) uses an aligned 2x2 grid of samples within each pixel. 4x MSAA typically uses four positions that cover four different X positions and four different Y positions. This results in higher quality edge anti-aliasing.

[0006] Texture filtering is a technique that allows scaling, rotating, and otherwise warping an image by applying a filtering function to the color values stored at multiple texels in the array. Typical filtering functions multiply each of a set of texel values by a filter weight and then sum the results. The set of texture values that are filtered are called the “footprint” of the filter function.

[0007] One type of texture filtering, called bilinear filtering or bilinear interpolation, generally requires a rectangular or square grid of sample positions. This technique linearly interpolates between two horizontal pairs of values at top and bottom of a rectangle and then linearly interpolates between those two values. While it is conceivable that a non-rectangular grid could be used, the calculation complexity would be sufficiently large to deter most practical applications.

[0008] When texture filtering is applied to multi-sampled images, a problem arises because multi-samples in conventional multi-sample images are not arranged within the pixel in a rectangular way. Another problem that arises with respect to conventional multi-sampling is that, in some situations with particular depicted objects, block artifacts will be created. These block artifacts are the result of samples within a pixel being rendered with the same color value when they are in the interior of a triangle or other rendered primitive. Typically multi-sample values only differ in a pixel when they are part of different primitives.

[0009] The standard way to texture filter a multi-sample image is to first convert it to a single-sample image. This is referred to as a “resolve” operation. A typical resolve operation simply averages the sample values within a pixel. Expressed as a filter function, a 4-sample resolve multiplies each sample value by 0.25 and then sums them. This blends color values across primitive edges and simply returns the common sample value for pixels entirely in the interior of a triangle or other primitive.

#### Brief Description Of The Drawings

[0010] Some embodiments are described with respect to the following figures:

Figure 1 is a depiction of a conventional multi-sampled pixel;

Figure 2 is a depiction of a linear function that is computed at locations in a 2x2 array of pixels;

Figure 3 is a depiction of conventional per-sample multi-sample, anti-alias (MSAA) shading;

Figure 4 is a depiction of conventional per-pixel multi-sampling, anti-aliasing (MSAA) shading;

Figure 5 is a flow chart for one embodiment;

Figure 6 is a schematic diagram for one embodiment;

Figure 7 is a block diagram of a processing system according to one embodiment;

Figure 8 is a block diagram of a processor according to one embodiment;

Figure 9 is a block diagram of a graphics processor according to one embodiment;

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

Figure 11 is a block diagram of another embodiment of a graphics processor;

Figure 12 is a depiction thread execution logic according to one embodiment;

Figure 13 is a block diagram of a graphics processor instruction format according to some embodiments;

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

Figure 15A is a block diagram of a graphics processor command format according to some embodiments;

Figure 15B is a block diagram illustrating a graphics processor command sequence according to some embodiments;

Figure 16 is a depiction of an exemplary graphics software architecture according to some embodiments;

Figure 17 is a block diagram illustrating an IP core development system according to some embodiments; and

Figure 18 is a block diagram showing an exemplary system on chip integrated circuit according to some embodiments.

#### Detailed Description

[0011] In accordance with some embodiments, multi-sampling may be used together with texture filtering and particularly texture filtering that generally uses rectangular grids of samples. This is accomplished by performing the texture filtering before doing the resolve, while conventionally the resolve is done and then the texture filtering is done. In addition, each sample is filtered separately, as if a series of single-sample images were specified that each only contained the value of that sample at each pixel.

[0012] Thus, any technique for texture filtering that is suitable for images with single samples can be used in connection with multi-sampling. Also, in many embodiments, it is possible to use these techniques with per-pixel shading and per-sample shading, as well as with arrangements that switch between per-pixel and per-sample shading. In some embodiments, this may reduce the blurring produced by resolving before filtering. Embodiments also eliminate the blockiness that can result if multi-sample images are filtered as if they were single-sample images and the filter footprint covers multiple samples of the same pixel.

[0013] Figure 1 shows a pixel having four samples which are not arranged in a rectangular grid. There are many reasons for this arrangement of samples but it creates a problem in that now, for texture filtering such as texture filtering that uses bilinear filtering or interpolation, the desirable rectangular grid is absent.

[0014] Figure 2 shows a linear function with values computed at locations within a 2x2 quad of four pixels. Values are shown at a 4x4 grid of positions within each pixel and also at the center of each pixel. In this linear function, the samples increase by four going vertically from pixel to pixel and by eight going horizontally from pixel to pixel. This function is used to demonstrate that embodiments of the described multi-sample filtering method correctly reconstruct a linear function, and thus produce the same results as existing single-sample filter methods when applied to a linear

function. The values 8.5, 16.5, 4.5, and 12.5 at the center of the four pixels are the values of the function at the pixel centers.

[0015] Figure 3 shows an arrangement of a 2x2 quad of four pixels which may be identified such that the upper left pixel has the number 9 at the upper left sample position, the upper right pixel has the sample 17 in its upper left position, the pixel on the lower left has a sample 5 in its upper left position and the pixel in the lower right has the sample 13 in its upper right position. The numbers provided in the samples simply show the application of the linear function in Figure 2 to the sample positions. Thus in per-sample shading, each sample has its own value.

[0016] In contrast, as shown in Figure 4, per-pixel multi-sampling anti-aliasing shading generally gives each of the four sample positions the same value when all four are covered by the primitive. In one case, this same value may be determined by performing the filter function at the center position of each pixel. Thus for the upper left pixel shown in Figure 4, each of the sample positions have the value 8.5 which is the result of computing the linear function in Figure 3 at the center of each pixel.

[0017] In accordance with some embodiments, filtering may be performed before resolving to a single value per pixel by treating the multi-sample image as a set of single-sample images, one for each sample in the pixel in the multi-sample image. This is possible since sample positions are typically the same in each pixel. Therefore, an image created from choosing sample N in each pixel forms a rectangular grid, allowing all existing single-sample filtering methods to be used. However, other embodiments can be conceived where the sample positions in adjacent pixels are jittered by a small amount, e.g. by 1/8 of a pixel in X and/or Y. In such a case, the grid formed by an image formed by a single sample is not rectangular, but is close enough to rectangular that it can be treated as such for filtering.

[0018] Single-sample filter functions typically assume that the sample position is at the center of each pixel. To apply these functions to per-sample filtering, it is necessary to offset the filter position. For example, if a sample is at location (1/8,

3/8) relative to the center of each pixel, where each pixel has a size of 1 unit, then that offset must be subtracted from the filter position to compensate. Embodiments may perform this either by biasing the filter values before being sent for hardware processing or it may occur in fixed function hardware that implements filtering using this invention.

[0019] Typical single-sample filtering defines a footprint around the filter position that represents the texels to combine in the filter. This occurs in the usual way when filtering a single sample position according to this invention. For example, bilinear filtering selects the four pixels nearest to the filter position. Embodiments of this method would select four pixels nearest to the offset filter position. Since each sample uses its own offset, the four nearest pixels can be different for different samples.

[0020] The final step is to average together the results of filtering each individual sample position. A typical number of samples per pixel is four. In that case, this method would filter each of the four samples individually and then average together the four results to produce the filter result.

[0021] Figure 5 shows a multi-sampled sequence that may be implemented in software, firmware and/or hardware. In software and firmware embodiments it may be implemented by computer executed instructions stored in one or more non-transitory computer readable media such as magnetic, optical or semiconductor storage. In some embodiments, such a sequence may be implemented by a graphics processor.

[0022] The multi-sample sequence 10, shown in Figure 5, starts by setting an accumulator to zero, as shown in block 10. Once the sample has been obtained (block 12), the sample position relative to the pixel center is subtracted from the texture filter position, as shown in block 14. Then, single sample texture filtering is performed at that position, using the specified sample at each pixel as indicated in block 16. For example, bilinear filtering finds the pixel centers nearest to the filter position and linearly interpolates them. Then the result from the filter is added to the accumulator, as shown in block 18. Of course it should be understood that any

texture filtering technique that is applicable to a single sample texture can be used in connection with the techniques described herein. A check at diamond 20 indicates if this is the last sample. If not, the flow iterates until all the samples have been processed.

[0023] After the last sample has been completed, the resulting texture filter values are resolved for the  $N$  samples, as indicated in block 22. A typical way to resolve the results is to sum the individual filter results and divide by the number of samples. In the example described above  $N = 4$ . The accumulator result may be divided by  $N$  in one embodiment. Other resolve techniques may also be used.

[0024] Thus, by doing the filtering then the resolve, the reverse of the conventional order, a variety of advantages may be obtained in some embodiments. The blockiness conventionally occurring with some images that are multi-sampled can be reduced or even substantially eliminated. In addition, it is possible to filter whether the rendering process computes a different value per sample or computes a single value for all of the samples in a pixel that are within the primitive. As a result it is possible with these techniques to switch between per-sample and per-pixel multi-sample, anti-aliasing.

[0025] It can be shown that when the pixels being filtered are fully covered by a primitive and the rendered per-sample or per-pixel values are a linear function, the same final result is produced for both per-pixel and per-sample rendering, and for both normal pre-resolve filtering and this post-resolve filtering method. As a result, in these cases, changing to this algorithm does not alter the result of filtering. For non-linear functions and for partially covered pixels, both per-pixel rendering and pre-resolve filtering introduce blurriness that may be reduced by this algorithm.

[0026] It is known that where multi-sampling has weaknesses, one can use super-sampling instead. But the problem is that super-sampling creates a lot more work which means slower operation and higher power consumption for the graphics processor. Therefore this method allows super-sampling (which renders a separate value per sample) and multi-sampling (which renders a single value per pixel covered by a primitive) to be mixed in the same image, the application choosing the

former where the latter is insufficient quality, e.g. because of a highly non-linear rendering function. This method allows both techniques to be combined in the same image without forcing the super-sampled pixels to be resolved prior to filtering.

[0027] This method can be implemented in software, firmware or hardware. Thus, referring to Figure 6, the render engine 30 includes a texture filter 32 coupled to a resolve engine, including an accumulator register 36 and a divide circuit 38. Typical numbers of samples per pixel are powers of two, so the divide circuit can typically be implemented with a bit shift. The texture filter receives the sample positions, offset by subtractor 30 and performs any filter function that is defined on single-sample pixels, applying it to the specified sample.

[0028] Then the result is added at 34 to the accumulator 36. This is repeated for each sample until all samples are processed. Then a resolve is completed at block 38. For example, the accumulator result may be divided by the  $N$ , the number of samples.

[0029] As used herein, single sampling refers to the normal rendering method where a single sample position is rendered at the center of each pixel. Super-sampling generally involves rendering each sample separately. As a result, within super-sampling, both rendering time and memory bandwidth are times larger. Typically super-samples are organized in a grid within a pixel. This allows a  $4x$  super-sampled image to be filtered as if it was a single sample image that is twice as large in each dimension.

[0030] In accordance with some embodiments, the resulting filter has a footprint that, when measured in pixels, is the same size as the filter footprint for single sampling. This prevents artifacts in the presence of fragment compression. Fragment compression stores the computed multi-sampled results in an index per-sample to indicate which rendered color covers that sample. If a pixel is fully covered, just one color is stored, along with an index per-sample selecting that color. This may dramatically reduce the required rendering time and memory bandwidth. For example  $8x$  multi-sampling may require  $2x$  of bandwidth of single sampling and one-quarter of the bandwidth of  $8x$  super-sampling. When measured in samples, this is a

wider filter than could be used with super-sampling. However in many cases, a wider filter is desirable because it may permit smoothing of the data.

[0031] The techniques described herein are applicable in some cases even if different pixels use different sample patterns. Some multi-sampling systems vary the sample pattern for adjacent pixels to increase the degree of anti-aliasing over a multi-pixel region. This is problematic since it can cause sawtooth artifacts along the edges. One way to limit such artifacts is to use the same pattern in each pixel, except offset by a small amount. For example for the four sample patterns, pixels in a  $2x$  region may use the same pattern offset by  $(-1/16, -1/16)$ ,  $(+1/16, -1/16)$ ,  $(+1/16, -1/16)$ , and  $(+1/16, +1/16)$ . This scheme results in sample plane where samples vary by  $+$  or  $-$   $1/16$ th of the pixel center. This should be a small enough variation to allow bilinear interpolation to be performed, assuming that all of the samples in the plane are at the same position.

[0032] **Figure 7** is a block diagram of a processing system 100, according to an embodiment. In various embodiments the system 100 includes one or more processors 102 and one or more graphics processors 108, and may be a single processor desktop system, a multiprocessor workstation system, or a server system having a large number of processors 102 or processor cores 107. In one embodiment, the system 100 is a processing platform incorporated within a system-on-a-chip (SoC) integrated circuit for use in mobile, handheld, or embedded devices.

[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 discreet external graphics processor, such as the external graphics processor 112.

[0039] **Figure 8** 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 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. 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 cache 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 9** 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 10** 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 9**. 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.

[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 11** is a block diagram of another embodiment of a graphics processor 500. Elements of **Figure 11** having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such.

[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 12** illustrates thread execution logic 600 including an array of processing elements employed in some embodiments of a GPE. Elements of **Figure 12** 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 11**) dispatches vertex processing, tessellation, or geometry processing threads to thread execution logic 600 (**Figure 12**). 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 13** 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 14** is a block diagram of another embodiment of a graphics processor 800. Elements of **Figure 14** having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such.

[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 15A** is a block diagram illustrating a graphics processor command format 900 according to some embodiments. **Figure 15B** is a block diagram illustrating a graphics processor command sequence 910 according to an embodiment. The solid lined boxes in **Figure 15A** 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 15A** 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 15B** 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.

[00100] 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.

[00101] 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.

[00102] 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.

[00103] 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.

[00104] 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.

[00105] 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.

[00106] 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.

[00107] **Figure 16** 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.

[00108] 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.

[00109] 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.

[00110] 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.

[00111] 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.

[00112] **Figure 17** 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 using a simulation model 1112. The simulation model 1112 may include functional, behavioral, and/or timing simulations. A register transfer level (RTL) design can then be created or

synthesized from the simulation model 1112. 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.

[00113] 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 3<sup>rd</sup> 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.

[00114] **Figure 18** 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 I<sup>2</sup>S/I<sup>2</sup>C 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.

[00115] 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.

[00116] The following clauses and/or examples pertain to further embodiments:

One example embodiment may be a method comprising treating a multi-sample image as a set of single-sample images, one single sample image for each sample in the multi-sample image, for a first of N samples, subtracting a sample position relative to a pixel center from a texture filter position, performing single sample texture filtering at the texture filter position to generate a result, combining the results into an accumulator, repeating the steps above for N-1 samples, and resolving the N texture filtering results. The method may also include wherein texture filtering includes bilinear filtering. The method may also include implementing per sample multi-sample anti-aliasing. The method may also include implementing per pixel multi-sampling anti-aliasing. The method may also include processing both per-pixel and per-sample multi-sample anti-aliasing. The method may also include receiving samples with a linear function and providing a texture filtering result that preserves said linear function. The method may also include receiving samples with a non-linear function and providing a texture filtering result that preserves said non-linear function. The method may also include wherein resolving includes summing the N-per sample results and dividing by N.

[00117] Another example embodiment may be one or more non-transitory computer readable media storing instructions executed by a processor to perform a sequence comprising treating a multi-sample image as a set of single-sample images, one single sample image for each sample in the multi-sample image, for a first of N samples, subtracting a sample position from a texture filter position, performing single sample texture filtering at the texture filter position to generate a result, combining the results into an accumulator, repeating the steps above for N-1 samples, and resolving the N texture filtering results. The media may include

wherein texture filtering includes bilinear filtering. The media may include said sequence including implementing per sample multi-sample anti-aliasing. The media may include said sequence including implementing per pixel multi-sampling anti-aliasing. The media may include said sequence including processing both per-pixel and per-sample multi-sample anti-aliasing. The media may include said sequence including receiving samples with a linear function and providing a texture filtering result that preserves said linear function. The media may include said sequence including receiving samples with a non-linear function and providing a texture filtering result that preserves said non-linear function. The media may include wherein resolving includes summing the N-per sample results and dividing by N.

[00118] In another example embodiment may be an apparatus comprising a subtractor to subtract a sample position from a texture filter position for a first of N samples, a texture filter to perform single sample texture filtering at the texture filter position, an accumulator to add the result of the filter and a resolve unit to resolve the N texture filtering results. The apparatus may include wherein  $N = 4$ . The apparatus may include wherein said texture filter is a bilinear filtering. The apparatus may include wherein said apparatus to implement per sample multi-sample anti-aliasing. The apparatus may include wherein said apparatus to implement per pixel multi-sampling anti-aliasing. The apparatus may include wherein said apparatus to process both per-pixel and per-sample multi-sample anti-aliasing. The apparatus may include wherein said apparatus to receive samples with a linear function and provide a texture filtering result that preserves said linear function. The apparatus may include wherein said apparatus to receive samples with a non-linear function and provide a texture filtering result that preserves said non-linear function. The apparatus may include wherein the resolve unit to sum the N-sample results and divide by N. The apparatus may include wherein said resolve unit uses a bit shift to divide by N.

[00119] 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.

[00120] 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.

[00121] 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 treating a multi-sample image as a set of single-sample images, one single  
3 sample image for each sample in the multi-sample image;  
4 for a first of N samples, subtracting a sample position relative to a pixel center  
5 from a texture filter position;  
6 performing single sample texture filtering at the texture filter position to  
7 generate a result;  
8 combining the results into an accumulator;  
9 repeating the steps above for N-1 samples; and  
10 resolving the N texture filtering results.
- 1 2. The method of claim 1 wherein texture filtering includes bilinear filtering.
- 1 3. The method of claim 1 including implementing per sample multi-sample anti-  
2 aliasing.
- 1 4. The method of claim 1 including implementing per pixel multi-sampling anti-  
2 aliasing.
- 1 5. The method of claim 1 including processing both per-pixel and per-sample  
2 multi-sample anti-aliasing.
- 1 6. The method of claim 1 including receiving samples with a linear function and  
2 providing a texture filtering result that preserves said linear function.
- 1 7. The method of claim 1 including receiving samples with a non-linear function  
2 and providing a texture filtering result that preserves said non-linear function.
- 1 8. The method of claim 1 wherein resolving includes summing the N-per sample  
2 results and dividing by N.

- 1 9. One or more non-transitory computer readable media storing instructions  
2 executed by a processor to perform a sequence comprising:  
3 treating a multi-sample image as a set of single-sample images, one single  
4 sample image for each sample in the multi-sample image;  
5 for a first of N samples, subtracting a sample position from a texture filter  
6 position;  
7 performing single sample texture filtering at the texture filter position to  
8 generate a result;  
9 combining the results into an accumulator;  
10 repeating the steps above for N-1 samples; and  
11 resolving the N texture filtering results.
- 1 10. The media of claim 9 wherein texture filtering includes bilinear filtering.
- 1 11. The media of claim 9, said sequence including implementing per sample  
2 multi-sample anti-aliasing.
- 1 12. The media of claim 9, said sequence including implementing per pixel multi-  
2 sampling anti-aliasing.
- 1 13. The media of claim 9, said sequence including processing both per-pixel and  
2 per-sample multi-sample anti-aliasing.
- 1 14. The media of claim 9, said sequence including receiving samples with a linear  
2 function and providing a texture filtering result that preserves said linear function.
- 1 15. The media of claim 9, said sequence including receiving samples with a non-  
2 linear function and providing a texture filtering result that preserves said non-linear  
3 function.

1 16. The media of claim 9 wherein resolving includes summing the N-per sample  
2 results and dividing by N.

1 17. An apparatus comprising:  
2 a subtractor to subtract a sample position from a texture filter position for a  
3 first of N samples;  
4 a texture filter to perform single sample texture filtering at the texture filter  
5 position;  
6 an accumulator to add the result of the filter; and  
7 a resolve unit to resolve the N texture filtering results.

1 18. The apparatus of claim 17 wherein said texture filter is a bilinear filtering.

1 19. The apparatus of claim 17 wherein said apparatus to implement per sample  
2 multi-sample anti-aliasing.

1 20. The apparatus of claim 17 wherein said apparatus to implement per pixel  
2 multi-sampling anti-aliasing.

1 21. The apparatus of claim 17 wherein said apparatus to process both per-pixel  
2 and per-sample multi-sample anti-aliasing.

1 22. The apparatus of claim 17 wherein said apparatus to receive samples with a  
2 linear function and provide a texture filtering result that preserves said linear  
3 function.

1 23. The apparatus of claim 17 wherein said apparatus to receive samples with a  
2 non-linear function and provide a texture filtering result that preserves said non-linear  
3 function.

1 24. The apparatus of claim 17 wherein the resolve unit to sum the N-sample  
2 results and divide by N.

- 1 25. The apparatus of claim 24 wherein said resolve unit uses a bit shift to divide
- 2 by N.

1 / 15




FIG. 1

|   |   |     |    |    |    |    |      |
|---|---|-----|----|----|----|----|------|
| 7 | 9 | 11  | 13 | 15 | 17 | 19 | 21   |
| 6 | 8 | 10  | 12 | 14 | 16 | 18 | 20   |
| 5 | 7 | 8.5 | 9  | 11 | 13 | 15 | 16.5 |
| 4 | 6 | 8   | 10 | 12 | 14 | 16 | 18   |
| 3 | 5 | 7   | 9  | 11 | 13 | 15 | 17   |
| 2 | 4 | 6   | 8  | 10 | 12 | 14 | 16   |
| 1 | 3 | 4.5 | 5  | 7  | 9  | 11 | 12.5 |
| 0 | 2 | 4   | 6  | 8  | 10 | 12 | 14   |

FIG. 2

2 / 15

|   |   |   |    |    |    |    |    |
|---|---|---|----|----|----|----|----|
|   | 9 |   |    |    | 17 |    |    |
|   |   |   | 12 |    |    |    | 20 |
| 5 |   |   |    | 13 |    |    |    |
|   |   | 8 |    |    |    | 16 |    |
|   | 5 |   |    |    | 13 |    |    |
|   |   |   | 8  |    |    |    | 16 |
| 1 |   |   |    | 9  |    |    |    |
|   |   | 4 |    |    |    | 12 |    |

FIG. 3

|     |     |     |     |      |      |      |      |
|-----|-----|-----|-----|------|------|------|------|
|     | 8.5 |     |     |      | 16.5 |      |      |
|     |     |     | 8.5 |      |      |      | 16.5 |
| 8.5 |     |     |     | 16.5 |      |      |      |
|     |     | 8.5 |     |      |      | 16.5 |      |
|     | 4.5 |     |     |      | 12.5 |      |      |
|     |     |     | 4.5 |      |      |      | 12.5 |
| 4.5 |     |     |     | 12.5 |      |      |      |
|     |     | 4.5 |     |      |      | 12.5 |      |

FIG. 4

3 / 15




FIG. 5




FIG. 6

4 / 15




FIG. 7

PROCESSOR 200




FIG. 8




FIG. 9




FIG. 10




FIG. 11




FIG. 12




FIG. 13

11 / 15




FIG. 14

12 / 15

GRAPHICS PROCESSOR COMMAND FORMAT  
900

FIG. 15A


GRAPHICS PROCESSOR COMMAND SEQUENCE  
910

FIG. 15B

13 / 15




FIG. 16




FIG. 17




FIG. 18

## INTERNATIONAL SEARCH REPORT

International application No.  
**PCT/US2016/036184****A. CLASSIFICATION OF SUBJECT MATTER****G06T 1/20(2006.01)I, G06T 3/40(2006.01)I, G06T 7/40(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 1/20; G09G 5/00; G06T 5/00; G06T 3/40; G06T 7/40

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) &amp; keywords: multi-sampling, texture, filtering, sample position, resolving

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

| Category* | Citation of document, with indication, where appropriate, of the relevant passages                                     | Relevant to claim No. |
|-----------|------------------------------------------------------------------------------------------------------------------------|-----------------------|
| Y         | US 2010-0141677 A1 (KEVIN HUNTER) 10 June 2010<br>See paragraphs [0010]–[0034]; claim 1; and figures 1–4.              | 1–25                  |
| Y         | US 2013-0106902 A1 (ADVANCED MICRO DEVICES, INC.) 02 May 2013<br>See paragraphs [0017], [0040]; claim 1; and figure 2. | 1–25                  |
| A         | US 2006-0077209 A1 (RUI M. BASTOS et al.) 13 April 2006<br>See paragraphs [0024]–[0040]; and figures 2–3C.             | 1–25                  |
| A         | US 2006-0197778 A1 (JAMES R. PETERSON et al.) 07 September 2006<br>See paragraphs [0022]–[0028]; and figures 1–2.      | 1–25                  |
| A         | US 7502035 B1 (HENRY PACKARD MORETON) 10 March 2009<br>See column 3, line 66 – column 5, line 26; and figures 3–4.     | 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

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

Date of the actual completion of the international search

02 September 2016 (02.09.2016)

Date of mailing of the international search report

**02 September 2016 (02.09.2016)**

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

AHN, Jeong Hwan

Telephone No. +82-42-481-8633



**INTERNATIONAL SEARCH REPORT**

Information on patent family members

International application No.

**PCT/US2016/036184**

| Patent document<br>cited in search report | Publication<br>date | Patent family<br>member(s)                                                                                                                                                                      | Publication<br>date                                                                                                                      |
|-------------------------------------------|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| US 2010-0141677 A1                        | 10/06/2010          | US 2003-0043169 A1<br>US 2006-0203005 A1<br>US 7145577 B2<br>US 7675526 B2<br>US 8068119 B2                                                                                                     | 06/03/2003<br>14/09/2006<br>05/12/2006<br>09/03/2010<br>29/11/2011                                                                       |
| US 2013-0106902 A1                        | 02/05/2013          | CN 102016916 A<br>CN 102016916 B<br>EP 2274725 A1<br>JP 05520286 B2<br>JP 2011-516981 A<br>KR 10-1574279 B1<br>KR 10-2011-0002858 A<br>US 2009-0256848 A1<br>US 8345063 B2<br>WO 2009-124248 A1 | 13/04/2011<br>13/08/2014<br>19/01/2011<br>11/06/2014<br>26/05/2011<br>04/12/2015<br>10/01/2011<br>15/10/2009<br>01/01/2013<br>08/10/2009 |
| US 2006-0077209 A1                        | 13/04/2006          | KR 10-0913459 B1<br>KR 10-2007-0061916 A<br>US 2007-0008336 A1<br>US 7425966 B2<br>US 7576751 B2<br>WO 2006-041953 A2<br>WO 2006-041953 A3                                                      | 25/08/2009<br>14/06/2007<br>11/01/2007<br>16/09/2008<br>18/08/2009<br>20/04/2006<br>28/12/2006                                           |
| US 2006-0197778 A1                        | 07/09/2006          | US 2002-0140706 A1                                                                                                                                                                              | 03/10/2002                                                                                                                               |
| US 7502035 B1                             | 10/03/2009          | None                                                                                                                                                                                            |                                                                                                                                          |