



(51) International Patent Classification:

G06F 17/11 (2006.01) G06T 1/20 (2006.01)
G06F 7/483 (2006.01) G06F 9/38 (2006.01)

(21) International Application Number:

PCT/US2016/035347

(22) International Filing Date:

1 June 2016 (01.06.2016)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

14/749,838 25 June 2015 (25.06.2015) US

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

(72) Inventors: MAIYURAN, Subramaniam; 1915 Black Slate Court, Gold River, California 95670 (US). GARCIA PABON, Jorge F.; 1900 Prairie City Road, Folsom, California 95630 (US). GARG, Ashutosh; 1900 Prairie City Road, Folsom, California 95630 (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).


Declarations under Rule 4.17:

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

Published:

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

(54) Title: HANDLING INSTRUCTIONS THAT REQUIRE ADDING RESULTS OF A PLURALITY OF MULTIPLICATIONS



HANDLING INSTRUCTIONS THAT REQUIRE ADDING RESULTS OF A PLURALITY OF MULTIPLICATIONS

Background

[0001] This relates generally to graphics processors and particularly to handling instructions that perform the addition of results of multiple multiplications.

[0002] A number of equations must be evaluated in graphics processors despite the fact that they involve several multiplications and additions of multiplied terms. Examples of such compound equations include the plane equation, the dot product, and linear interpolation. For example, the plane equation has the form $A * X + B * Y + C$, where A , B , C , X and Y are floating point numbers.

[0003] The entities $A * X$, $B * Y$ and C , each involve one “term” instruction. Thus the plane equation is an example of an equation with three term instructions. An example of an equation with four term instructions is the dot product DP4 calculation.

[0004] Precision is important, particularly in commonly used equations. In addition it is desirable to reduce the power consumption of instruction execution and increase throughput.

Brief Description Of The Drawings

[0005] Some embodiments are described with respect to the following figures:

Figure 1 is a schematic depiction of one embodiment;

Figure 2 is a simplified schematic depiction of part of another embodiment;

Figure 3 is a schematic depiction of still another embodiment;

Figure 4 is a depiction of a first scenario;

Figure 5 is a depiction of the second scenario;

Figure 6 is a depiction of a third scenario;

Figure 7 is a depiction of a fourth scenario;

Figure 8 is a depiction of a fifth scenario;

Figure 9 is a schematic depiction of another embodiment useful by four-term instructions;

Figure 10 is a depiction of a four-term scenario;

Figure 11 is a depiction of another four-term scenario;

Figure 12 is a flow chart for a three term equation; and

Figure 13 is a flow chart for a four term equation.

Figure 14 is a block diagram of a processing system according to one embodiment;

Figure 15 is a block diagram of a processor according to one embodiment;

Figure 16 is a block diagram of a graphics processor according to one embodiment;

Figure 17 is a block diagram of a graphics processing engine according to one embodiment;

Figure 18 is a block diagram of another embodiment of a graphics processor;

Figure 19 is a depiction thread execution logic according to one embodiment;

Figure 20 is a block diagram of a graphics processor instruction format according to some embodiments;

Figure 21 is a block diagram of another embodiment of a graphics processor;

Figure 22A is a block diagram of a graphics processor command format according to some embodiments;

Figure 22B is a block diagram illustrating a graphics processor command sequence according to some embodiments;

Figure 23 is a depiction of an exemplary graphics software architecture according to some embodiments;

Figure 24 is a block diagram illustrating an IP core development system according to some embodiments; and

Figure 25 is a block diagram showing an exemplary system on chip integrated circuit according to some embodiments.

Detailed Description

[0006] Graphics processors are commonly called upon to calculate compound equations that involve both multiplication and addition of multiplication results involving floating point values. For example, when interpolating attributes such as color and texture coordinates, a plane equation (PLN) is analyzed where $PLN = AX + BY + C$, all in floating point numbers. Generally these calculations are performed in the execution units (EUs) of a graphics processing unit (GPU).

[0007] Ideally, the error arising from 3 and 4 term instruction calculations is no more than one least significant bit.

[0008] A base 2 or binary floating point number is represented as:

$$(-1)^s \cdot (1 + M) \cdot 2^{e-127},$$

where the sign s is a value from 0 or 1, the mantissa M is an integer and the exponent e is an integer. All zeros is the smallest exponent and all ones is the largest exponent. If the mantissa is non-zero, then the floating point number goes to “not a number” (NaN). If the mantissa is zero, it goes to infinity.

[0009] Examples of other instructions that can be implemented using some embodiments are a dot product, DP2, which equals $A0 * B0 + A1 * B1$, linear interpolation, $LRP = A * X + B * (1.0 - X)$, dot product DP3 which is equal to $A0 * B0 + A1 * B1 + A2 * B2$ and dot product DP4 which is equal to $A0 * B0 + A1 * B1 + A2 * B2 + A3 * B3$.

[0010] In order to perform three or more term calculations, such as the addition of three terms of a plane equation, all the terms need to have the same exponent e . In order for them all to have the same exponent e , the largest exponent term is used as a reference. The upper bits of an adder are used for the largest exponent term. As used herein, the upper bits or “Hi bits” are the more significant bits and lower bits or “Lo bits” are less significant bits.

[0011] The exponents of the smaller exponent terms (e.g. two smaller terms with a three term equation) are right shifted by an amount dependent on a determination made during a sort operation. The differences between the larger exponent term and each of the other term’s exponents are also calculated.

[0012] A problem can arise when the two largest exponent terms get cancelled when added because one is positive and the other negative and they have the same magnitude. Then the third or smaller term is shifted and dropped from the adder. The erroneous result is then zero, instead of the correct result which is the smaller (third) term.

[0013] In some embodiments, the amount of hardware may be reduced. In some cases, a maximum error of one least significant bit is obtained for a plane instruction.

[0014] An embodiment is depicted in which the same hardware can be used for evaluating any of the compound equations set forth above. However in other embodiments, specialized hardware may be used that only handles one specific equation or less than all of the above-described equations.

[0015] The execution of all these equations is divided into two parts. The first part, which is not discussed in detail here, includes the multipliers that determine the three or more terms of the equations. Conventional multiplier hardware can be used for this part. The second part, discussed in more detail here, is the adder part. For the DP2 instructions, two results are added, for the LRP, DP3 and PLN instructions, three results are added, and for the DP4 instruction, four floating point numbers are added.

[0016] Overlap mantissas occur where exponents of summands differ by the mantissa length. Thus hardware may be able to store, not only the result of adding the overlap mantissas of the upper bits, but all the values that result from adding the different numbers, all without losing information that can cause a precision impact in some embodiments.

[0017] Referring to **Figure 1**, the exponents are supplied to pins along the left column and the mantissas are supplied to pins arranged horizontally across the top of Figure 1. There are two channels in this example, channel 0 and channel 1, so two plane equations can be evaluated at the same time. Two three-term instructions can be executed in parallel or one four-term instruction may be executed in such an embodiment. However, the number of channels may be more or less than two channels.

[0018] In Figure 1 each exponent pin 10 provides an input to a sort block 12. Each mantissa pin 16 provides an input to a different shift block 20. In the case of a single three-term instruction, only one sort block is needed. When two three-term instructions are done in parallel, two sort blocks may be used including sort 0 and sort 1 at the top and bottom of Figure 1.

[0019] Thus the channel 0 uses the top three exponent pins 10 (C0, c1, c0), sort 0 block 12, a multiplexer 18, and the shift blocks 24 (shift0, shift1, shift2), the mantissa pins 16 (C0, c1, and c0) and the five adders: add4 Hi, add0 Hi (upper bits), add1 Hi (upper bits), add0 Lo (upper bits) and Add1 Lo (upper bits).

[0020] The other channel, channel 1, executed at the same time as channel 0, uses the lower three exponent pin 10 (c2, c3, C1), the rightmost mantissa pins 14 (c2, c3, C1), two sort blocks, two multiplexers, shift3, shift4 and shift 5, as well as Add2 Lo, Add2 Hi, Add3 Lo and Add3 Hi.

Bypass0 is used by the channel 0 while bypass 1 is used by the channel 1. Case mux 0 is used for channel 0 and the output of the mantissa 0 is on the connected pin. Case mux1 is used for the channel 1 and its output mantissa 1 is on the connected pin. The adders adder(0) and adder(1) that add the outputs of the shifters and are stored in the registers. For three terms, the addition works as follows:

$$\text{ADDER0} = \text{ADD4Hi} + \text{ADD0Hi} + \text{ADD1Hi}$$

$$\text{ADDER1} = \text{ADD2Lo} + \text{ADD3LO} + \text{ADD4Lo}$$

In case of 4 terms, the addition works as follows:

$$\text{ADDER0} = \text{ADD0Hi} + \text{ADD1Hi} + \text{ADD2Hi} + \text{ADD3Hi}$$

$$\text{ADDER1} = \text{ADD0Lo} + \text{ADD1Lo} + \text{ADD2Lo} + \text{ADD3Lo}$$

$$\text{Carry Adder} = \{\text{ADDER0} + \text{ADDER1 Carry}, \text{ADDER1}\}$$

To add 4 terms, more bits are used to reduce the precision error. In this example, the 3 term case size is doubled for the 4 term case.

[0021] Each of the sort blocks sorts the input exponent values for one channel. A channel of input values may correspond to the three exponents of an equation such as a plane equation. A sort block determines which one of the exponents is the largest exponent and which one is the smallest. A sort block also determines the difference between exponents of terms A and B as well as the difference between exponents of terms A and C. Then it appropriately selects an input to be connected, via multiplexer 18, to a shift block 20 and the extent of shift, based on the nature of the determinations made in the sort block. The sort block also controls the use or non-use of the bypass blocks.

[0022] For example, depending on the results of the exponent differences identified in sorting, a multiplexer 18 selects the appropriate input to pass on to the connected shift block 20.

[0023] In the sort block, based on the determination of the type of transaction, the amount of shifting in a shift block 20 is determined. Similarly based on the determination of the sort block, the bypass is set to operate or not operate to bypass adders and go to the output.

[0024] Thus in one embodiment, for a plane equation on channel 0, $A * X$ is exponent (Exp) Src0, $B * Y$ is Exp Src1 and C is Exp SrcC0, each on a different pin 10. Channel 1 has $A * X$ at Exp Src2, $B * Y$ at Exp Src3 and C at Exp SrcC1.

[0025] For a two channel dot product (DP) 2 calculation, channel 0 has $A0 * B0$ at Exp Src0, $A1 * B1$ at Exp Src1, and Exp SrcC0 = 0. Channel 1 has $A0 * B0$ at Exp Src2, $A1 * B1$ at Exp Src3 and Exp SrcC1 is 0.

[0026] The LRP calculation has in channel 0, $A * X$ as Exp Src0, $B * (1.0 - X)$ as Exp Src1 and Exp SrcC0 and Exp SrcC1 are all zero. In channel 1, $A * X$ is Exp Src2, $B * (1.0 - X)$ as Exp Src3 and SrcC1 is zero.

[0027] For the four term instruction dot product DP3, $A0 * B0$ is Exp Src0, $A1 * B1$ is Exp Src1 and $A2 * B2$ is Exp Src2, and Exps Src3, SrcC0 and SrcC1 are all zero.

[0028] For the four term dot product DP4 calculation, $A0 * B0$ is Exp Src0, $A1 * B1$ is Exp Src1, $A2 * B2$ is Exp Src2, $A3 * B3$ is a Exp Src3, while SrcC0 and SrcC1 are zero.

[0029] Referring now to **Figure 2**, showing the operation of a sort block 12, to add all the terms, the largest exponent is identified. Based on that identification, all the mantissas are shifted to the right and then they are added. There are three sort blocks to handle two three-term instructions or one four-term instruction.

[0030] Figure 2 shows that, based on the exponent values of the three terms, there are different scenarios. Figure 2 also shows how the shift values are calculated and also how to decide each scenario. In other words once the three values are

determined for A, B and C and $A - B$ is determined and $A - C$ is determined, then the apparatus can determine all the necessary settings of other components.

[0031] Specifically, a sort block 12 compares the three exponent value inputs (for a three term equation) and sorts them to produce three ordered values. For example if A is the largest exponent, then $A > B > C$. The sort with four inputs produces $A \geq B \geq C \geq D$.

[0032] For the three term inputs, the difference between the largest value and the other two terms determines the scenario type. Figures 4-11 show the possible scenarios when adding different floating point numbers. Each of Figures 4-11 show the possible solution space and also the result of a sticky bit. The sticky bit is a logical OR of all the remaining bits of the adder. A sticky bit is used to calculate a rounding result. N is the number of bits of the adder.

[0033] Channel0 is in dashed lines in **Figure 3** and channel1 is shown shaded in Figure 3.

[0034] Thus, as shown in **Figure 4**, in scenario 1, if $A \geq B$ and $B \geq C$, and $C + 3 \geq A$, then all the terms have similar exponents and the three terms are simply added. Only the adder result is used without a bypass.

[0035] In scenario 2, shown in **Figure 5**, if $A \geq B \geq C$ and $B + 3 \geq A$ and $C + N + 3 \geq A$, the result is obtained by adding the three terms. But the three terms can cancel out and the result should then be the lower bits of C . Therefore the bypass has the lower bits of the C input and the result of the adder is concatenated with the bypass. This bypass prevents dropping the lower bits of C from the final result. The bypass is selected by the sort block.

[0036] Scenario 3, shown in **Figure 6**, occurs when $A \geq B \geq C$ and $B + 3 \geq A$ and $A > C + N + 3$. If A has the same magnitude as B but opposite sign, the result is C . Otherwise the result is $A + B$. Therefore if the adder result is zero, the output is taken from the bypass as determined in the sort block. If the adder result is not zero, then the output is taken from the adder to get the output $A + B$.

[0037] In scenario 4, shown in **Figure 7**, if $A > B + 3$ which is greater than or equal to C , A cannot be canceled. The result is then taken from the output of the adder.

[0038] In scenario 5, shown in **Figure 8**, if $A > B + N + 3 \geq C$, then B and C are too small, and the result is A . The A input uses a bypass (as selected by the sort block) and the adder is not even used.

[0039] If one or more inputs have erroneous values, the blocks are bypassed and only a signal may be sent saying that the output is erroneous. Because the blocks are bypassed, the inputs of the shifters do not change and the previous values remain. Avoiding the toggles saves power.

[0040] With four term instructions, all three sort blocks 12 are used as shown in **Figure 9**. For four term instructions, the first scenario is shown in **Figure 10**. If $A > (B + N + 3)$ and $B \geq C \geq D$, A uses the bypass and goes directly to the output.

[0041] The second scenario, shown in **Figure 11**, occurs if $A \leq (B + N + 3)$ and $B \geq C \geq D$. The sort block then provides the inputs to the shifters and then all four values are added, concatenating the outputs of both adders (which adders) to provide a better precision.

[0042] Sequences shown in **Figures 12 and 13**, 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. Generally they may be implemented in many embodiments in graphics processors.

[0043] The sequence 100, shown in Figure 12, begins by determining whether an NaN or infinite input is detected as indicated in diamond 102. If so, an output is NaN or infinity as indicated in block 104, generally indicating that the flow cannot proceed.

[0044] Otherwise, the exponents are sorted as indicated in block 106. This generally means that the exponents are sorted by size and determinations are made that affect the other components or sequences. More particularly as indicated in block 108, the inputs to the multiplexers are selected, the amount of shift may be

determined and, whether or not to use bypass may all be based on the results of the sort.

[0045] Then a check at diamond 110 determines whether $A \geq B \geq C$ and $C + 3 \geq A$. If so, the terms are simply added as indicated in block 112.

[0046] Otherwise, a check at diamond 114 determines whether $A \geq B > C$ and whether $B + 3 \geq A$ and further whether $C + N + 3 \geq A$. If so, the lower bits of C are output by bypassing the addition stage as indicated in block 116.

[0047] If the result of diamond 114 is no, then a check at diamond 118 determines $A \geq B > C$, whether $B + 3 \geq A$ and $A > C + N + 3$. If so, a check at diamond 120 determines whether $A = B$ with opposite signs. If so, C is outputted as indicated in block 122. Otherwise $A + B$ is outputted as indicated in block 124.

[0048] If the check at diamond 118 is negative, then a check at diamond 126 determines whether $A > B + C \geq C$. If so, the output comes from the adder as indicated in block 128.

[0049] If the check at diamond 126 is negative, then a check at diamond 130 determines whether $A > B + 3 \geq C$. If so, the result is A using the bypass as indicated in block 132. If not, then an error condition may apply.

[0050] In the case of a four term instruction, the flow, as shown in Figure 13, is the same up through block 108 as shown in Figure 12. After block 108 a check at diamond 140 determines whether $A > (B + N + 3)$ and whether $B \geq C \geq D$. If so, A is the output using bypass as indicated in block 142.

[0051] If the check at diamond 140 is negative then a check at diamond 144 determines whether $A \leq (B + N + 3)$ and whether $B \geq C \geq D$. If so, then all four values are added and the results of both adders are concatenated as indicated in block 146. If the check at diamond 144 is negative then an error may be indicated.

[0052] 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.

[0053] **Figure 14** 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.

[0054] 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.

[0055] 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 as a Digital Signal Processor (DSP).

[0056] 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.

[0057] 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.

[0058] 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.

[0059] In some embodiments, ICH 130 enables peripherals to connect to memory device 120 and processor 102 via a high-speed I/O bus. The I/O peripherals include, but are not limited to, an audio controller 146, a firmware interface 128, a wireless transceiver 126 (e.g., Wi-Fi, Bluetooth), a data storage device 124 (e.g., hard disk drive, flash memory, etc.), and a legacy I/O controller 140 for coupling legacy (e.g., Personal System 2 (PS/2)) devices to the system. One or more Universal Serial Bus (USB) controllers 142 connect input devices, such as keyboard and mouse 144 combinations. A network controller 134 may also couple to ICH 130. In some embodiments, a high-performance network controller (not shown) couples to processor bus 110. It will be appreciated that the system 100 shown is exemplary and not limiting, as other types of data processing systems that are differently configured may also be used. For example, the I/O controller hub 130 may be integrated within the one or more processor 102, or the memory controller hub 116 and I/O controller hub 130 may be integrated into a discrete external graphics processor, such as the external graphics processor 112.

[0060] **Figure 15** 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 15** having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such. 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.

[0061] 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.

[0062] 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).

[0063] 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.

[0064] 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.

[0065] 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.

[0066] 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.

[0067] 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.

[0068] **Figure 16** 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.

[0069] 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) 421MVC-1, and Joint Photographic Experts Group (JPEG) formats such as JPEG, and Motion JPEG (MJPEG) formats.

[0070] 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.

[0071] 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.

[0072] 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.

[0073] 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.

[0074] **Figure 17** 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 16**. Elements of **Figure 17** having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such.

[0075] 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.

[0076] 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.

[0077] 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).

[0078] 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.

[0079] 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.

[0080] 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.

[0081] **Figure 18** is a block diagram of another embodiment of a graphics processor 500. Elements of **Figure 18** 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 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.

[0083] 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.

[0084] 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.

[0085] **Figure 19** illustrates thread execution logic 600 including an array of processing elements employed in some embodiments of a GPE. Elements of **Figure 19** 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.

[0086] 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.

[0087] 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).

[0088] 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.

[0089] 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.

[0090] 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.

[0091] 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 18**) dispatches vertex processing, tessellation, or geometry processing threads to thread execution logic 600 (**Figure 19**). In some embodiments, thread dispatcher 604 can also process runtime thread spawning requests from the executing shader programs.

[0092] 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.

[0093] 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.

[0094] **Figure 20** 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.

[0095] 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.

[0096] 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.

[0097] 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.

[0098] 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.

[0099] 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.

[0100] 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.

[0101] **Figure 21** is a block diagram of another embodiment of a graphics processor 800. Elements of **Figure 21** 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.

[0102] 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.

[0103] 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.

[0104] 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.

[0105] 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.

[0106] 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.

[0107] 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.

[0108] 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.

[0109] 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.

[0110] 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.

[0111] 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.

[0112] 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.

[0113] **Figure 22A** is a block diagram illustrating a graphics processor command format 900 according to some embodiments. **Figure 22B** is a block diagram illustrating a graphics processor command sequence 910 according to an embodiment. The solid lined boxes in **Figure 22A** 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 22A** 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.

[0114] 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.

[0115] The flow diagram in **Figure 22B** 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.

[0116] 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.

[0117] In some embodiments, a pipeline select command 913 is used when a command sequence requires the graphics processor to explicitly switch between pipelines. In some embodiments, a pipeline select command 913 is required only once within an execution context before issuing pipeline commands unless the context is to issue commands for both pipelines. In some embodiments, a pipeline flush command is 912 is required immediately before a pipeline switch via the pipeline select command 913.

[0118] 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.

[0119] 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.

[0120] 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.

[0121] 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.

[0122] 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.

[0123] 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.

[0124] 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.

[0125] 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.

[0126] 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.

[0127] **Figure 23** 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.

[0128] 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.

[0129] 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.

[0130] 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.

[0131] 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.

[0132] **Figure 24** 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.

[0133] The RTL design 1115 or equivalent may be further synthesized by the design facility into a hardware model 1120, which may be in a hardware description language (HDL), or some other representation of physical design data. The HDL may be further simulated or tested to verify the IP core design. The IP core design can be stored for delivery to a 3rd party fabrication facility 1165 using non-volatile memory 1140 (e.g., hard disk, flash memory, or any non-volatile storage medium). Alternatively, the IP core design may be transmitted (e.g., via the Internet) over a wired connection 1150 or wireless connection 1160. The fabrication facility 1165 may then fabricate an integrated circuit that is based at least in part on the IP core design. The fabricated integrated circuit can be configured to perform operations in accordance with at least one embodiment described herein.

[0134] **Figure 25** is a block diagram illustrating an exemplary system on a chip integrated circuit 1200 that may be fabricated using one or more IP cores, according to an embodiment. The exemplary integrated circuit includes one or more application processors 1205 (e.g., CPUs), at least one graphics processor 1210, and may additionally include an image processor 1215 and/or a video processor 1220, any of which may be a modular IP core from the same or multiple different design facilities. The integrated circuit includes peripheral or bus logic including a USB controller 1225, UART controller 1230, an SPI/SDIO controller 1235, and an I²S/I²C controller 1240. Additionally, the integrated circuit can include a display device 1245 coupled to one or more of a high-definition multimedia interface (HDMI) controller 1250 and a mobile industry processor interface (MIPI) display interface 1255. Storage may be provided by a flash memory subsystem 1260 including flash

memory and a flash memory controller. Memory interface may be provided via a memory controller 1265 for access to SDRAM or SRAM memory devices. Some integrated circuits additionally include an embedded security engine 1270.

[0135] 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.

[0136] The following clauses and/or examples pertain to further embodiments:

One example embodiment may be a method comprising receiving at least three multiplied floating point terms, adding the three terms by sorting the floating point numbers based on exponent size, and bypassing addition of lower bits of a smallest exponent term to present the lower bits as an output. The method may also include receiving both three and four term equations. The method may also include calculating two three term equations at the same time. The method may also include handling dot products, plane equations and linear interpolation in the same hardware. The method may also include sorting exponents based on size and determining differences between largest and smaller exponents. The method may also include selecting an input to a multiplexer for outputting, determining the amount of shift and determining whether or not to bypass an adder based on the sorting. The method may also include determining for exponents A , B , and C of a three term equation whether $A \geq B > C$ and whether $B + 3 \geq A$ and whether $C + N + 3 \geq A$ and if so, bypassing an adder and outputting the lower bits of C . The method may also include determining for exponents A , B , and C of a three term equation whether $A \geq B > C$ and whether $B + 3 \geq A$ and whether $A > C + N + 3$ and, if so, determining if A has the same magnitude as B but opposite sign. The method may also include taking as the result C and otherwise taking as the result $A + B$. The method may also include determining if $A > (B + N + 3) \geq C$ and if so bypassing A to the output.

[0137] Another example embodiment may be one or more non-transitory computer readable media storing instructions executed by a processor to perform a sequence comprising receiving at least three multiplied floating point terms, sorting the floating point numbers based on exponent size, and bypassing addition of lower bits of a

smallest exponent term to present the lower bits as an output. The media may include said sequence including receiving both three and four term equations. The media may include said sequence including calculating two three term equations at the same time. The media may include said sequence including handling dot products, plane equations and linear interpolation in the same hardware. The media may include said sequence including sorting exponents based on size and determining differences between largest and smaller exponents. The media may include said sequence including selecting an input to a multiplexer for outputting, determining the amount of shift and determining whether or not to bypass an adder based on the sorting. The media may include said sequence including determining for exponents A , B , and C of a three term equation whether $A \geq B > C$ and whether $B + 3 \geq A$ and whether $C + N + 3 \geq A$ and if so, bypassing an adder and outputting the lower bits of C . The media may include said sequence including determining for exponents A , B , and C of a three term equation whether $A \geq B > C$ and whether $B + 3 \geq A$ and whether $A > C + N + 3$ and, if so, determining if A has the same magnitude as B but opposite sign. The media may include said sequence including if so, taking as the result C and otherwise taking as the result $A + B$. The media may include said sequence including determining if $A > (B + N + 3) \geq C$ and if so bypassing A to the output.

[0138] In another example embodiment may be an apparatus comprising a sort unit to obtain at least three multiple floating point terms, sort the floating point numbers based on exponent size, and a switch coupled to said sort unit to bypass addition of lower bits of a smallest exponent term to present the lower bits as an output. The apparatus may include said apparatus to receive both three and four term equations. The apparatus may include said apparatus to calculate two three term equations at the same time. The apparatus may include said apparatus to handle dot products, plane equations and linear interpolation in the same hardware. The apparatus may include said apparatus to sort exponents based on size and determining differences between largest and smaller exponents. The apparatus may include a shift block and a multiplexer to select an input from the unit to output, determine the amount of shift and determine whether or not to bypass all based on an indication from the sort unit. The apparatus may include said sort unit to determine for exponents A , B , and

C of a three term equation whether $A \geq B > C$ and whether $B + 3 \geq A$ and whether $C + N + 3 \geq A$ and if so, bypassing an adder and outputting the lower bits of C . The apparatus may include said sort unit to determine for exponents A , B , and C of a three term equation whether $A \geq B > C$ and whether $B + 3 \geq A$ and whether $A > C + N + 3$ and, if so, determining if A has the same magnitude as B but opposite sign. The apparatus may include said sort unit to take as the result C and otherwise taking as the result $A + B$. The apparatus may include said sort unit to determine if $A > (B + N + 3) \geq C$ and if so bypassing A to the output.

[0139] 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.

[0140] 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. A method comprising:
receiving at least three multiplied floating point terms;
sorting the floating point numbers based on exponent size; and
bypassing addition of lower bits of a smallest exponent term to present the lower bits as an output.
2. The method of claim 1 including receiving both three and four term equations.
3. The method of claim 1 including calculating two three term equations at the same time.
4. The method of claim 1 including handling dot products, plane equations and linear interpolation in the same hardware.
5. The method of claim 1 including sorting exponents based on size and determining differences between largest and smaller exponents.
6. The method of claim 1 including selecting an input to a multiplexer for outputting, determining the amount of shift and determining whether or not to bypass an adder based on the sorting.
7. The method of claim 1 including determining for exponents A , B , and C of a three term equation whether $A \geq B > C$ and whether $B + 3 \geq A$ and whether $C + N + 3 \geq A$ and if so, bypassing an adder and outputting the lower bits of C .
8. The method of claim 1 including determining for exponents A , B , and C of a three term equation whether $A \geq B > C$ and whether $B + 3 \geq A$ and whether $A > C + N + 3$ and, if so, determining if A has the same magnitude as B but opposite sign.

9. The method of claim 8 including if so, taking as the result C and otherwise taking as the result $A + B$.
10. The method of claim 1 including determining if $A > (B + N + 3) \geq C$ and if so bypassing A to the output.
11. One or more non-transitory computer readable media storing instructions executed by a processor to perform a sequence comprising:
 - receiving at least three multiplied floating point terms;
 - sorting the floating point numbers based on exponent size; and
 - bypassing addition of lower bits of a smallest exponent term to present the lower bits as an output.
12. The media of claim 11, said sequence including receiving both three and four term equations.
13. The media of claim 11, said sequence including calculating two three term equations at the same time.
14. The media of claim 11, said sequence including handling dot products, plane equations and linear interpolation in the same hardware.
15. The media of claim 11, said sequence including sorting exponents based on size and determining differences between largest and smaller exponents.
16. The media of claim 11, said sequence including selecting an input to a multiplexer for outputting, determining the amount of shift and determining whether or not to bypass an adder based on the sorting.
17. The media of claim 11, said sequence including determining for exponents A, B , and C of a three term equation whether $A \geq B > C$ and whether $B + 3 \geq A$ and

whether $C + N + 3 \geq A$ and if so, bypassing an adder and outputting the lower bits of C .

18. The media of claim 11, said sequence including determining for exponents A , B , and C of a three term equation whether $A \geq B > C$ and whether $B + 3 \geq A$ and whether $A > C + N + 3$ and, if so, determining if A has the same magnitude as B but opposite sign.

19. The media of claim 18, said sequence including if so, taking as the result C and otherwise taking as the result $A + B$.

20. The media of claim 11, said sequence including determining if $A > (B + N + 3) \geq C$ and if so bypassing A to the output.

21. An apparatus comprising:

a sort unit to obtain at least three multiplied floating point terms and to sort the floating point numbers based on exponent size; and

a switch coupled to said sort unit to bypass addition of lower bits of a smallest exponent term and to present the lower bits as an output.

22. The apparatus of claim 21, said apparatus to receive both three and four term equations.

23. The apparatus of claim 21, said apparatus to calculate two three term equations at the same time.

24. The apparatus of claim 21, said apparatus to handle dot products, plane equations and linear interpolation in the same hardware.

25. The apparatus of claim 21, said apparatus to sort exponents based on size and determining differences between largest and smaller exponents.

1 / 21




FIG. 1




FIG. 2

3 / 21




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8

6 / 21




FIG. 9




FIG. 10




FIG. 11

8 / 21




FIG. 12

9 / 21




FIG. 13

10 / 21




FIG. 14




FIG. 15




FIG. 16




FIG. 17




FIG. 18







FIG. 19

GRAPHICS CORE INSTRUCTION
FORMATS

700



OPCODE DECODE

740




FIG. 20

17 / 21




FIG. 21

18 / 21

GRAPHICS PROCESSOR COMMAND FORMAT
900

FIG. 22A


GRAPHICS PROCESSOR COMMAND SEQUENCE
910

FIG. 22B

19 / 21




FIG. 23




FIG. 24




FIG. 25

A. CLASSIFICATION OF SUBJECT MATTER**G06F 17/11(2006.01)I, G06F 7/483(2006.01)I, G06T 1/20(2006.01)I, G06F 9/38(2006.01)I**

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 17/11; G06F 7/487; G06F 7/485; G06F 7/38; G06F 7/483; G06T 1/20; G06F 9/38

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: floating point, multiply, adder, sort, exponent size, smallest, bypass, and similar terms.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	RIDHI SETH, 'Design and Implementation of Low power High Speed Floating Point Adder and Multiplier,' PhD diss., THAPAR UNIVERSITY PATIALA, June 2010 (http://dspace.thapar.edu:8080/jspui/handle/10266/1333) See page 18 and figure 2.2.	1-25
A	ANDRE GUNTORO et al., 'Low-Latency VLSI Architecture of a 3-Input Floating-Point Adder,' In: Proceedings of Asia Pacific Conference on Circuits and Systems (APCCAS 2008), pp. 180-183, 30 November-03 December 2008 See pages 181-182 and figure 3.	1-25
A	EKTA et al., 'Low Power Binary Floating Multiplier using Bypassing Technique,' International Journal of Emerging Technology and Advanced Engineering, Vol. 4, No. 9, September 2014 (http://www.ijetae.com/files/Volume4Issue9/IJETAE_0914_25.pdf) See section IV and figure 3.	1-25
A	US 2014-0188966 A1 (SAMEH GALAL et al.) 03 July 2014 See paragraphs [0030]-[0035] and figures 2-3.	1-25
A	US 6,996,596 B1 (YING-WAI HO et al.) 07 February 2006 See column 7, line 1 - column 9, line 7; and figure 2.	1-25



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

24 August 2016 (24.08.2016)

Date of mailing of the international search report

24 August 2016 (24.08.2016)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

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

Facsimile No. +82-42-481-8578

Authorized officer

LEE, EUN KYU

Telephone No. +82-42-481-3580



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2016/035347Patent document
cited in search reportPublication
datePatent family
member(s)Publication
date

US 2014-0188966 A1

03/07/2014

US 8892619 B2

18/11/2014

US 6996596 B1

07/02/2006

None