




- (51) International Patent Classification:
G06K 9/66 (2006.01) G06N 3/08 (2006.01)
- (21) International Application Number:
PCT/CN2017/079719
- (22) International Filing Date:
07 April 2017 (07.04.2017)
- (25) Filing Language: English
- (26) Publication Language: English
- (71) Applicant: INTEL CORPORATION [US/US]; 2200 Mission College Blvd., Santa Clara, California 95054 (US).
- (72) Inventors; and
- (71) Applicants (for BZ only): GUO, Yiwen [CN/CN]; 8F, Raycom Infotech Park A, No. 2 KeXueYuan South Road, ZhongGuanCun, HaiDian District, Beijing, 100190 (CN). HOU, Yuqing [CN/CN]; 8F, Raycom Infotech Park A, No. 2 KeXueYuan South Road, ZhongGuanCun, HaiDian District, Beijing, 100190 (CN). YAO, Anbang [CN/CN]; 8F, Raycom Infotech Park A, No. 2 KeXueYuan South Road, ZhongGuanCun, HaiDian District, Beijing, 100190 (CN). CAI, Dongqi [CN/CN]; 8F, Raycom Infotech Park A, No. 2 KeXueYuan South Road, ZhongGuanCun, HaiDian District, Beijing, 100190 (CN). WANG, Libin [CN/CN]; 9F, Raycom Infotech Park A, No. 2 KeXueYuan South

Road, ZhongGuanCun, HaiDian District, Beijing, 100190 (CN). XU, Lin [CN/CN]; 8F, Raycom Infotech Park A, No. 2 KeXueYuan South Road, ZhongGuanCun, HaiDian District, Beijing, 100190 (CN). HU, Ping [CN/CN]; 8F, Raycom Infotech Park A, No. 2 KeXueYuan South Road, ZhongGuanCun, HaiDian District, Beijing, 100190 (CN). WANG, Shandong [CN/CN]; 8F, Raycom Infotech Park A, No. 2 KeXueYuan South Road, ZhongGuanCun, HaiDian District, Beijing, 100190 (CN). CHENG, Wenhua [CN/CN]; 8F, Raycom Infotech Park A, No. 2 KeXueYuan South Road, ZhongGuanCun, HaiDian District, Beijing, 100190 (CN). CHEN, Yurong [CN/CN]; Room 503, Unit 1, Building 4, QingFengHuaJingYuan, HaiDian District, Beijing, 100085 (CN).

(74) Agent: NTD PATENT AND TRADEMARK AGENCY LIMITED; 10th Floor, Block A, Investment Plaza, 27 Jinrongdajie, Xicheng District, Beijing 100033 (CN).

(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, DJ, 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, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM,

(54) Title: METHODS AND SYSTEMS FOR BUDGETED AND SIMPLIFIED TRAINING OF DEEP NEURAL NETWORKS



(57) Abstract: Methods and systems for budgeted and simplified training of deep neural networks (DNNs) are disclosed. In one example, a trainer is to train a DNN using a plurality of training sub-images derived from a down-sampled training image. A tester is to test the trained DNN using a plurality of testing sub-images derived from a down-sampled testing image. In another example, in a recurrent deep Q-network (RDQN) having a local attention mechanism located between a convolutional neural network (CNN) and a long-short time memory (LSTM), a plurality of feature maps are generated by the CNN from an input image. Hard-attention is applied by the local attention mechanism to the generated plurality of feature maps by selecting a subset of the generated feature maps. Soft attention is applied by the local attention mechanism to the selected subset of generated feature maps by providing weights to the selected subset of generated feature maps in obtaining weighted feature maps. The weighted feature maps are stored in the LSTM. A Q value is calculated for different actions based on the weighted feature maps stored in the LSTM.

FIG. 15



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:

- *of inventorship (Rule 4.17(iv))*

Published:

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

METHODS AND SYSTEMS FOR BUDGETED AND SIMPLIFIED TRAINING OF DEEP NEURAL NETWORKS

FIELD

5 [0001] Embodiments of the invention are in the field of data processing including image processing, graphics processing and machine learning. More particularly, embodiments of the invention relate to methods and systems for budgeted and simplified training of deep neural networks (DNNs).

BACKGROUND

10 [0002] Current parallel graphics data processing includes systems and methods developed to perform specific operations on graphics data such as, for example, linear interpolation, tessellation, rasterization, texture mapping, depth testing, etc. Traditionally, graphics processors used fixed function computational units to process graphics data; however, more recently, portions of graphics processors have been made programmable,
15 enabling such processors to support a wider variety of operations for processing vertex and fragment data.

[0003] To further increase performance, graphics processors typically implement processing techniques such as pipelining that attempt to process, in parallel, as much graphics data as possible throughout the different parts of the graphics pipeline. Parallel
20 graphics processors with single instruction, multiple thread (SIMT) architectures are designed to maximize the amount of parallel processing in the graphics pipeline. In an SIMT architecture, groups of parallel threads attempt to execute program instructions synchronously together as often as possible to increase processing efficiency. A general overview of software and hardware for SIMT architectures can be found in Shane Cook,
25 *CUDA Programming* Chapter 3, pages 37-51 (2013).

[0004] Machine learning has been successful at solving many kinds of tasks. The computations that arise when training and using machine learning algorithms (e.g., neural networks) lend themselves naturally to efficient parallel implementations. Accordingly, parallel processors such as general-purpose graphic processing units (GPGPUs) have
30 played a significant role in the practical implementation of deep neural networks. Parallel graphics processors with single instruction, multiple thread (SIMT) architectures

are designed to maximize the amount of parallel processing in the graphics pipeline. In an SIMT architecture, groups of parallel threads attempt to execute program instructions synchronously together as often as possible to increase processing efficiency. The efficiency provided by parallel machine learning algorithm implementations allows the use of high capacity networks and enables those networks to be trained on larger datasets.

5 **[0005]** Deep neural networks (DNNs) can perform deep machine learning useful in computer vision and image recognition applications because of its feature recognition capabilities. DNNs include layers with nodes which are organized into a set of “filters,” which can act as feature detectors. The output of each set of filters is propagated to nodes in successive layers of the network. DNN processing can be computational intensive at each layer having a number of nodes with a number of parameters to be computed for image detection and processing applications. Furthermore, training and learning for DNNs can be extensive requiring large amounts of training data and use of parameters if the number layers and nodes are deep. Thus, what is needed is improved training and learning techniques for DNNs.

10

15

BRIEF DESCRIPTION N OF THE DRAWINGS

- [0006] The appended drawings illustrate examples and are, therefore, exemplary embodiments and not considered to be limiting in scope.
- [0007] FIG. 1 is a block diagram illustrating a computer system configured to
5 implement one or more aspects of exemplary embodiments described herein.
- [0008] FIG. 2A-2D illustrate a parallel processor components according to an exemplary embodiment.
- [0009] FIGS. 3A-3B are block diagrams of graphics multiprocessors according to exemplary embodiments.
- 10 [0010] FIG. 4A-4F illustrate an exemplary architecture in which a plurality of Graphic Processing Units (GPUs) are communicatively coupled to a plurality of multi-core processors.
- [0011] FIG. 5 illustrates a graphics processing pipeline according to an exemplary embodiment.
- 15 [0012] FIG. 6 illustrates a machine learning software stack according to an exemplary embodiment.
- [0013] FIG. 7 illustrates a highly-parallel general-purpose graphics processing unit according to an exemplary embodiment.
- [0014] FIG. 8 illustrates a multi-GPU computing system according to an
20 exemplary embodiment.
- [0015] FIGS. 9A-9B illustrate layers of exemplary deep neural networks.
- [0016] FIG. 10 illustrates an exemplary recurrent neural network.
- [0017] FIG. 11 illustrates exemplary embodiment of training and deployment of a deep neural network.
- 25 [0018] FIG. 12 is an exemplary block diagram illustrating distributed learning.
- [0019] FIG. 13 illustrates an exemplary inferencing system on a chip (SOC) suitable for performing inferencing using a trained model.
- [0020] FIG. 14 is an exemplary block diagram of a basic training and learning architecture having a Deep Neural Network (DNN) training and learning system to
30 receive training data for learning and training a DNN.

- [0021] FIG. 15 illustrates an exemplary process for selecting image data for a learning machine.
- [0022] FIG. 16A is an exemplary block diagram of a learning system which can
5 implement budgeted training and learning techniques for a DNN.
- [0023] FIGS. 16B-16C illustrate exemplary flow diagrams of operations to train and test a DNN.
- [0024] FIG. 17A is a block diagram of an improved Deep Recurrent Q-Network (DRQN) architecture according to an exemplary embodiment.
- 10 [0025] FIG. 17B illustrates a flow diagram of an exemplary operation for the improved DRQN of FIG. 17A.
- [0026] FIG. 18A is a block diagram of an improved DRQN architecture according to another exemplary embodiment.
- [0027] FIG. 18B illustrates a flow diagram of an exemplary operation for the
15 improved DRQN of FIG. 18A.
- [0028] FIG. 19 illustrates a block diagram of a processing system according to an exemplary embodiment.
- [0029] FIG. 20 illustrates an exemplary block diagram of an embodiment of a processor having one or more processor cores, an integrated memory controller, and an
20 integrated graphics processor.
- [0030] FIG. 21 illustrates an exemplary block diagram of a graphics processor.
- [0031] FIG. 22 illustrates a block diagram of a graphics processing engine of a graphics processor according to exemplary embodiments.
- [0032] FIG. 23 illustrates a block diagram of another exemplary embodiment of a
25 graphics processor.
- [0033] FIG. 24 illustrates thread execution logic including an array of processing elements employed in exemplary embodiments of a graphics processing engine (GPE).
- [0034] FIG. 25 illustrates a block diagram of a graphics processor instruction formats according to exemplary embodiments.

[0035] FIG. 26 illustrates a block diagram of an exemplary embodiment of a graphics processor.

[0036] FIG. 27A illustrates a block diagram of a graphics processor command format according to an exemplary embodiment.

5 [0037] FIG. 27B illustrates a block diagram of a graphics processor command sequence according to an exemplary embodiment.

[0038] FIG. 28 illustrates exemplary graphics software architecture for a data processing system according exemplary embodiments.

[0039] FIG. 29 illustrates a block diagram of an IP core development system that
10 may be used to manufacture an integrated circuit (IC) to perform operations according to an exemplary embodiment.

[0040] FIG. 30 illustrates a block diagram of an exemplary system on a chip IC that may be fabricated using or one more IP cores according to an exemplary embodiment.

15 [0041] FIG. 31 illustrates a block diagram of an exemplary graphics processor on a system on a chip IC that may be fabricated using one or more IP cores according to an exemplary embodiment.

[0042] FIG. 32 illustrates a block diagram of an exemplary additional graphics
20 processor of a system on a chip IC that may be fabricated using one or more IP cores according to an exemplary embodiment.

DETAILED DESCRIPTION

[0043] In some embodiments, a graphics processing unit (GPU) is communicatively coupled to host/processor cores to accelerate graphics operations, machine-learning operations, pattern analysis operations, and various general purpose GPU (GPGPU) functions. The GPU may be communicatively coupled to the host processor/cores over a bus or another interconnect (e.g., a high-speed interconnect such as PCIe or NVLink). In other embodiments, the GPU may be integrated on the same package or chip as the cores and communicatively coupled to the cores over an internal processor bus/interconnect (i.e., internal to the package or chip). Regardless of the manner in which the GPU is connected, the processor cores may allocate work to the GPU in the form of sequences of commands/instructions contained in a work descriptor. The GPU then uses dedicated circuitry/logic for efficiently processing these commands/instructions.

[0044] In some embodiments, an image capturing device is a standalone device to capture an input image. The image capturing device, however, can be part of or a subcomponent of another computing device requiring image capturing capabilities such as a portable or hand-held computing device with a digital camera to capture images.

[0045] In the following description, numerous specific details are set forth to provide a more thorough understanding. However, it will be apparent that embodiments described herein may be practiced without one or more of these specific details. In other instances, well-known features have not been described to avoid obscuring the details of the exemplary embodiments.

Computing System Overview

[0046] FIG. 1 is a block diagram illustrating a computing system 100 configured to implement one or more aspects of exemplary embodiments described herein. The computing system 100 includes a processing subsystem 101 having one or more processor(s) 102 and a system memory 104 communicating via an interconnection path that may include a memory hub 105. The memory hub 105 may be a separate component within a chipset component or may be integrated within the one or more processor(s) 102. The memory hub 105 couples with an I/O subsystem 111 via a communication link 106.

The I/O subsystem 111 includes an I/O hub 107 that can enable the computing system 100 to receive input from one or more input device(s) 108. Additionally, the I/O hub 107 can enable a display controller, which may be included in the one or more processor(s) 102, to provide outputs to one or more display device(s) 110A. In one embodiment, the one or more display device(s) 110A coupled with the I/O hub 107 can include a local, internal, or embedded display device.

[0047] In one embodiment, the processing subsystem 101 includes one or more parallel processor(s) 112 coupled to memory hub 105 via a bus or other communication link 113. The communication link 113 may be one of any number of standards based communication link technologies or protocols, such as, but not limited to PCI Express, or may be a vendor specific communications interface or communications fabric. In one embodiment, the one or more parallel processor(s) 112 form a computationally focused parallel or vector processing system that include a large number of processing cores and/or processing clusters, such as a many integrated core (MIC) processor. In one embodiment, the one or more parallel processor(s) 112 form a graphics processing subsystem that can output pixels to one of the one or more display device(s) 110A coupled via the I/O Hub 107. The one or more parallel processor(s) 112 can also include a display controller and display interface (not shown) to enable a direct connection to one or more display device(s) 110B.

[0048] Within the I/O subsystem 111, a system storage unit 114 can connect to the I/O hub 107 to provide a storage mechanism for the computing system 100. An I/O switch 116 can be used to provide an interface mechanism to enable connections between the I/O hub 107 and other components, such as a network adapter 118 and/or wireless network adapter 119 that may be integrated into the platform, and various other devices that can be added via one or more add-in device(s) 120. The network adapter 118 can be an Ethernet adapter or another wired network adapter. The wireless network adapter 119 can include one or more of a Wi-Fi, Bluetooth, near field communication (NFC), or other network device that includes one or more wireless radios.

[0049] The computing system 100 can include other components not explicitly shown, including USB or other port connections, optical storage drives, video capture

devices, and the like, may also be connected to the I/O hub 107. Communication paths interconnecting the various components in FIG. 1 may be implemented using any suitable protocols, such as PCI (Peripheral Component Interconnect) based protocols (e.g., PCI-Express), or any other bus or point-to-point communication interfaces and/or protocol(s),
5 such as the NV-Link high-speed interconnect, or interconnect protocols known in the art.

[0050] In one embodiment, the one or more parallel processor(s) 112 incorporate circuitry optimized for graphics and video processing, including, for example, video output circuitry, and constitutes a graphics processing unit (GPU). In another embodiment, the one or more parallel processor(s) 112 incorporate circuitry optimized
10 for general purpose processing, while preserving the underlying computational architecture, described in greater detail herein. In yet another embodiment, components of the computing system 100 may be integrated with one or more other system elements on a single integrated circuit. For example, the one or more parallel processor(s), 112 memory hub 105, processor(s) 102, and I/O hub 107 can be integrated into a system on
15 chip (SoC) integrated circuit. Alternatively, the components of the computing system 100 can be integrated into a single package to form a system in package (SIP) configuration. In one embodiment, at least a portion of the components of the computing system 100 can be integrated into a multi-chip module (MCM), which can be interconnected with other multi-chip modules into a modular computing system.

[0051] It will be appreciated that the computing system 100 shown herein is illustrative and that variations and modifications are possible. The connection topology, including the number and arrangement of bridges, the number of processor(s) 102, and the number of parallel processor(s) 112, may be modified as desired. For instance, in
20 some embodiments, system memory 104 is connected to the processor(s) 102 directly rather than through a bridge, while other devices communicate with system memory 104 via the memory hub 105 and the processor(s) 102. In other alternative topologies, the parallel processor(s) 112 are connected to the I/O hub 107 or directly to one of the one or more processor(s) 102, rather than to the memory hub 105. In other embodiments, the I/O hub 107 and memory hub 105 may be integrated into a single chip. Some
25

embodiments may include two or more sets of processor(s) 102 attached via multiple sockets, which can couple with two or more instances of the parallel processor(s) 112.

[0052] Some of the particular components shown herein are optional and may not be included in all implementations of the computing system 100. For example, any number of add-in cards or peripherals may be supported, or some components may be eliminated. Furthermore, some architectures may use different terminology for components similar to those illustrated in FIG. 1. For example, the memory hub 105 may be referred to as a Northbridge in some architectures, while the I/O hub 107 may be referred to as a Southbridge.

10 **[0053]** **FIG. 2A** illustrates a parallel processor 200 according to an exemplary embodiment. The various components of the parallel processor 200 may be implemented using one or more integrated circuit devices, such as programmable processors, application specific integrated circuits (ASICs), or field programmable gate arrays (FPGA). The illustrated parallel processor 200 is a variant of the one or more parallel processor(s) 112 shown in **FIG. 1** according to an exemplary embodiment.

15 **[0054]** In one embodiment, the parallel processor 200 includes a parallel processing unit 202. The parallel processing unit includes an I/O unit 204 that enables communication with other devices, including other instances of the parallel processing unit 202. The I/O unit 204 may be directly connected to other devices. In one embodiment, the I/O unit 204 connects with other devices via the use of a hub or switch interface, such as memory hub 105. The connections between the memory hub 105 and the I/O unit 204 form a communication link 113. Within the parallel processing unit 202, the I/O unit 204 connects with a host interface 206 and a memory crossbar 216, where the host interface 206 receives commands directed to performing processing operations and the memory crossbar 216 receives commands directed to performing memory operations.

20 **[0055]** When the host interface 206 receives a command buffer via the I/O unit 204, the host interface 206 can direct work operations to perform those commands to a front end 208. In one embodiment, the front end 208 couples with a scheduler 210, which is configured to distribute commands or other work items to a processing cluster array 212. In one embodiment, the scheduler 210 ensures that the processing cluster

25
30

array 212 is properly configured and in a valid state before tasks are distributed to the processing clusters of the processing cluster array 212. In one embodiment, the scheduler 210 is implemented via firmware logic executing on a microcontroller. The microcontroller implemented scheduler 210 is configurable to perform complex scheduling and work distribution operations at coarse and fine granularity, enabling rapid preemption and context switching of threads executing on the processing array 212. In one embodiment, the host software can provide workloads for scheduling on the processing array 212 via one of multiple graphics processing doorbells. The workloads can then be automatically distributed across the processing array 212 by the scheduler 210 logic within the scheduler microcontroller.

[0056] The processing cluster array 212 can include up to “N” processing clusters (e.g., cluster 214A, cluster 214B, through cluster 214N). Each cluster 214A-214N of the processing cluster array 212 can execute a large number of concurrent threads. The scheduler 210 can allocate work to the clusters 214A-214N of the processing cluster array 212 using various scheduling and/or work distribution algorithms, which may vary depending on the workload arising for each type of program or computation. The scheduling can be handled dynamically by the scheduler 210, or can be assisted in part by compiler logic during compilation of program logic configured for execution by the processing cluster array 212. In one embodiment, different clusters 214A-214N of the processing cluster array 212 can be allocated for processing different types of programs or for performing different types of computations.

[0057] The processing cluster array 212 can be configured to perform various types of parallel processing operations. In one embodiment, the processing cluster array 212 is configured to perform general-purpose parallel compute operations. For example, the processing cluster array 212 can include logic to execute processing tasks including filtering of video and/or audio data, performing modeling operations, including physics operations, and performing data transformations.

[0058] In one embodiment, the processing cluster array 212 is configured to perform parallel graphics processing operations. In embodiments in which the parallel processor 200 is configured to perform graphics processing operations, the processing

cluster array 212 can include additional logic to support the execution of such graphics processing operations, including, but not limited to texture sampling logic to perform texture operations, as well as tessellation logic and other vertex processing logic. Additionally, the processing cluster array 212 can be configured to execute graphics processing related shader programs such as, but not limited to vertex shaders, tessellation
5 shaders, geometry shaders, and pixel shaders. The parallel processing unit 202 can transfer data from system memory via the I/O unit 204 for processing. During processing the transferred data can be stored to on-chip memory (e.g., parallel processor memory 222) during processing, then written back to system memory.

10 **[0059]** In one embodiment, when the parallel processing unit 202 is used to perform graphics processing, the scheduler 210 can be configured to divide the processing workload into approximately equal sized tasks, to better enable distribution of the graphics processing operations to multiple clusters 214A-214N of the processing cluster array 212. In some embodiments, portions of the processing cluster array 212 can
15 be configured to perform different types of processing. For example, a first portion may be configured to perform vertex shading and topology generation, a second portion may be configured to perform tessellation and geometry shading, and a third portion may be configured to perform pixel shading or other screen space operations, to produce a rendered image for display. Intermediate data produced by one or more of the clusters
20 214A-214N may be stored in buffers to allow the intermediate data to be transmitted between clusters 214A-214N for further processing.

[0060] During operation, the processing cluster array 212 can receive processing tasks to be executed via the scheduler 210, which receives commands defining processing tasks from front end 208. For graphics processing operations, processing tasks can
25 include indices of data to be processed, e.g., surface (patch) data, primitive data, vertex data, and/or pixel data, as well as state parameters and commands defining how the data is to be processed (e.g., what program is to be executed). The scheduler 210 may be configured to fetch the indices corresponding to the tasks or may receive the indices from the front end 208. The front end 208 can be configured to ensure the processing cluster

array 212 is configured to a valid state before the workload specified by incoming command buffers (e.g., batch-buffers, push buffers, etc.) is initiated.

[0061] Each of the one or more instances of the parallel processing unit 202 can couple with parallel processor memory 222. The parallel processor memory 222 can be accessed via the memory crossbar 216, which can receive memory requests from the processing cluster array 212 as well as the I/O unit 204. The memory crossbar 216 can access the parallel processor memory 222 via a memory interface 218. The memory interface 218 can include multiple partition units (e.g., partition unit 220A, partition unit 220B, through partition unit 220N) that can each couple to a portion (e.g., memory unit) of parallel processor memory 222. In one implementation, the number of partition units 220A-220N is configured to be equal to the number of memory units, such that a first partition unit 220A has a corresponding first memory unit 224A, a second partition unit 220B has a corresponding memory unit 224B, and an Nth partition unit 220N has a corresponding Nth memory unit 224N. In other embodiments, the number of partition units 220A-220N may not be equal to the number of memory devices.

[0062] In various embodiments, the memory units 224A-224N can include various types of memory devices, including dynamic random access memory (DRAM) or graphics random access memory, such as synchronous graphics random access memory (SGRAM), including graphics double data rate (GDDR) memory. In one embodiment, the memory units 224A-224N may also include 3D stacked memory, including but not limited to high bandwidth memory (HBM). Persons skilled in the art will appreciate that the specific implementation of the memory units 224A-224N can vary, and can be selected from one of various conventional designs. Render targets, such as frame buffers or texture maps may be stored across the memory units 224A-224N, allowing partition units 220A-220N to write portions of each render target in parallel to efficiently use the available bandwidth of parallel processor memory 222. In some embodiments, a local instance of the parallel processor memory 222 may be excluded in favor of a unified memory design that utilizes system memory in conjunction with local cache memory.

[0063] In one embodiment, any one of the clusters 214A-214N of the processing cluster array 212 can process data that will be written to any of the memory units 224A-

224N within parallel processor memory 222. The memory crossbar 216 can be configured to transfer the output of each cluster 214A-214N to any partition unit 220A-220N or to another cluster 214A-214N, which can perform additional processing operations on the output. Each cluster 214A-214N can communicate with the memory interface 218 through the memory crossbar 216 to read from or write to various external memory devices. In one embodiment, the memory crossbar 216 has a connection to the memory interface 218 to communicate with the I/O unit 204, as well as a connection to a local instance of the parallel processor memory 222, enabling the processing units within the different processing clusters 214A-214N to communicate with system memory or other memory that is not local to the parallel processing unit 202. In one embodiment, the memory crossbar 216 can use virtual channels to separate traffic streams between the clusters 214A-214N and the partition units 220A-220N.

[0064] While a single instance of the parallel processing unit 202 is illustrated within the parallel processor 200, any number of instances of the parallel processing unit 202 can be included. For example, multiple instances of the parallel processing unit 202 can be provided on a single add-in card, or multiple add-in cards can be interconnected. The different instances of the parallel processing unit 202 can be configured to inter-operate even if the different instances have different numbers of processing cores, different amounts of local parallel processor memory, and/or other configuration differences. For example, and in one embodiment, some instances of the parallel processing unit 202 can include higher precision floating point units relative to other instances. Systems incorporating one or more instances of the parallel processing unit 202 or the parallel processor 200 can be implemented in a variety of configurations and form factors, including but not limited to desktop, laptop, or handheld personal computers, servers, workstations, game consoles, and/or embedded systems.

[0065] **FIG. 2B** is a block diagram of a partition unit 220 according to an exemplary embodiment. In one embodiment, the partition unit 220 is an instance of one of the partition units 220A-220N of **FIG. 2A**. As illustrated, the partition unit 220 includes an L2 cache 221, a frame buffer interface 225, and a ROP 226 (raster operations unit). The L2 cache 221 is a read/write cache that is configured to perform load and store

operations received from the memory crossbar 216 and ROP 226. Read misses and urgent write-back requests are output by L2 cache 221 to frame buffer interface 225 for processing. Updates can also be sent to the frame buffer via the frame buffer interface 225 for processing. In one embodiment, the frame buffer interface 225 interfaces with
5 one of the memory units in parallel processor memory, such as the memory units 224A-224N of FIG. 2 (e.g., within parallel processor memory 222).

[0066] In graphics applications, the ROP 226 is a processing unit that performs raster operations such as stencil, z test, blending, and the like. The ROP 226 then outputs processed graphics data that is stored in graphics memory. In some embodiments, the
10 ROP 226 includes compression logic to compress depth or color data that is written to memory and decompress depth or color data that is read from memory. The compression logic can be lossless compression logic that makes use of one or more of multiple compression algorithms. The type of compression that is performed by the ROP 226 can vary based on the statistical characteristics of the data to be compressed. For example, in
15 one embodiment, delta color compression is performed on depth and color data on a per-tile basis.

[0067] In some embodiments, the ROP 226 is included within each processing cluster (e.g., cluster 214A-214N of **FIG. 2A**) instead of within the partition unit 220. In such embodiment, read and write requests for pixel data are transmitted over the memory
20 crossbar 216 instead of pixel fragment data. The processed graphics data may be displayed on a display device, such as one of the one or more display device(s) 110 of FIG. 1, routed for further processing by the processor(s) 102, or routed for further processing by one of the processing entities within the parallel processor 200 of **FIG. 2A**.

[0068] **FIG. 2C** is a block diagram of a processing cluster 214 within a parallel
25 processing unit according to an exemplary embodiment. In one embodiment, the processing cluster is an instance of one of the processing clusters 214A-214N of **FIG. 2A**. The processing cluster 214 can be configured to execute many threads in parallel, where the term “thread” refers to an instance of a particular program executing on a particular set of input data. In some embodiments, single-instruction, multiple-data (SIMD)
30 instruction issue techniques are used to support parallel execution of a large number of

threads without providing multiple independent instruction units. In other embodiments, single-instruction, multiple-thread (SIMT) techniques are used to support parallel execution of a large number of generally synchronized threads, using a common instruction unit configured to issue instructions to a set of processing engines within each
5 one of the processing clusters. Unlike a SIMD execution regime, where all processing engines typically execute identical instructions, SIMT execution allows different threads to more readily follow divergent execution paths through a given thread program. Persons skilled in the art will understand that a SIMD processing regime represents a functional subset of a SIMT processing regime.

10 **[0069]** Operation of the processing cluster 214 can be controlled via a pipeline manager 232 that distributes processing tasks to SIMT parallel processors. The pipeline manager 232 receives instructions from the scheduler 210 of **FIG. 2A** and manages execution of those instructions via a graphics multiprocessor 234 and/or a texture unit 236. The illustrated graphics multiprocessor 234 is an exemplary instance of a SIMT
15 parallel processor. However, various types of SIMT parallel processors of differing architectures may be included within the processing cluster 214. One or more instances of the graphics multiprocessor 234 can be included within a processing cluster 214. The graphics multiprocessor 234 can process data and a data crossbar 240 can be used to distribute the processed data to one of multiple possible destinations, including other
20 shader units. The pipeline manager 232 can facilitate the distribution of processed data by specifying destinations for processed data to be distributed via the data crossbar 240.

[0070] Each graphics multiprocessor 234 within the processing cluster 214 can include an identical set of functional execution logic (e.g., arithmetic logic units, load-store units, etc.). The functional execution logic can be configured in a pipelined manner
25 in which new instructions can be issued before previous instructions are complete. The functional execution logic supports a variety of operations including integer and floating point arithmetic, comparison operations, Boolean operations, bit-shifting, and computation of various algebraic functions. In one embodiment, the same functional-unit hardware can be leveraged to perform different operations and any combination of
30 functional units may be present.

[0071] The instructions transmitted to the processing cluster 214 constitutes a thread. A set of threads executing across the set of parallel processing engines is a thread group. A thread group executes the same program on different input data. Each thread within a thread group can be assigned to a different processing engine within a graphics multiprocessor 234. A thread group may include fewer threads than the number of processing engines within the graphics multiprocessor 234. When a thread group includes fewer threads than the number of processing engines, one or more of the processing engines may be idle during cycles in which that thread group is being processed. A thread group may also include more threads than the number of processing engines within the graphics multiprocessor 234. When the thread group includes more threads than the number of processing engines within the graphics multiprocessor 234, processing can be performed over consecutive clock cycles. In one embodiment, multiple thread groups can be executed concurrently on a graphics multiprocessor 234.

[0072] In one embodiment, the graphics multiprocessor 234 includes an internal cache memory to perform load and store operations. In one embodiment, the graphics multiprocessor 234 can forego an internal cache and use a cache memory (e.g., L1 cache 308) within the processing cluster 214. Each graphics multiprocessor 234 also has access to L2 caches within the partition units (e.g., partition units 220A-220N of FIG. 2A) that are shared among all processing clusters 214 and may be used to transfer data between threads. The graphics multiprocessor 234 may also access off-chip global memory, which can include one or more of local parallel processor memory and/or system memory. Any memory external to the parallel processing unit 202 may be used as global memory. Embodiments in which the processing cluster 214 includes multiple instances of the graphics multiprocessor 234 can share common instructions and data, which may be stored in the L1 cache 308.

[0073] Each processing cluster 214 may include an MMU 245 (memory management unit) that is configured to map virtual addresses into physical addresses. In other embodiments, one or more instances of the MMU 245 may reside within the memory interface 218 of FIG. 2A. The MMU 245 includes a set of page table entries (PTEs) used to map a virtual address to a physical address of a tile (talk more about tiling)

and optionally a cache line index. The MMU 245 may include address translation lookaside buffers (TLB) or caches that may reside within the graphics multiprocessor 234 or the L1 cache or processing cluster 214. The physical address is processed to distribute surface data access locality to allow efficient request interleaving among partition units.

- 5 The cache line index may be used to determine whether a request for a cache line is a hit or miss.

[0074] In graphics and computing applications, a processing cluster 214 may be configured such that each graphics multiprocessor 234 is coupled to a texture unit 236 for performing texture mapping operations, e.g., determining texture sample positions, reading texture data, and filtering the texture data. Texture data is read from an internal texture L1 cache (not shown) or in some embodiments from the L1 cache within graphics multiprocessor 234 and is fetched from an L2 cache, local parallel processor memory, or system memory, as needed. Each graphics multiprocessor 234 outputs processed tasks to the data crossbar 240 to provide the processed task to another processing cluster 214 for further processing or to store the processed task in an L2 cache, local parallel processor memory, or system memory via the memory crossbar 216. A preROP 242 (pre-raster operations unit) is configured to receive data from graphics multiprocessor 234, direct data to ROP units, which may be located with partition units as described herein (e.g., partition units 220A-220N of FIG. 2A). The preROP 242 unit can perform optimizations for color blending, organize pixel color data, and perform address translations.

10
15
20

[0075] It will be appreciated that the core architecture described herein is illustrative and that variations and modifications are possible. Any number of processing units, e.g., graphics multiprocessor 234, texture units 236, preROPs 242, etc., may be included within a processing cluster 214. Further, while only one processing cluster 214 is shown, a parallel processing unit as described herein may include any number of instances of the processing cluster 214. In one embodiment, each processing cluster 214 can be configured to operate independently of other processing clusters 214 using separate and distinct processing units, L1 caches, etc.

25

[0076] FIG. 2D shows a graphics multiprocessor 234 according to one exemplary embodiment. In such embodiment, the graphics multiprocessor 234 couples with the

30

pipeline manager 232 of the processing cluster 214. The graphics multiprocessor 234 has an execution pipeline including but not limited to an instruction cache 252, an instruction unit 254, an address mapping unit 256, a register file 258, one or more general purpose graphics processing unit (GPGPU) cores 262, and one or more load/store units 266. The
5 GPGPU cores 262 and load/store units 266 are coupled with cache memory 272 and shared memory 270 via a memory and cache interconnect 268.

[0077] In one embodiment, the instruction cache 252 receives a stream of instructions to execute from the pipeline manager 232. The instructions are cached in the instruction cache 252 and dispatched for execution by the instruction unit 254. The
10 instruction unit 254 can dispatch instructions as thread groups (e.g., warps), with each thread of the thread group assigned to a different execution unit within GPGPU core 262. An instruction can access any of a local, shared, or global address space by specifying an address within a unified address space. The address mapping unit 256 can be used to translate addresses in the unified address space into a distinct memory address that can be
15 accessed by the load/store units 266.

[0078] The register file 258 provides a set of registers for the functional units of the graphics multiprocessor 324. The register file 258 provides temporary storage for operands connected to the data paths of the functional units (e.g., GPGPU cores 262, load/store units 266) of the graphics multiprocessor 324. In one embodiment, the register
20 file 258 is divided between each of the functional units such that each functional unit is allocated a dedicated portion of the register file 258. In one embodiment, the register file 258 is divided between the different warps being executed by the graphics multiprocessor 324.

[0079] The GPGPU cores 262 can each include floating point units (FPUs) and/or
25 integer arithmetic logic units (ALUs) that are used to execute instructions of the graphics multiprocessor 324. The GPGPU cores 262 can be similar in architecture or can differ in architecture, according to embodiments. For example, and in one embodiment, a first portion of the GPGPU cores 262 include a single precision FPU and an integer ALU while a second portion of the GPGPU cores include a double precision FPU. In one
30 embodiment, the FPUs can implement the IEEE 754-2008 standard for floating point

arithmetic or enable variable precision floating point arithmetic. The graphics multiprocessor 324 can additionally include one or more fixed function or special function units to perform specific functions such as copy rectangle or pixel blending operations. In one embodiment one or more of the GPGPU cores can also include fixed or special function logic.

[0080] In one embodiment, the GPGPU cores 262 include SIMD logic capable of performing a single instruction on multiple sets of data. In one embodiment GPGPU cores 262 can physically execute SIMD4, SIMD8, and SIMD16 instructions and logically execute SIMD1, SIMD2, and SIMD32 instructions. The SIMD instructions for the GPGPU cores can be generated at compile time by a shader compiler or automatically generated when executing programs written and compiled for single program multiple data (SPMD) or SIMT architectures. Multiple threads of a program configured for the SIMT execution model can be executed via a single SIMD instruction. For example, and in one embodiment, eight SIMT threads that perform the same or similar operations can be executed in parallel via a single SIMD8 logic unit.

[0081] The memory and cache interconnect 268 is an interconnect network that connects each of the functional units of the graphics multiprocessor 324 to the register file 258 and to the shared memory 270. In one embodiment, the memory and cache interconnect 268 is a crossbar interconnect that allows the load/store unit 266 to implement load and store operations between the shared memory 270 and the register file 258. The register file 258 can operate at the same frequency as the GPGPU cores 262, thus data transfer between the GPGPU cores 262 and the register file 258 is very low latency. The shared memory 270 can be used to enable communication between threads that execute on the functional units within the graphics multiprocessor 234. The cache memory 272 can be used as a data cache for example, to cache texture data communicated between the functional units and the texture unit 236. The shared memory 270 can also be used as a program managed cache. Threads executing on the GPGPU cores 262 can programmatically store data within the shared memory in addition to the automatically cached data that is stored within the cache memory 272.

[0082] FIGs. 3A-3B illustrate additional graphics multiprocessors according to exemplary embodiments. The illustrated graphics multiprocessors 325, 350 are variants of the graphics multiprocessor 234 of FIGS. 2C-2D. The illustrated graphics multiprocessors 325, 350 can be configured as a streaming multiprocessor (SM) capable of simultaneous execution of a large number of execution threads.

[0083] FIG. 3A shows a graphics multiprocessor 325 according to an additional exemplary embodiment. The graphics multiprocessor 325 includes multiple additional instances of execution resource units relative to the graphics multiprocessor 234 of FIGS. 2C-2D. For example, the graphics multiprocessor 325 can include multiple instances of the instruction unit 332A-332B, register file 334A-334B, and texture unit(s) 344A-344B. The graphics multiprocessor 325 also includes multiple sets of graphics or compute execution units (e.g., GPGPU core 336A-336B, GPGPU core 337A-337B, GPGPU core 338A-338B) and multiple sets of load/store units 340A-340B. In one embodiment, the execution resource units have a common instruction cache 330, texture and/or data cache memory 342, and shared memory 346.

[0084] The various components can communicate via an interconnect fabric 327. In one embodiment, the interconnect fabric 327 includes one or more crossbar switches to enable communication between the various components of the graphics multiprocessor 325. In one embodiment, the interconnect fabric 327 is a separate, high-speed network fabric layer upon which each component of the graphics multiprocessor 325 is stacked. The components of the graphics multiprocessor 325 communicate with remote components via the interconnect fabric 327. For example, the GPGPU cores 336A-336B, 337A-337B, and 3378A-338B can each communicate with shared memory 346 via the interconnect fabric 327. The interconnect fabric 327 can arbitrate communication within the graphics multiprocessor 325 to ensure a fair bandwidth allocation between components.

[0085] FIG. 3B shows a graphics multiprocessor 350 according to an additional exemplary embodiment. The graphics processor includes multiple sets of execution resources 356A-356D, where each set of execution resource includes multiple instruction units, register files, GPGPU cores, and load store units, as illustrated in FIG. 2D and FIG.

3A. The execution resources 356A-356D can work in concert with texture unit(s) 360A-360D for texture operations, while sharing an instruction cache 354, and shared memory 362. In one embodiment, the execution resources 356A-356D can share an instruction cache 354 and shared memory 362, as well as multiple instances of a texture and/or data
5 cache memory 358A-358B. The various components can communicate via an interconnect fabric 352 similar to the interconnect fabric 327 of FIG. 3A.

[0086] Persons skilled in the art will understand that the architecture described in FIGS. 1, 2A-2D, and 3A-3B are descriptive and not limiting as to the scope of the present embodiments, which are exemplary. Thus, the techniques described herein may
10 be implemented on any properly configured processing unit, including, without limitation, one or more mobile application processors, one or more desktop or server central processing units (CPUs) including multi-core CPUs, one or more parallel processing units, such as the parallel processing unit 202 of FIG. 2A, as well as one or more graphics processors or special purpose processing units, without departure from the scope
15 of the embodiments described herein.

[0087] In some embodiments, a parallel processor or GPGPU as described herein is communicatively coupled to host/processor cores to accelerate graphics operations, machine-learning operations, pattern analysis operations, and various general purpose GPU (GPGPU) functions. The GPU may be communicatively coupled to the host
20 processor/cores over a bus or other interconnect (e.g., a high speed interconnect such as PCIe or NVLink). In other embodiments, the GPU may be integrated on the same package or chip as the cores and communicatively coupled to the cores over an internal processor bus/interconnect (i.e., internal to the package or chip). Regardless of the manner in which the GPU is connected, the processor cores may allocate work to the
25 GPU in the form of sequences of commands/instructions contained in a work descriptor. The GPU then uses dedicated circuitry/logic for efficiently processing these commands/instructions.

Techniques for GPU to Host Processor Interconnection

[0088] FIG. 4A illustrates an exemplary architecture in which a plurality of GPUs
30 410-413 are communicatively coupled to a plurality of multi-core processors 405-406

over high-speed links 440-443 (e.g., buses, point-to-point interconnects, etc.). In one embodiment, the high-speed links 440-443 support a communication throughput of 4GB/s, 30GB/s, 80GB/s or higher, depending on the implementation. Various interconnect protocols may be used including, but not limited to, PCIe 4.0 or 5.0 and
5 NVLink 2.0. However, the underlying principles of the invention are not limited to any particular communication protocol or throughput.

[0089] In addition, and in one embodiment, two or more of the GPUs 410-413 are interconnected over high-speed links 444-445, which may be implemented using the same or different protocols/links than those used for high-speed links 440-443. Similarly,
10 two or more of the multi-core processors 405-406 may be connected over high speed link 433 which may be symmetric multi-processor (SMP) buses operating at 20GB/s, 30GB/s, 120GB/s or higher. Alternatively, all communication between the various system components shown in **FIG. 4A** may be accomplished using the same protocols/links (e.g., over a common interconnection fabric). As mentioned, however, the underlying
15 principles of the invention are not limited to any particular type of interconnect technology.

[0090] In one embodiment, each multi-core processor 405-406 is communicatively coupled to a processor memory 401-402, via memory interconnects 430-431, respectively, and each GPU 410-413 is communicatively coupled to GPU
20 memory 420-423 over GPU memory interconnects 450-453, respectively. The memory interconnects 430-431 and 450-453 may utilize the same or different memory access technologies. By way of example, and not limitation, the processor memories 401-402 and GPU memories 420-423 may be volatile memories such as dynamic random access
25 memories (DRAMs) (including stacked DRAMs), Graphics DDR SDRAM (GDDR) (e.g., GDDR5, GDDR6), or High Bandwidth Memory (HBM) and/or may be non-volatile memories such as 3D XPoint or Nano-Ram. In one embodiment, some portion of the memories may be volatile memory and another portion may be non-volatile memory (e.g., using a two-level memory (2LM) hierarchy).

[0091] As described below, although the various processors 405-406 and GPUs
30 410-413 may be physically coupled to a particular memory 401-402, 420-423,

respectively, a unified memory architecture may be implemented in which the same virtual system address space (also referred to as the “effective address” space) is distributed among all of the various physical memories. For example, processor memories 401-402 may each comprise 64GB of the system memory address space and
5 GPU memories 420-423 may each comprise 32GB of the system memory address space (resulting in a total of 256GB addressable memory in this example).

[0092] FIG. 4B illustrates additional details for an interconnection between a multi-core processor 407 and a graphics acceleration module 446 in accordance with one exemplary embodiment. The graphics acceleration module 446 may include one or more
10 GPU chips integrated on a line card which is coupled to the processor 407 via the high-speed link 440. Alternatively, the graphics acceleration module 446 may be integrated on the same package or chip as the processor 407.

[0093] The illustrated processor 407 includes a plurality of cores 460A-460D, each with a translation lookaside buffer 461A-461D and one or more caches 462A-462D.
15 The cores may include various other components for executing instructions and processing data which are not illustrated to avoid obscuring the underlying principles of the invention (e.g., instruction fetch units, branch prediction units, decoders, execution units, reorder buffers, etc.). The caches 462A-462D may comprise level 1 (L1) and level 2 (L2) caches. In addition, one or more shared caches 426 may be included in the
20 caching hierarchy and shared by sets of the cores 460A-460D. For example, one embodiment of the processor 407 includes 24 cores, each with its own L1 cache, twelve shared L2 caches, and twelve shared L3 caches. In this embodiment, one of the L2 and L3 caches are shared by two adjacent cores. The processor 407 and the graphics accelerator integration module 446 connect with system memory 441, which may include
25 processor memories 401-402

[0094] Coherency is maintained for data and instructions stored in the various caches 462A-462D, 456 and system memory 441 via inter-core communication over a coherence bus 464. For example, each cache may have cache coherency logic/circuitry associated therewith to communicate to over the coherence bus 464 in response to
30 detected reads or writes to particular cache lines. In one implementation, a cache

snooping protocol is implemented over the coherence bus 464 to snoop cache accesses. Cache snooping/coherency techniques are well understood by those of skill in the art and will not be described in detail here to avoid obscuring the underlying principles of the invention.

5 [0095] In one embodiment, a proxy circuit 425 communicatively couples the graphics acceleration module 446 to the coherence bus 464, allowing the graphics acceleration module 446 to participate in the cache coherence protocol as a peer of the cores. In particular, an interface 435 provides connectivity to the proxy circuit 425 over high-speed link 440 (e.g., a PCIe bus, NVLink, etc.) and an interface 437 connects the
10 graphics acceleration module 446 to the link 440.

[0096] In one implementation, an accelerator integration circuit 436 provides cache management, memory access, context management, and interrupt management services on behalf of a plurality of graphics processing engines 431, 432, N of the graphics acceleration module 446. The graphics processing engines 431, 432, N may
15 each comprise a separate graphics processing unit (GPU). Alternatively, the graphics processing engines 431, 432, N may comprise different types of graphics processing engines within a GPU such as graphics execution units, media processing engines (e.g., video encoders/decoders), samplers, and blit engines. In other words, the graphics acceleration module may be a GPU with a plurality of graphics processing engines 431-
20 432, N or the graphics processing engines 431-432, N may be individual GPUs integrated on a common package, line card, or chip.

[0097] In one embodiment, the accelerator integration circuit 436 includes a memory management unit (MMU) 439 for performing various memory management functions such as virtual-to-physical memory translations (also referred to as effective-to-
25 real memory translations) and memory access protocols for accessing system memory 441. The MMU 439 may also include a translation lookaside buffer (TLB) (not shown) for caching the virtual/effective to physical/real address translations. In one implementation, a cache 438 stores commands and data for efficient access by the graphics processing engines 431-432, N. In one embodiment, the data stored in cache
30 438 and graphics memories 433-434, N is kept coherent with the core caches 462A-462D,

456 and system memory 411. As mentioned, this may be accomplished via proxy circuit 425 which takes part in the cache coherency mechanism on behalf of cache 438 and memories 433-434, N (e.g., sending updates to the cache 438 related to modifications/accesses of cache lines on processor caches 462A-462D, 456 and receiving updates from the cache 438).

[0098] A set of registers 445 store context data for threads executed by the graphics processing engines 431-432, N and a context management circuit 448 manages the thread contexts. For example, the context management circuit 448 may perform save and restore operations to save and restore contexts of the various threads during contexts switches (e.g., where a first thread is saved and a second thread is stored so that the second thread can be execute by a graphics processing engine). For example, on a context switch, the context management circuit 448 may store current register values to a designated region in memory (e.g., identified by a context pointer). It may then restore the register values when returning to the context. In one embodiment, an interrupt management circuit 447 receives and processes interrupts received from system devices.

[0099] In one implementation, virtual/effective addresses from a graphics processing engine 431 are translated to real/physical addresses in system memory 411 by the MMU 439. One embodiment of the accelerator integration circuit 436 supports multiple (e.g., 4, 8, 16) graphics accelerator modules 446 and/or other accelerator devices. The graphics accelerator module 446 may be dedicated to a single application executed on the processor 407 or may be shared between multiple applications. In one embodiment, a virtualized graphics execution environment is presented in which the resources of the graphics processing engines 431-432, N are shared with multiple applications or virtual machines (VMs). The resources may be subdivided into “slices” which are allocated to different VMs and/or applications based on the processing requirements and priorities associated with the VMs and/or applications.

[00100] Thus, the accelerator integration circuit acts as a bridge to the system for the graphics acceleration module 446 and provides address translation and system memory cache services. In addition, the accelerator integration circuit 436 may provide

virtualization facilities for the host processor to manage virtualization of the graphics processing engines, interrupts, and memory management.

[00101] Because hardware resources of the graphics processing engines 431-432, N are mapped explicitly to the real address space seen by the host processor 407, any host processor can address these resources directly using an effective address value. One function of the accelerator integration circuit 436, in one embodiment, is the physical separation of the graphics processing engines 431-432, N so that they appear to the system as independent units.

[00102] As mentioned, in the illustrated embodiment, one or more graphics memories 433-434, M are coupled to each of the graphics processing engines 431-432, N, respectively. The graphics memories 433-434, M store instructions and data being processed by each of the graphics processing engines 431-432, N. The graphics memories 433-434, M may be volatile memories such as DRAMs (including stacked DRAMs), GDDR memory (e.g., GDDR5, GDDR6), or HBM, and/or may be non-volatile memories such as 3D XPoint or Nano-Ram.

[00103] In one embodiment, to reduce data traffic over link 440, biasing techniques are used to ensure that the data stored in graphics memories 433-434, M is data which will be used most frequently by the graphics processing engines 431-432, N and preferably not used by the cores 460A-460D (at least not frequently). Similarly, the biasing mechanism attempts to keep data needed by the cores (and preferably not the graphics processing engines 431-432, N) within the caches 462A-462D, 456 of the cores and system memory 411.

[00104] FIG. 4C illustrates another exemplary embodiment in which the accelerator integration circuit 436 is integrated within the processor 407. In this embodiment, the graphics processing engines 431-432, N communicate directly over the high-speed link 440 to the accelerator integration circuit 436 via interface 437 and interface 435 (which, again, may be utilize any form of bus or interface protocol). The accelerator integration circuit 436 may perform the same operations as those described with respect to FIG. 4B, but potentially at a higher throughput given its close proximity to the coherency bus 462 and caches 462A-462D, 426.

[00105] One embodiment supports different programming models including a dedicated-process programming model (no graphics acceleration module virtualization) and shared programming models (with virtualization). The latter may include programming models which are controlled by the accelerator integration circuit 436 and
5 programming models which are controlled by the graphics acceleration module 446.

[00106] In one embodiment of the dedicated process model, graphics processing engines 431-432, N are dedicated to a single application or process under a single operating system. The single application can funnel other application requests to the graphics engines 431-432, N, providing virtualization within a VM/partition.

10 [00107] In the dedicated-process programming models, the graphics processing engines 431-432, N, may be shared by multiple VM/application partitions. The shared models require a system hypervisor to virtualize the graphics processing engines 431-432, N to allow access by each operating system. For single-partition systems without a hypervisor, the graphics processing engines 431-432, N are owned by the operating
15 system. In both cases, the operating system can virtualize the graphics processing engines 431-432, N to provide access to each process or application.

[00108] For the shared programming model, the graphics acceleration module 446 or an individual graphics processing engine 431-432, N selects a process element using a process handle. In one embodiment, process elements are stored in system memory 411
20 and are addressable using the effective address to real address translation techniques described herein. The process handle may be an implementation-specific value provided to the host process when registering its context with the graphics processing engine 431-432, N (that is, calling system software to add the process element to the process element linked list). The lower 16-bits of the process handle may be the offset of the process
25 element within the process element linked list.

[00109] FIG. 4D illustrates an exemplary accelerator integration slice 490. As used herein, a "slice" comprises a specified portion of the processing resources of the accelerator integration circuit 436. Application effective address space 482 within system memory 411 stores process elements 483. In one embodiment, the process
30 elements 483 are stored in response to GPU invocations 481 from applications 480

executed on the processor 407. A process element 483 contains the process state for the corresponding application 480. A work descriptor (WD) 484 contained in the process element 483 can be a single job requested by an application or may contain a pointer to a queue of jobs. In the latter case, the WD 484 is a pointer to the job request queue in the application's address space 482.

5
[00110] The graphics acceleration module 446 and/or the individual graphics processing engines 431-432, N can be shared by all or a subset of the processes in the system. Embodiments of the invention include an infrastructure for setting up the process state and sending a WD 484 to a graphics acceleration module 446 to start a job in a virtualized environment.

10
[00111] In one implementation, the dedicated-process programming model is implementation-specific. In this model, a single process owns the graphics acceleration module 446 or an individual graphics processing engine 431. Because the graphics acceleration module 446 is owned by a single process, the hypervisor initializes the accelerator integration circuit 436 for the owning partition and the operating system initializes the accelerator integration circuit 436 for the owning process at the time when the graphics acceleration module 446 is assigned.

15
[00112] In operation, a WD fetch unit 491 in the accelerator integration slice 490 fetches the next WD 484 which includes an indication of the work to be done by one of the graphics processing engines of the graphics acceleration module 446. Data from the WD 484 may be stored in registers 445 and used by the MMU 439, interrupt management circuit 447 and/or context management circuit 446 as illustrated. For example, one embodiment of the MMU 439 includes segment/page walk circuitry for accessing segment/page tables 486 within the OS virtual address space 485. The interrupt management circuit 447 may process interrupt events 492 received from the graphics acceleration module 446. When performing graphics operations, an effective address 493 generated by a graphics processing engine 431-432, N is translated to a real address by the MMU 439.

20
[00113] In one embodiment, the same set of registers 445 are duplicated for each graphics processing engine 431-432, N and/or graphics acceleration module 446 and may
25
30

be initialized by the hypervisor or operating system. Each of these duplicated registers may be included in an accelerator integration slice 490. Exemplary registers that may be initialized by the hypervisor are shown in **Table 1**.

Table 1 – Hypervisor Initialized Registers

1	Slice Control Register
2	Real Address (RA) Scheduled Processes Area Pointer
3	Authority Mask Override Register
4	Interrupt Vector Table Entry Offset
5	Interrupt Vector Table Entry Limit
6	State Register
7	Logical Partition ID
8	Real address (RA) Hypervisor Accelerator Utilization Record Pointer
9	Storage Description Register

5

[00114] Exemplary registers that may be initialized by the operating system are shown in **Table 2**.

Table 2 – Operating System Initialized Registers

1	Process and Thread Identification
2	Effective Address (EA) Context Save/Restore Pointer
3	Virtual Address (VA) Accelerator Utilization Record Pointer
4	Virtual Address (VA) Storage Segment Table Pointer
5	Authority Mask
6	Work descriptor

10 **[00115]** In one embodiment, each WD 484 is specific to a particular graphics acceleration module 446 and/or graphics processing engine 431-432, N. It contains all the information a graphics processing engine 431-432, N requires to do its work or it can

be a pointer to a memory location where the application has set up a command queue of work to be completed.

5 [00116] FIG. 4E illustrates additional details for one exemplary embodiment of a shared model. This embodiment includes a hypervisor real address space 498 in which a process element list 499 is stored. The hypervisor real address space 498 is accessible via a hypervisor 496 which virtualizes the graphics acceleration module engines for the operating system 495.

10 [00117] The shared programming models allow for all or a subset of processes from all or a subset of partitions in the system to use a graphics acceleration module 446. There are two programming models where the graphics acceleration module 446 is shared by multiple processes and partitions: time-sliced shared and graphics directed shared.

15 [00118] In this model, the system hypervisor 496 owns the graphics acceleration module 446 and makes its function available to all operating systems 495. For a graphics acceleration module 446 to support virtualization by the system hypervisor 496, the graphics acceleration module 446 may adhere to the following requirements: 1) An application's job request must be autonomous (that is, the state does not need to be maintained between jobs), or the graphics acceleration module 446 must provide a context save and restore mechanism. 2) An application's job request is guaranteed by the graphics acceleration module 446 to complete in a specified amount of time, including any translation faults, or the graphics acceleration module 446 provides the ability to preempt the processing of the job. 3) The graphics acceleration module 446 must be guaranteed fairness between processes when operating in the directed shared programming model.

25 [00119] In one embodiment, for the shared model, the application 480 is required to make an operating system 495 system call with a graphics acceleration module 446 type, a work descriptor (WD), an authority mask register (AMR) value, and a context save/restore area pointer (CSRP). The graphics acceleration module 446 type describes the targeted acceleration function for the system call. The graphics acceleration module
30 446 type may be a system-specific value. The WD is formatted specifically for the

graphics acceleration module 446 and can be in the form of a graphics acceleration module 446 command, an effective address pointer to a user-defined structure, an effective address pointer to a queue of commands, or any other data structure to describe the work to be done by the graphics acceleration module 446. In one embodiment, the AMR value is the AMR state to use for the current process. The value passed to the operating system is similar to an application setting the AMR. If the accelerator integration circuit 436 and graphics acceleration module 446 implementations do not support a User Authority Mask Override Register (UAMOR), the operating system may apply the current UAMOR value to the AMR value before passing the AMR in the hypervisor call. The hypervisor 496 may optionally apply the current Authority Mask Override Register (AMOR) value before placing the AMR into the process element 483. In one embodiment, the CSRP is one of the registers 445 containing the effective address of an area in the application's address space 482 for the graphics acceleration module 446 to save and restore the context state. This pointer is optional if no state is required to be saved between jobs or when a job is preempted. The context save/restore area may be pinned system memory.

[00120] Upon receiving the system call, the operating system 495 may verify that the application 480 has registered and been given the authority to use the graphics acceleration module 446. The operating system 495 then calls the hypervisor 496 with the information shown in **Table 3**.

Table 3 – OS to Hypervisor Call Parameters

1	A work descriptor (WD)
2	An Authority Mask Register (AMR) value (potentially masked).
3	An effective address (EA) Context Save/Restore Area Pointer (CSRP)
4	A process ID (PID) and optional thread ID (TID)
5	A virtual address (VA) accelerator utilization record pointer (AURP)

6	The virtual address of the storage segment table pointer (SSTP)
7	A logical interrupt service number (LISN)

[00121] Upon receiving the hypervisor call, the hypervisor 496 verifies that the operating system 495 has registered and been given the authority to use the graphics acceleration module 446. The hypervisor 496 then puts the process element 483 into the process element linked list for the corresponding graphics acceleration module 446 type. The process element may include the information shown in **Table 4**.

Table 4 – Process Element Information

1	A work descriptor (WD)
2	An Authority Mask Register (AMR) value (potentially masked).
3	An effective address (EA) Context Save/Restore Area Pointer (CSRP)
4	A process ID (PID) and optional thread ID (TID)
5	A virtual address (VA) accelerator utilization record pointer (AURP)
6	The virtual address of the storage segment table pointer (SSTP)
7	A logical interrupt service number (LISN)
8	Interrupt vector table, derived from the hypervisor call parameters.
9	A state register (SR) value
10	A logical partition ID (LPID)
11	A real address (RA) hypervisor accelerator utilization record pointer
12	The Storage Descriptor Register (SDR)

[00122] In one embodiment, the hypervisor initializes a plurality of accelerator integration slice 490 registers 445.

[00123] As illustrated in **FIG. 4F**, one exemplary embodiment of the invention employs a unified memory addressable via a common virtual memory address space used to access the physical processor memories 401-402 and GPU memories 420-423. In this

implementation, operations executed on the GPUs 410-413 utilize the same virtual/effective memory address space to access the processors memories 401-402 and vice versa, thereby simplifying programmability. In one embodiment, a first portion of the virtual/effective address space is allocated to the processor memory 401, a second
5 portion to the second processor memory 402, a third portion to the GPU memory 420, and so on. The entire virtual/effective memory space (sometimes referred to as the effective address space) is thereby distributed across each of the processor memories 401-402 and GPU memories 420-423, allowing any processor or GPU to access any physical memory with a virtual address mapped to that memory.

10 **[00124]** In one embodiment, bias/coherence management circuitry 494A-494E within one or more of the MMUs 439A-439E ensures cache coherence between the caches of the host processors (e.g., 405) and the GPUs 410-413 and implements biasing techniques indicating the physical memories in which certain types of data should be stored. While multiple instances of bias/coherence management circuitry 494A-494E are
15 illustrated in **FIG. 4F**, the bias/coherence circuitry may be implemented within the MMU of one or more host processors 405 and/or within the accelerator integration circuit 436.

[00125] One embodiment allows GPU-attached memory 420-423 to be mapped as part of system memory, and accessed using shared virtual memory (SVM) technology, but without suffering the typical performance drawbacks associated with full system
20 cache coherence. The ability to GPU-attached memory 420-423 to be accessed as system memory without onerous cache coherence overhead provides a beneficial operating environment for GPU offload. This arrangement allows the host processor 405 software to setup operands and access computation results, without the overhead of tradition I/O DMA data copies. Such traditional copies involve driver calls, interrupts and memory
25 mapped I/O (MMIO) accesses that are all inefficient relative to simple memory accesses. At the same time, the ability to access GPU attached memory 420-423 without cache coherence overheads can be critical to the execution time of an offloaded computation. In cases with substantial streaming write memory traffic, for example, cache coherence overhead can significantly reduce the effective write bandwidth seen by a GPU 410-413.

The efficiency of operand setup, the efficiency of results access, and the efficiency of GPU computation all play a role in determining the effectiveness of GPU offload.

[00126] In one implementation, the selection of between GPU bias and host processor bias is driven by a bias tracker data structure. A bias table may be used, for example, which may be a page-granular structure (i.e., controlled at the granularity of a memory page) that includes 1 or 2 bits per GPU-attached memory page. The bias table may be implemented in a stolen memory range of one or more GPU-attached memories 420-423, with or without a bias cache in the GPU 410-413 (e.g., to cache frequently/recently used entries of the bias table). Alternatively, the entire bias table may be maintained within the GPU.

[00127] In one implementation, the bias table entry associated with each access to the GPU-attached memory 420-423 is accessed prior the actual access to the GPU memory, causing the following operations. First, local requests from the GPU 410-413 that find their page in GPU bias are forwarded directly to a corresponding GPU memory 420-423. Local requests from the GPU that find their page in host bias are forwarded to the processor 405 (e.g., over a high-speed link as discussed above). In one embodiment, requests from the processor 405 that find the requested page in host processor bias complete the request like a normal memory read. Alternatively, requests directed to a GPU-biased page may be forwarded to the GPU 410-413. The GPU may then transition the page to a host processor bias if it is not currently using the page.

[00128] The bias state of a page can be changed either by a software-based mechanism, a hardware-assisted software-based mechanism, or, for a limited set of cases, a purely hardware-based mechanism.

[00129] One mechanism for changing the bias state employs an API call (e.g. OpenCL), which, in turn, calls the GPU's device driver which, in turn, sends a message (or enqueues a command descriptor) to the GPU directing it to change the bias state and, for some transitions, perform a cache flushing operation in the host. The cache flushing operation is required for a transition from host processor 405 bias to GPU bias, but is not required for the opposite transition.

[00130] In one embodiment, cache coherency is maintained by temporarily rendering GPU-biased pages uncacheable by the host processor 405. To access these pages, the processor 405 may request access from the GPU 410 which may or may not grant access right away, depending on the implementation. Thus, to reduce communication between the processor 405 and GPU 410 it is beneficial to ensure that GPU-biased pages are those which are required by the GPU but not the host processor 405 and vice versa.

Graphics Processing Pipeline

[00131] FIG. 5 illustrates a graphics processing pipeline 500 according to an exemplary embodiment. In one embodiment, a graphics processor can implement the illustrated graphics processing pipeline 500. The graphics processor can be included within the parallel processing subsystems as described herein, such as the parallel processor 200 of FIG. 2A, which, in one embodiment, is a variant of the parallel processor(s) 112 of FIG. 1. The various parallel processing systems can implement the graphics processing pipeline 500 via one or more instances of the parallel processing unit (e.g., parallel processing unit 202 of FIG. 2A) as described herein. For example, a shader unit (e.g., graphics multiprocessor 234 of FIGS. 2C-2D) may be configured to perform the functions of one or more of a vertex processing unit 504, a tessellation control processing unit 508, a tessellation evaluation processing unit 512, a geometry processing unit 516, and a fragment/pixel processing unit 524. The functions of data assembler 502, primitive assemblers 506, 514, 518, tessellation unit 510, rasterizer 522, and raster operations unit 526 may also be performed by other processing engines within a processing cluster (e.g., processing cluster 214 of FIG. 2A) and a corresponding partition unit (e.g., partition unit 220A-220N of FIG. 2). The graphics processing pipeline 500 may also be implemented using dedicated processing units for one or more functions. In one embodiment, one or more portions of the graphics processing pipeline 500 can be performed by parallel processing logic within a general purpose processor (e.g., CPU). In one embodiment, one or more portions of the graphics processing pipeline 500 can access on-chip memory (e.g., parallel processor memory 222 as in FIG.

2A) via a memory interface 528, which may be an instance of the memory interface 218 of FIG. 2A.

[00132] In one embodiment, the data assembler 502 is a processing unit that collects vertex data for surfaces and primitives. The data assembler 502 then outputs the vertex data, including the vertex attributes, to the vertex processing unit 504. The vertex processing unit 504 is a programmable execution unit that executes vertex shader programs, lighting and transforming vertex data as specified by the vertex shader programs. The vertex processing unit 504 reads data that is stored in cache, local or system memory for use in processing the vertex data and may be programmed to transform the vertex data from an object-based coordinate representation to a world space coordinate space or a normalized device coordinate space.

[00133] A first instance of a primitive assembler 506 receives vertex attributes from the vertex processing unit 504. The primitive assembler 506 reads stored vertex attributes as needed and constructs graphics primitives for processing by tessellation control processing unit 508. The graphics primitives include triangles, line segments, points, patches, and so forth, as supported by various graphics processing application programming interfaces (APIs).

[00134] The tessellation control processing unit 508 treats the input vertices as control points for a geometric patch. The control points are transformed from an input representation from the patch (e.g., the patch's bases) to a representation that is suitable for use in surface evaluation by the tessellation evaluation processing unit 512. The tessellation control processing unit 508 can also compute tessellation factors for edges of geometric patches. A tessellation factor applies to a single edge and quantifies a view-dependent level of detail associated with the edge. A tessellation unit 510 is configured to receive the tessellation factors for edges of a patch and to tessellate the patch into multiple geometric primitives such as line, triangle, or quadrilateral primitives, which are transmitted to a tessellation evaluation processing unit 512. The tessellation evaluation processing unit 512 operates on parameterized coordinates of the subdivided patch to generate a surface representation and vertex attributes for each vertex associated with the geometric primitives.

[00135] A second instance of a primitive assembler 514 receives vertex attributes from the tessellation evaluation processing unit 512, reading stored vertex attributes as needed, and constructs graphics primitives for processing by the geometry processing unit 516. The geometry processing unit 516 is a programmable execution unit that
5 executes geometry shader programs to transform graphics primitives received from primitive assembler 514 as specified by the geometry shader programs. In one embodiment, the geometry processing unit 516 is programmed to subdivide the graphics primitives into one or more new graphics primitives and calculate parameters used to rasterize the new graphics primitives.

10 [00136] In some embodiments, the geometry processing unit 516 can add or delete elements in the geometry stream. The geometry processing unit 516 outputs the parameters and vertices specifying new graphics primitives to primitive assembler 518. The primitive assembler 518 receives the parameters and vertices from the geometry processing unit 516 and constructs graphics primitives for processing by a viewport scale, cull, and clip unit 520. The geometry processing unit 516 reads data that is stored in
15 parallel processor memory or system memory for use in processing the geometry data. The viewport scale, cull, and clip unit 520 performs clipping, culling, and viewport scaling and outputs processed graphics primitives to a rasterizer 522.

[00137] The rasterizer 522 can perform depth culling and other depth-based
20 optimizations. The rasterizer 522 also performs scan conversion on the new graphics primitives to generate fragments and output those fragments and associated coverage data to the fragment/pixel processing unit 524. The fragment/pixel processing unit 524 is a programmable execution unit that is configured to execute fragment shader programs or pixel shader programs. The fragment/pixel processing unit 524 transforming fragments
25 or pixels received from rasterizer 522, as specified by the fragment or pixel shader programs. For example, the fragment/pixel processing unit 524 may be programmed to perform operations included but not limited to texture mapping, shading, blending, texture correction and perspective correction to produce shaded fragments or pixels that are output to a raster operations unit 526. The fragment/pixel processing unit 524 can
30 read data that is stored in either the parallel processor memory or the system memory for

use when processing the fragment data. Fragment or pixel shader programs may be configured to shade at sample, pixel, tile, or other granularities depending on the sampling rate configured for the processing units.

[00138] The raster operations unit 526 is a processing unit that performs raster operations including, but not limited to stencil, z test, blending, and the like, and outputs pixel data as processed graphics data to be stored in graphics memory (e.g., parallel processor memory 222 as in FIG. 2A, and/or system memory 104 as in FIG. 1, to be displayed on the one or more display device(s) 110 or for further processing by one of the one or more processor(s) 102 or parallel processor(s) 112. In some embodiments, the raster operations unit 526 is configured to compress z or color data that is written to memory and decompress z or color data that is read from memory.

Machine Learning Overview

[00139] A machine learning algorithm is an algorithm that can learn based on a set of data. Embodiments of machine learning algorithms can be designed to model high-level abstractions within a data set. For example, image recognition algorithms can be used to determine which of several categories to which a given input belong; regression algorithms can output a numerical value given an input; and pattern recognition algorithms can be used to generate translated text or perform text to speech and/or speech recognition.

[00140] An exemplary type of machine learning algorithm is a neural network. There are many types of neural networks; a simple type of neural network is a feedforward network. A feedforward network may be implemented as an acyclic graph in which the nodes are arranged in layers. Typically, a feedforward network topology includes an input layer and an output layer that are separated by at least one hidden layer. The hidden layer transforms input received by the input layer into a representation that is useful for generating output in the output layer. The network nodes are fully connected via edges to the nodes in adjacent layers, but there are no edges between nodes within each layer. Data received at the nodes of an input layer of a feedforward network are propagated (i.e., “fed forward”) to the nodes of the output layer via an activation function that calculates the states of the nodes of each successive layer in the network based on

coefficients (“weights”) respectively associated with each of the edges connecting the layers. Depending on the specific model being represented by the algorithm being executed, the output from the neural network algorithm can take various forms.

[00141] Before a machine learning algorithm can be used to model a particular
5 problem, the algorithm is trained using a training data set. Training a neural network involves selecting a network topology, using a set of training data representing a problem being modeled by the network, and adjusting the weights until the network model performs with a minimal error for all instances of the training data set. For example, during a supervised learning training process for a neural network, the output produced
10 by the network in response to the input representing an instance in a training data set is compared to the “correct” labeled output for that instance, an error signal representing the difference between the output and the labeled output is calculated, and the weights associated with the connections are adjusted to minimize that error as the error signal is backward propagated through the layers of the network. The network is considered
15 “trained” when the errors for each of the outputs generated from the instances of the training data set are minimized.

[00142] The accuracy of a machine learning algorithm can be affected significantly by the quality of the data set used to train the algorithm. The training process can be computationally intensive and may require a significant amount of time on a conventional
20 general-purpose processor. Accordingly, parallel processing hardware is used to train many types of machine learning algorithms. This is particularly useful for optimizing the training of neural networks, as the computations performed in adjusting the coefficients in neural networks lend themselves naturally to parallel implementations. Specifically, many machine learning algorithms and software applications have been adapted to make
25 use of the parallel processing hardware within general-purpose graphics processing devices.

[00143] FIG. 6 is a generalized diagram of a machine learning software stack 600. A machine learning application 602 can be configured to train a neural network using a training dataset or to use a trained deep neural network to implement machine
30 intelligence. The machine learning application 602 can include training and inference

functionality for a neural network and/or specialized software that can be used to train a neural network before deployment. The machine learning application 602 can implement any type of machine intelligence including but not limited to image recognition, mapping and localization, autonomous navigation, speech synthesis, medical imaging, or language translation.

5 [00144] Hardware acceleration for the machine learning application 602 can be enabled via a machine learning framework 604. The machine learning framework 604 can provide a library of machine learning primitives. Machine learning primitives are basic operations that are commonly performed by machine learning algorithms. Without
10 the machine learning framework 604, developers of machine learning algorithms would be required to create and optimize the main computational logic associated with the machine learning algorithm, then re-optimize the computational logic as new parallel processors are developed. Instead, the machine learning application can be configured to perform the necessary computations using the primitives provided by the machine
15 learning framework 604. Exemplary primitives include tensor convolutions, activation functions, and pooling, which are computational operations that are performed while training a convolutional neural network (CNN). The machine learning framework 604 can also provide primitives to implement basic linear algebra subprograms performed by many machine-learning algorithms, such as matrix and vector operations.

20 [00145] The machine learning framework 604 can process input data received from the machine learning application 602 and generate the appropriate input to a compute framework 606. The compute framework 606 can abstract the underlying instructions provided to the GPGPU driver 608 to enable the machine learning framework 604 to take advantage of hardware acceleration via the GPGPU hardware 610
25 without requiring the machine learning framework 604 to have intimate knowledge of the architecture of the GPGPU hardware 610. Additionally, the compute framework 606 can enable hardware acceleration for the machine learning framework 604 across a variety of types and generations of the GPGPU hardware 610.

GPGPU Machine Learning Acceleration

[00146] FIG. 7 illustrates a highly-parallel general-purpose graphics processing unit 700 according to an exemplary embodiment. In one embodiment, the general-purpose processing unit (GPGPU) 700 can be configured to be particularly efficient in processing the type of computational workloads associated with training deep neural networks. Additionally, the GPGPU 700 can be linked directly to other instances of the GPGPU to create a multi-GPU cluster to improve training speed for particularly deep neural networks.

[00147] The GPGPU 700 includes a host interface 702 to enable a connection with a host processor. In one embodiment, the host interface 702 is a PCI Express interface. However, the host interface can also be a vendor specific communications interface or communications fabric. The GPGPU 700 receives commands from the host processor and uses a global scheduler 704 to distribute execution threads associated with those commands to a set of compute clusters 706A-H. The compute clusters 706A-H share a cache memory 708. The cache memory 708 can serve as a higher-level cache for cache memories within the compute clusters 706A-H.

[00148] The GPGPU 700 includes memory 714A-B coupled with the compute clusters 706A-H via a set of memory controllers 712A-B. In various embodiments, the memory 714A-B can include various types of memory devices including dynamic random access memory (DRAM) or graphics random access memory, such as synchronous graphics random access memory (SGRAM), including graphics double data rate (GDDR) memory. In one embodiment, the memory units 224A-N may also include 3D stacked memory, including but not limited to high bandwidth memory (HBM).

[00149] In one embodiment, each compute cluster 706A-H includes a set of graphics multiprocessors, such as the graphics multiprocessor 400 of FIG. 4A. The graphics multiprocessors of the compute cluster multiple types of integer and floating point logic units that can perform computational operations at a range of precisions including suited for machine learning computations. For example, and in one embodiment, at least a subset of the floating-point units in each of the compute clusters 706A-H can be configured to perform 16-bit or 32-bit floating point operations, while a

different subset of floating point units can be configured to perform 64-bit floating point operations.

[00150] Multiple instances of the GPGPU 700 can be configured to operate as a compute cluster. The communication mechanism used by the compute cluster for synchronization and data exchange varies across embodiments. In one embodiment, the multiple instances of the GPGPU 700 communicate over the host interface 702. In one embodiment, the GPGPU 700 includes an I/O hub 708 that couples the GPGPU 700 with a GPU link 710 that enables a direct connection to other instances of the GPGPU. In one embodiment, the GPU link 710 is coupled to a dedicated GPU-to-GPU bridge that enables communication and synchronization between multiple instances of the GPGPU 700. In one embodiment, the GPU link 710 couples with a high speed interconnect to transmit and receive data to other GPGPUs or parallel processors. In one embodiment, the multiple instances of the GPGPU 700 are located in separate data processing systems and communicate via a network device that is accessible via the host interface 702. In one embodiment, the GPU link 710 can be configured to enable a connection to a host processor in addition to or as an alternative to the host interface 702.

[00151] While the illustrated configuration of the GPGPU 700 can be configured to train neural networks, one embodiment provides alternate configuration of the GPGPU 700 that can be configured for deployment within a high performance or low power inferencing platform. In an inferencing configuration, the GPGPU 700 includes fewer of the compute clusters 706A-H relative to the training configuration. Additionally, memory technology associated with the memory 714A-B may differ between inferencing and training configurations. In one embodiment, the inferencing configuration of the GPGPU 700 can support inferencing specific instructions. For example, an inferencing configuration can provide support for one or more 8-bit integer dot product instructions, which are commonly used during inferencing operations for deployed neural networks.

[00152] FIG. 8 illustrates a multi-GPU computing system 800 according to an exemplary embodiment. The multi-GPU computing system 800 can include a processor 802 coupled to multiple GPGPUs 806A-D via a host interface switch 804. The host interface switch 804, in one embodiment, is a PCI express switch device that couples the

processor 802 to a PCI express bus over which the processor 802 can communicate with the set of GPGPUs 806A-D. Each of the multiple GPGPUs 806A-D can be an instance of the GPGPU 700 of FIG. 7. The GPGPUs 806A-D can interconnect via a set of high-speed point to point GPU to GPU links 816. The high-speed GPU to GPU links can connect to each of the GPGPUs 806A-D via a dedicated GPU link, such as the GPU link 710 as in FIG. 7. The P2P GPU links 816 enable direct communication between each of the GPGPUs 806A-D without requiring communication over the host interface bus to which the processor 802 is connected. With GPU-to-GPU traffic directed to the P2P GPU links, the host interface bus remains available for system memory access or to communicate with other instances of the multi-GPU computing system 800, for example, via one or more network devices. While in the illustrated embodiment the GPGPUs 806A-D connect to the processor 802 via the host interface switch 804, in one embodiment the processor 802 includes direct support for the P2P GPU links 816 and can connect directly to the GPGPUs 806A-D.

15 Machine Learning Neural Network Implementations

[00153] The computing architecture provided by embodiments described herein can be configured to perform the types of parallel processing that is particularly suited for training and deploying neural networks for machine learning. A neural network can be generalized as a network of functions having a graph relationship. As is well-known in the art, there are a variety of types of neural network implementations used in machine learning. One exemplary type of neural network is the feedforward network, as previously described.

[00154] A second exemplary type of neural network is the Convolutional Neural Network (CNN). A CNN is a specialized feedforward neural network for processing data having a known, grid-like topology, such as image data. Accordingly, CNNs are commonly used for computer vision and image recognition applications, but they also may be used for other types of pattern recognition such as speech and language processing. The nodes in the CNN input layer are organized into a set of “filters” (feature detectors inspired by the receptive fields found in the retina), and the output of each set of filters is propagated to nodes in successive layers of the network. The

computations for a CNN include applying the convolution mathematical operation to each filter to produce the output of that filter. Convolution is a specialized kind of mathematical operation performed by two functions to produce a third function that is a modified version of one of the two original functions. In convolutional network terminology, the first function to the convolution can be referred to as the input, while the second function can be referred to as the convolution kernel. The output may be referred to as the feature map. For example, the input to a convolution layer can be a multidimensional array of data that defines the various color components of an input image. The convolution kernel can be a multidimensional array of parameters, where the parameters are adapted by the training process for the neural network.

[00155] Recurrent neural networks (RNNs) are a family of feedforward neural networks that include feedback connections between layers. RNNs enable modeling of sequential data by sharing parameter data across different parts of the neural network. The architecture for a RNN includes cycles. The cycles represent the influence of a present value of a variable on its own value at a future time, as at least a portion of the output data from the RNN is used as feedback for processing subsequent input in a sequence. This feature makes RNNs particularly useful for language processing due to the variable nature in which language data can be composed.

[00156] The figures described below present exemplary feedforward, CNN, and RNN networks, as well as describe a general process for respectively training and deploying each of those types of networks. It will be understood that these descriptions are exemplary and non-limiting as to any specific embodiment described herein and the concepts illustrated can be applied generally to deep neural networks and machine learning techniques in general.

[00157] The exemplary neural networks described above can be used to perform deep learning. Deep learning is machine learning using deep neural networks. The deep neural networks used in deep learning are artificial neural networks composed of multiple hidden layers, as opposed to shallow neural networks that include only a single hidden layer. Deeper neural networks are generally more computationally intensive to train.

However, the additional hidden layers of the network enable multistep pattern recognition that results in reduced output error relative to shallow machine learning techniques.

[00158] Deep neural networks used in deep learning typically include a front-end network to perform feature recognition coupled to a back-end network which represents a
5 mathematical model that can perform operations (e.g., object classification, speech recognition, etc.) based on the feature representation provided to the model. Deep learning enables machine learning to be performed without requiring hand crafted feature engineering to be performed for the model. Instead, deep neural networks can learn
10 features based on statistical structure or correlation within the input data. The learned features can be provided to a mathematical model that can map detected features to an output. The mathematical model used by the network is generally specialized for the specific task to be performed, and different models will be used to perform different task.

[00159] Once the neural network is structured, a learning model can be applied to the network to train the network to perform specific tasks. The learning model describes
15 how to adjust the weights within the model to reduce the output error of the network. Backpropagation of errors is a common method used to train neural networks. An input vector is presented to the network for processing. The output of the network is compared to the desired output using a loss function and an error value is calculated for each of the neurons in the output layer. The error values are then propagated backwards until each
20 neuron has an associated error value which roughly represents its contribution to the original output. The network can then learn from those errors using an algorithm, such as the stochastic gradient descent algorithm, to update the weights of the of the neural network.

[00160] FIGS. 9A-B illustrate an exemplary convolutional neural network. FIG.
25 9A illustrates various layers within a CNN. As shown in FIG. 9A, an exemplary CNN used to model image processing can receive input 902 describing the red, green, and blue (RGB) components of an input image. The input 902 can be processed by multiple convolutional layers (e.g., convolutional layer 904, convolutional layer 906). The output from the multiple convolutional layers may optionally be processed by a set of fully
30 connected layers 908. Neurons in a fully connected layer have full connections to all

activations in the previous layer, as previously described for a feedforward network. The output from the fully connected layers 908 can be used to generate an output result from the network. The activations within the fully connected layers 908 can be computed using matrix multiplication instead of convolution. Not all CNN implementations make use of fully connected layers 906. For example, in some implementations, the convolutional layer 906 can generate output for the CNN.

[00161] The convolutional layers are sparsely connected, which differs from traditional neural network configuration found in the fully connected layers 908. Traditional neural network layers are fully connected, such that every output unit interacts with every input unit. However, the convolutional layers are sparsely connected because the output of the convolution of a field is input (instead of the respective state value of each of the nodes in the field) to the nodes of the subsequent layer, as illustrated. The kernels associated with the convolutional layers perform convolution operations, the output of which is sent to the next layer. The dimensionality reduction performed within the convolutional layers is one aspect that enables the CNN to scale to process large images.

[00162] FIG. 9B illustrates exemplary computation stages within a convolutional layer of a CNN. Input to a convolutional layer 912 of a CNN can be processed in three stages of a convolutional layer 914. The three stages can include a convolution stage 916, a detector stage 918, and a pooling stage 920. The convolution layer 914 can then output data to a successive convolutional layer. The final convolutional layer of the network can generate output feature map data or provide input to a fully connected layer, for example, to generate a classification value for the input to the CNN.

[00163] In the convolution stage 916, the convolutional layer 914 can perform several convolutions in parallel to produce a set of linear activations. The convolution stage 916 can include an affine transformation, which is any transformation that can be specified as a linear transformation plus a translation. Affine transformations include rotations, translations, scaling, and combinations of these transformations. The convolution stage computes the output of functions (e.g., neurons) that are connected to specific regions in the input, which can be determined as the local region associated with

the neuron. The neurons compute a dot product between the weights of the neurons and the region in the local input to which the neurons are connected. The output from the convolution stage 916 defines a set of linear activations that are processed by successive stages of the convolutional layer 914.

5 [00164] The linear activations can be processed by a detector stage 918. In the detector stage 918, each linear activation is processed by a non-linear activation function. The non-linear activation function increases the nonlinear properties of the overall network without affecting the receptive fields of the convolution layer. Several types of non-linear activation functions may be used. One particular type is the rectified linear unit (ReLU), which uses an activation function defined as $f(x) = \max(0, x)$, such that
10 the activation is threshold at zero.

[00165] The pooling stage 920 uses a pooling function that replaces the output of the convolutional layer 906 with a summary statistic of the nearby outputs. The pooling function can be used to introduce translation invariance into the neural network, such that
15 small translations to the input do not change the pooled outputs. Invariance to local translation can be useful in scenarios where the presence of a feature in the input data is more important than the precise location of the feature. Various types of pooling functions can be used during the pooling stage 920, including max pooling, average pooling, and l2-norm pooling. Additionally, some CNN implementations do not include
20 a pooling stage. Instead, such implementations substitute an additional convolution stage having an increased stride relative to previous convolution stages.

[00166] The output from the convolutional layer 914 can then be processed by the next layer 922. The next layer 922 can be an additional convolutional layer or one of the fully connected layers 908. For example, the first convolutional layer 904 of **FIG. 9A**
25 can output to the second convolutional layer 906, while the second convolutional layer can output to a first layer of the fully connected layers 908.

[00167] **FIG. 10** illustrates an exemplary recurrent neural network 1000. In a recurrent neural network (RNN), the previous state of the network influences the output of the current state of the network. RNNs can be built in a variety of ways using a variety
30 of functions. The use of RNNs generally revolves around using mathematical models to

predict the future based on a prior sequence of inputs. For example, an RNN may be used to perform statistical language modeling to predict an upcoming word given a previous sequence of words. The illustrated RNN 1000 can be described as having an input layer 1002 that receives an input vector, hidden layers 1004 to implement a recurrent function, a feedback mechanism 1005 to enable a 'memory' of previous states, and an output layer 1006 to output a result. The RNN 1000 operates based on time-steps. The state of the RNN at a given time step is influenced based on the previous time step via the feedback mechanism 1005. For a given time step, the state of the hidden layers 1004 is defined by the previous state and the input at the current time step. An initial input (x_1) at a first time step can be processed by the hidden layer 1004. A second input (x_2) can be processed by the hidden layer 1004 using state information that is determined during the processing of the initial input (x_1). A given state can be computed as $s_t = f(Ux_t + Ws_{t-1})$, where U and W are parameter matrices. The function f is generally a nonlinearity, such as the hyperbolic tangent function (Tanh) or a variant of the rectifier function $f(x) = \max(0, x)$. However, the specific mathematical function used in the hidden layers 1004 can vary depending on the specific implementation details of the RNN 1000.

[00168] In addition to the basic CNN and RNN networks described, variations on those networks may be enabled. One example RNN variant is the long short term memory (LSTM) RNN. LSTM RNNs are capable of learning long-term dependencies that may be necessary for processing longer sequences of language. A variant on the CNN is a convolutional deep belief network, which has a structure similar to a CNN and is trained in a manner similar to a deep belief network. A deep belief network (DBN) is a generative neural network that is composed of multiple layers of stochastic (random) variables. DBNs can be trained layer-by-layer using greedy unsupervised learning. The learned weights of the DBN can then be used to provide pre-train neural networks by determining an optimal initial set of weights for the neural network.

[00169] FIG. 11 illustrates an exemplary training and deployment of a deep neural network. Once a given network has been structured for a task the neural network is trained using a training dataset 1102. Various training frameworks 1104 have been

developed to enable hardware acceleration of the training process. For example, the machine learning framework 604 of FIG. 6 may be configured as a training framework 604. The training framework 604 can hook into an untrained neural network 1106 and enable the untrained neural net to be trained using the parallel processing resources
5 described herein to generate a trained neural net 1108.

[00170] To start the training process the initial weights may be chosen randomly or by pre-training using a deep belief network. The training cycle then be performed in either a supervised or unsupervised manner.

[00171] Supervised learning is a learning method in which training is performed as a mediated operation, such as when the training dataset 1102 includes input paired with
10 the desired output for the input, or where the training dataset includes input having known output and the output of the neural network is manually graded. The network processes the inputs and compares the resulting outputs against a set of expected or desired outputs. Errors are then propagated back through the system. The training
15 framework 1104 can adjust to adjust the weights that control the untrained neural network 1106. The training framework 1104 can provide tools to monitor how well the untrained neural network 1106 is converging towards a model suitable to generating correct answers based on known input data. The training process occurs repeatedly as the weights of the network are adjusted to refine the output generated by the neural network.
20 The training process can continue until the neural network reaches a statistically desired accuracy associated with a trained neural net 1108. The trained neural network 1108 can then be deployed to implement any number of machine learning operations.

[00172] Unsupervised learning is a learning method in which the network attempts to train itself using unlabeled data. Thus, for unsupervised learning the training dataset
25 1102 will include input data without any associated output data. The untrained neural network 1106 can learn groupings within the unlabeled input and can determine how individual inputs are related to the overall dataset. Unsupervised training can be used to generate a self-organizing map, which is a type of trained neural network 1107 capable of performing operations useful in reducing the dimensionality of data. Unsupervised

training can also be used to perform anomaly detection, which allows the identification of data points in an input dataset that deviate from the normal patterns of the data.

[00173] Variations on supervised and unsupervised training may also be employed. Semi-supervised learning is a technique in which in the training dataset 1102 includes a mix of labeled and unlabeled data of the same distribution. Incremental learning is a variant of supervised learning in which input data is continuously used to further train the model. Incremental learning enables the trained neural network 1108 to adapt to the new data 1112 without forgetting the knowledge instilled within the network during initial training.

10 [00174] Whether supervised or unsupervised, the training process for particularly deep neural networks may be too computationally intensive for a single compute node. Instead of using a single compute node, a distributed network of computational nodes can be used to accelerate the training process.

[00175] **FIG. 12** is an exemplary block diagram illustrating distributed learning. Distributed learning is a training model that uses multiple distributed computing nodes to perform supervised or unsupervised training of a neural network. The distributed computational nodes can each include one or more host processors and one or more of the general-purpose processing nodes, such as the highly-parallel general-purpose graphics processing unit 700 as in **FIG. 7**. As illustrated, distributed learning can be performed model parallelism 1202, data parallelism 1204, or a combination of model and data parallelism 1204.

20 [00176] In model parallelism 1202, different computational nodes in a distributed system can perform training computations for different parts of a single network. For example, each layer of a neural network can be trained by a different processing node of the distributed system. The benefits of model parallelism include the ability to scale to particularly large models. Splitting the computations associated with different layers of the neural network enables the training of very large neural networks in which the weights of all layers would not fit into the memory of a single computational node. In some instances, model parallelism can be particularly useful in performing unsupervised training of large neural networks.

25
30

[00177] In data parallelism 1204, the different nodes of the distributed network have a complete instance of the model and each node receives a different portion of the data. The results from the different nodes are then combined. While different approaches to data parallelism are possible, data parallel training approaches all require a technique of combining results and synchronizing the model parameters between each node. Exemplary approaches to combining data include parameter averaging and update based data parallelism. Parameter averaging trains each node on a subset of the training data and sets the global parameters (e.g., weights, biases) to the average of the parameters from each node. Parameter averaging uses a central parameter server that maintains the parameter data. Update based data parallelism is similar to parameter averaging except that instead of transferring parameters from the nodes to the parameter server, the updates to the model are transferred. Additionally, update based data parallelism can be performed in a decentralized manner, where the updates are compressed and transferred between nodes.

15 [00178] Combined model and data parallelism 1206 can be implemented, for example, in a distributed system in which each computational node includes multiple GPUs. Each node can have a complete instance of the model with separate GPUs within each node are used to train different portions of the model.

[00179] Distributed training has increased overhead relative to training on a single machine. However, the parallel processors and GPGPUs described herein can each implement various techniques to reduce the overhead of distributed training, including techniques to enable high bandwidth GPU-to-GPU data transfer and accelerated remote data synchronization.

Exemplary Machine Learning Applications

25 [00180] Machine learning can be applied to solve a variety of technological problems, including but not limited to computer vision, autonomous driving and navigation, speech recognition, and language processing. Computer vision has traditionally been one of the most active research areas for machine learning applications. Applications of computer vision range from reproducing human visual abilities, such as recognizing faces, to creating new categories of visual abilities. For example, computer

vision applications can be configured to recognize sound waves from the vibrations induced in objects visible in a video. Parallel processor accelerated machine learning enables computer vision applications to be trained using significantly larger training dataset than previously feasible and enables inferencing systems to be deployed using
5 low power parallel processors.

[00181] Parallel processor accelerated machine learning has autonomous driving applications including lane and road sign recognition, obstacle avoidance, navigation, and driving control. Accelerated machine learning techniques can be used to train driving models based on datasets that define the appropriate responses to specific training input.

10 The parallel processors described herein can enable rapid training of the increasingly complex neural networks used for autonomous driving solutions and enables the deployment of low power inferencing processors in a mobile platform suitable for integration into autonomous vehicles.

[00182] Parallel processor accelerated deep neural networks have enabled machine learning approaches to automatic speech recognition (ASR). ASR includes the creation of a function that computes the most probable linguistic sequence given an input acoustic sequence. Accelerated machine learning using deep neural networks have enabled the replacement of the hidden Markov models (HMMs) and Gaussian mixture models (GMMs) previously used for ASR.

20 **[00183]** Parallel processor accelerated machine learning can also be used to accelerate natural language processing. Automatic learning procedures can make use of statistical inference algorithms to produce models that are robust to erroneous or unfamiliar input. Exemplary natural language processor applications include automatic machine translation between human languages.

25 **[00184]** The parallel processing platforms used for machine learning can be divided into training platforms and deployment platforms. Training platforms are generally highly parallel and include optimizations to accelerate multi-GPU single node training and multi-node, multi-GPU training. Exemplary parallel processors suited for training include the highly-parallel general-purpose graphics processing unit 700 of FIG.
30 700 and the multi-GPU computing system 800 of FIG. 800. On the contrary, deployed

machine learning platforms generally include lower power parallel processors suitable for use in products such as cameras, autonomous robots, and autonomous vehicles.

5 [00185] FIG. 13 illustrates an exemplary inferencing system on a chip (SOC) 1300 suitable for performing inferencing using a trained model. The SOC 1300 can integrate processing components including a media processor 1302, a vision processor 1304, a GPGPU 1306 and a multi-core processor 1308. The SOC 1300 can additionally include on-chip memory 1305 that can enable a shared on-chip data pool that is accessible by each of the processing components. The processing components can be optimized for low power operation to enable deployment to a variety of machine learning platforms, including autonomous vehicles and autonomous robots. For example, one 10 implementation of the SOC 1300 can be used as a portion of the main control system for an autonomous vehicle. Where the SOC 1300 is configured for use in autonomous vehicles the SOC is designed and configured for compliance with the relevant functional safety standards of the deployment jurisdiction.

15 [00186] During operation, the media processor 1302 and vision processor 1304 can work in concert to accelerate computer vision operations. The media processor 1302 can enable low latency decode of multiple high-resolution (e.g., 4K, 8K) video streams. The decoded video streams can be written to a buffer in the on-chip-memory 1305. The vision processor 1304 can then parse the decoded video and perform preliminary 20 processing operations on the frames of the decoded video in preparation of processing the frames using a trained image recognition model. For example, the vision processor 1304 can accelerate convolution operations for a CNN that is used to perform image recognition on the high-resolution video data, while back end model computations are performed by the GPGPU 1306.

25 [00187] The multi-core processor 1308 can include control logic to assist with sequencing and synchronization of data transfers and shared memory operations performed by the media processor 1302 and the vision processor 1304. The multi-core processor 1308 can also function as an application processor to execute software applications that can make use of the inferencing compute capability of the GPGPU 1306. 30 For example, at least a portion of the navigation and driving logic can be implemented in

software executing on the multi-core processor 1308. Such software can directly issue computational workloads to the GPGPU 1306 or the computational workloads can be issued to the multi-core processor 1308, which can offload at least a portion of those operations to the GPGPU 1306.

5 [00188] The GPGPU 1306 can include compute clusters such as a low power configuration of the compute clusters 706A-706H within the highly-parallel general-purpose graphics processing unit 700. The compute clusters within the GPGPU 1306 can support instruction that are specifically optimized to perform inferencing computations on a trained neural network. For example, the GPGPU 1306 can support instructions to
10 perform low precision computations such as 8-bit and 4-bit integer vector operations.

Budgeted and Simplified Training of Deep Neural Networks (DNNs)

[00189] FIG. 14 is an exemplary block diagram of a basic training and learning architecture 1400 having a deep neural network (DNN) training and learning system 1404 to receive training data 1402 for learning and training a DNN. Training data 1402 can be
15 input images having any number of samples of images (testing samples) or related data and information used. Such input images can be used budgeted and simplified training and learning by DNN training and learning system 1404. In some embodiments, DNN training and learning system 1404 can include or be implemented by or with the systems and processors disclosed and described in FIGS. 1-8 and 19-32. In other embodiments,
20 DNN training and learning system 1404 can be implemented using hardware acceleration as described in FIGS. 6 and 7. Training data (input images) 1402 can include image samples in any number of formats including Red Green Blue (RGB) format having R, G, and B channel values. Other examples include image samples in Color Space Pixel (YUV) having luminance, color, and chrominance channel values.

25 [00190] In exemplary embodiments, DNN training and learning system 1404 implements budgeted training and learning techniques for a DNN according to the description regarding FIGS. 15 and 16A-16C. For these exemplary embodiments, a small set of pixels are processed to train a DNN. In this way, input image size can be reduced such that large DNNs can be trained having less memory requirements.

[00191] In other exemplary embodiments, DNN training and learning system 1404 implements simplified training and learning techniques for a DNN such as a deep recurrent Q-network (DRQN) according to the description regarding **FIGS. 17A-17B** and **18A-18B**. For these exemplary embodiments, advantages of soft and hard attention mechanisms to train a DQRN are combined, which are differentiable and computational efficient.

(Budgeted Training of DNNs)

[00192] **FIG. 15** illustrates an exemplary process 1500 for selecting image data for a learning machine 1610. In this example, a large input image 1502 having a large array dimension width (W) and height (H) is sub-sampled. Large input image 1502 can include any type of image data and format, e.g., RGB format having red, green, and blue channels and values or YUV format having brightness, luminance, and color chrominance channels and values. The output of sub-sampling large input image 1502 includes N number of smaller sub-sampled images illustrated as sub-image 1 (1504-1) through sub-image N (1504-N). In this exemplary embodiment, sub-image 1 (1504-1) through sub-image N (1504-N) have a smaller dimension than the large input image 1502. Among the sub-images (e.g., sub-image 1 (1504-1) through sub-image N (1504-N)), one of them is randomly selected as randomly selected sub-image 1506 and is identified as selected sub-image k (1508). The selected sub-image k (1508) is input and processed by learning machine 1501 according to exemplary embodiments of **FIGS. 16A-16C**. In exemplary embodiments, learning machine 1501 can include a learning system 1600 described in **FIG. 16A** and N and K can be integers.

[00193] **FIG. 16A** is an exemplary block diagram of a learning system 1600 which can implement budgeted training and learning techniques for a DNN. In exemplary embodiments, learning system 1600 implements example training and learning techniques disclosed regarding **FIGS. 16B-16C**. Referring to **FIG. 16A**, image down sampler 1604 can down sample a large input image from image data 1602 into a number of smaller sub-images (e.g., sub-image 1 (1504-1) through sub-image N (1504-N) in **FIG. 15**). In one embodiment, the down sampled images have a lower resolution than the input image (e.g., input image 1502) from image data 1602 and are grouped in batches.

For example, one batch can include sub-images 1 and 2 and another batch can include sub-images 3 and 4 and so on in which sub-image N is placed in a batch with one or more other sub-images.

5 [00194] In an exemplary embodiment, batch selector 1606 selects one of the batches of sub-images and feeds it to trainer 1608. Batch selector 1606 can feed other batches to trainer 1608 including the last batch having sub-image N (1504-N). Trainer 1608 feeds selected batches of sub-images to a powerful neural network 1610, which can be any type of deep neural network (DNN), e.g., a deep convolutional neural network (CNN). In an exemplary embodiment, powerful neural network 1610 processes batches
10 of sub-images from trainer 1608 and updates parameters (e.g., updating weights at nodes of layers in the network) based on the processed sub-images and resulting classification by the network. In this way, the powerful neural network 1610 is trained using the training data 1602 to improve classification. The output of the powerful neural network 1610 can be a classification result, which can be input to tester 1616.

15 [00195] Tester 1616 can also process test images from test data 1612. In one embodiment, a large test image from test data 612 is down-sampled by image downsampler 1616 into smaller sub-images. Image downsampler 1616 can pass one of the smaller sub-images to tester 1616 which can test the sub-image using the trained powerful neural network 1610, which has been trained using the down-sampled images 1
20 (1504-1) through N (1504-N), to obtain a test result. Example of a test result can include a classification result for the tested sub-image. Tester 1616 forwards the test results to combiner 1618. In exemplary embodiments, combiner 1618 combines all the test results from training data 1602 and determines a final result 1619 using any number of methods, e.g., a voting method, which can be the best result among the test results. The exemplary
25 embodiments use sub-sampled images which have a lower image resolution than a full, large input image reducing computation and memory requirements while leveraging the capabilities of a powerful DNN.

[00196] FIGS. 16B-16C illustrate exemplary flow diagrams of operations 1620 and 1650 to train and test a DNN. Referring to operation 1620 of FIG. 16B, at operation
30 1622, a training image is randomly down sampled to generate low resolution sub-images.

For example, a training image can be divided into a number of sub-sections, and the sub-sections can be randomly selected to generate a sub-image having a lower resolution than the training image. In this way, a DNN can use a number of smaller images with lower resolution for training.

5 [00197] At operation 1624, the generated low resolution sub-images are divided into batches for training. In one example, the batches can include sub-images from different input images and the batches can be randomly selected for training. At operation 1626, the batches of sub-images are processed by a DNN (e.g., power neural network 1610) for training. For example, a first batch of sub-images can be processed by
10 the DNN to determine classification results and weights can be adjusted based on the classification results. The adjusted weights can thus be used for the next batch of sub-images in training the DNN.

[00198] Referring to operation 1650 of FIG. 16C, at operation 1652, a training image is randomly down sampled to generate low resolution sub-images for testing by
15 the DNN. The down-sampling of the testing image can be performed in the same way as down-sampling of the training image. At operation 1654, one of the down-sampled low resolution sub-images is input to the learning machine (or DNN) to process the sub-images to obtain a test result. The test result can be a classification result of the processed sub-images by the learning machine. Other down-sampled low resolution sub-
20 images can be input to the learning machine one at a time to obtain a number of test results. At operation 1656, the test results for the are combined and a final result can be selected based on a voting method or any other type of method to select a final result among a number of results.

(Simplified Training of a DQRN)

25 [00199] FIGS. 17A-18B illustrate exemplary embodiments of simplified training for a Deep Recurrent Q-Network (DRQN). A DRQN combines a Deep Q-Network (DQN) with Long-Short Term Memory (LSTM). A DQN is type of deep neural network (DNN) such as a deep convolutional neural network (CNN) with Q-learning. Q-Learning is a reinforcement learning technique that can be used to find an optimal action-selection
30 policy for any given state. DQN can provide human-level control policies on a variety of

tasks, however, drawbacks require long training time and lacks long term memory. DRQN combines a DQN with LSTM to acquire long term memory capabilities. DRQN can retain information from previous states and use it to predict better Q-values.

[00200] Exemplary embodiments of **FIGS. 17A-18B** provide a simplified and
5 improved DRQN to focus on relatively small informative regions of an input image thereby reducing the parameters and computational operations needed for training and testing a DNN. For example, DRQNs are disclosed with an improved local attention mechanism which forms part of a DRQN that can make decisions on a next action based
10 on a chosen region of an input image (i.e., attended region) rather than the entire region of an input image. Local attention includes two parts: “soft attention” and “hard attention.” Soft attention assigns soft weights over all inputs and is differentiable and can be trained using gradient based methods. Hard attention selects one patch of an image to attend to a time, which requires less parameters and is more computational efficient, but it is not differentiable and thus harder to train in comparison to soft attention.

15 [00201] In the disclosed examples, an improved local attention for a DRQN is provided that combines advantages of soft and hard attention and is differentiable and computation efficient. The improved local attention for a DRQN can be implemented in at least two exemplary embodiments: (1) the local attention mechanism (or network) is located between a convolutional neural network (CNN) and a LSTM (e.g., **FIGS. 17A-**
20 **17B**), and (2) a Hierarchical Memory Network (HMN) is located between a CNN and a LSTM (e.g., **FIGS. 18A-18B**).

[00202] **FIG. 17A** is a block diagram of an improved DRQN architecture 1700 according to an exemplary embodiment. Improved DRQN architecture 1700 includes location attention network 1708 located between convolutional neural network (CNN) 1704 and long-short time memory 1712. CNN 1704 can be configured and initialized by architecture configuration 1701. Training data 1702 can provide image data or an input image (e.g., training sample) to CNN 1704. CNN 1704 can process an input image and extract a feature map of the input image, which are fed to hard attention selector 1705 located within local attention network 1708. In an exemplary embodiment, hard attention
25 selector 1706 selects a subset of the feature maps, and discards others. In this example,
30

only subsets of feature maps of the input image are passed to soft attention assigner 1710, which assigns weights only for the selected subsets of feature maps by hard attention selector 1706 for training a DNN such as DRQN 1700. In this way, a smaller set of parameters (e.g., weights) are needed for training.

5 [00203] The outputs (e.g., weighted feature maps) of the soft attention assigner can be stored in LSTM 1712 in which hard attention selector 1706 can use stored information to select the next feature maps of input images. Q value estimation 1714 can also use the information stored in LSTM 172 to output a Q value for different actions based on the weighted feature maps to iteration manager 1713, which controls the iteration condition.

10 The iteration manager 1713 also manages the number of iterations or convergence of applying weights to selected feature maps and deriving resulting Q values (e.g., estimation on how valuable an action is). Iteration manager 1713 can output the final deep Q-network 1716.

[00204] FIG. 17B illustrates a flow diagram of an exemplary operation 1720 for the improved DRQN of FIG. 17A. In exemplary embodiments, at operation 1722, a current visual frame (input image) is received and a set of feature maps is extracted (e.g., extracted feature maps from CNN 1704). At operation 1724, a subset of feature maps is selected (e.g., hard attention selector 1706 selects the subset of feature maps). At operation 1726, weights are assigned over the selected subset of feature maps and weight parameters for non-selected feature maps can be set to zero '0'. At operation 1728, the Q-value is calculated for different actions based on the weighted feature maps (e.g., Q estimation 1714 calculates the Q-values). At operation 1730, a decision is made if a convergence or iteration number has been satisfied, and, if yes, the final deep Q-network 1716 is output at operation 1732. If no, operation 1720 returns to operation 1724 and repeats operations 1724 to 1728.

20

25

[00205] FIG. 18A is a block diagram of an improved DRQN architecture 1800 according to another exemplary embodiment. Improved DRQN architecture 1800 includes hierarchical memory network 1810 located between convolutional neural network (CNN) 1804 and long-short time memory (LSTM) 1816. In this example, CNN 1804 and LSTM 1816 are the input and output modules, respectively, for hierarchical

30

memory network 1810, which can also provide local attention capabilities. CNN 1704 can be configured and initialized by architecture configuration 1801. Training data 1802 can provide an input image (e.g., training sample) to CNN 1804. CNN 1704 can process an input image and extract a feature map of the input image.

5 [00206] Hierarchical memory network 1810 includes writer 1808, hierarchical memory 1806, and local-attention based reader 1812. In exemplary embodiments, hierarchical memory 1806 includes memory cells organized into groups and subgroups providing a hierarchical structure for the memory groups. In some embodiments, hierarchical memory 1806 is a fast memory non-volatile device. Writer 1808 writes or
10 stores extracted feature maps from CNN 1804 into hierarchical memory 1806. In some examples, writer 1808 stores feature maps sequentially or in any desired in hierarchical memory 1806. Local attention-based reader 1812 can read the feature maps in hierarchical memory 1806 and apply weights to the feature maps to provide weighted feature maps. In some examples, local attention-based reader 1812 reads only a selected
15 subset of groups of memory related feature maps requiring different weights or with any desired weights. In other examples, local attention-based reader 1812 can implement a search algorithm to exploit hierarchical memory structure of hierarchical memory 1806 to retrieve most relevant information and feature maps. Any type of search algorithm can also be employed for local attention-based reader 1812.

20 [00207] The outputs (e.g., weighted feature maps) of local attention-based reader 1812 can be stored in LSTM 1816 in which Q value estimation 1818 can use stored information in LSTM 1816 and training data 1802 to output a Q value for different actions based on the weighted feature maps to iteration manager 1714. Iteration manager 1814 manages the number of iterations or convergence of applying weights to selected
25 feature maps and deriving resulting Q values (e.g., estimation on how valuable an action is). Iteration manager 1814 can output the final deep Q-network 1820.

[00208] **FIG. 18B** illustrates a flow diagram of an exemplary operation 1850 for the improved DRQN of **FIG. 18A**. In exemplary embodiments, at operation 1852, a current visual frame (input image) is received and a set of feature maps is extracted (e.g.,
30 extracted feature maps from CNN 1804). At operation 1854, feature maps are written (or

stored) in hierarchical memory (e.g., writer 1808 writes feature maps from CNN 1804 into hierarchical memory 1806) within a hierarchical memory network (e.g., hierarchical memory network 1810). At operation 1856, selected subset of feature maps are read from hierarchical memory and soft-attention (e.g., by local-attention based reader 1812) is applied to them to generate weighted feature maps. At operation 1858, the Q value is calculated (e.g., by Q value estimation 1818) for different actions based on the weighted feature maps.. At operation 1860, a decision is made if a convergence or iteration number has been satisfied, and, if yes, the final deep Q-network 1820 at operation 1862. If no, operation 1850 returns to operation 1854 and repeats operations 1854 to 1858.

10

Graphics System Overview

[00209] FIG. 19 is a block diagram of a processing system 1900 according to an exemplary embodiment. In various embodiments, the system 1900 includes one or more processors 1902 and one or more graphics processors 1908, and may be a single processor desktop system, a multiprocessor workstation system, or a server system having a large number of processors 1902 or processor cores 107. In one embodiment, the system 1900 is a processing platform incorporated within a system-on-a-chip (SoC) integrated circuit for use in mobile, handheld, or embedded devices.

[00210] An embodiment of system 1900 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 1900 is a mobile phone, smart phone, tablet computing device or mobile Internet device. Data processing system 1900 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 25 embodiments, data processing system 1900 is a television or set top box device having one or more processors 1902 and a graphical interface generated by one or more graphics processors 1908.

[00211] In some embodiments, the one or more processors 1902 each include one or more processor cores 1907 to process instructions which, when executed, perform 30 operations for system and user software. In some embodiments, each of the one or more

processor cores 1907 is configured to process a specific instruction set 1909. In some embodiments, instruction set 1909 may facilitate Complex Instruction Set Computing (CISC), Reduced Instruction Set Computing (RISC), or computing via a Very Long Instruction Word (VLIW). Multiple processor cores 1907 may each process a different instruction set 1909, which may include instructions to facilitate the emulation of other instruction sets. Processor core 1907 may also include other processing devices, such as Digital Signal Processor (DSP).

[00212] In some embodiments, the processor 1902 includes cache memory 1904. Depending on the architecture, the processor 1902 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 1902. In some embodiments, the processor 1902 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 1907 using known cache coherency techniques. A register file 1906 is additionally included in processor 1902 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 1902.

[00213] In some embodiments, processor 1902 is coupled with a processor bus 1910 to transmit communication signals such as address, data, or control signals between processor 1902 and other components in system 1900. In one embodiment, the system 100 uses an exemplary 'hub' system architecture, including a memory controller hub 1916 and an Input Output (I/O) controller hub 1930. A memory controller hub 1916 facilitates communication between a memory device and other components of system 1900, while an I/O Controller Hub (ICH) 1930 provides connections to I/O devices via a local I/O bus. In one embodiment, the logic of the memory controller hub 1916 is integrated within the processor.

[00214] Memory device 1920 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 1920 can operate as system memory for the system 1900, to store data 1922 and instructions 1921 for use when the one or more processors 1902 executes an application or process. Memory controller hub 1916 also couples with an optional external graphics processor 1912, which may communicate with the one or more graphics processors 1908 in processors 1902 to perform graphics and media operations.

[00215] In some embodiments, ICH 1930 enables peripherals to connect to memory device 1920 and processor 1902 via a high-speed I/O bus. The I/O peripherals include, but are not limited to, an audio controller 1946, a firmware interface 1928, a wireless transceiver 1926 (e.g., Wi-Fi, Bluetooth), a data storage device 1924 (e.g., hard disk drive, flash memory, etc.), and a legacy I/O controller 1940 for coupling legacy (e.g., Personal System 2 (PS/2)) devices to the system. One or more Universal Serial Bus (USB) controllers 1942 connect input devices, such as keyboard and mouse 1944 combinations. A network controller 1934 may also couple with ICH 1930. In some embodiments, a high-performance network controller (not shown) couples with processor bus 1910. It will be appreciated that the system 1900 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 1930 may be integrated within the one or more processor 1902, or the memory controller hub 1916 and I/O controller hub 1930 may be integrated into a discreet external graphics processor, such as the external graphics processor 1912.

[00216] FIG. 20 is a block diagram of an exemplary embodiment of a processor 2000 having one or more processor cores 2002A-2002N, an integrated memory controller 2014, and an integrated graphics processor 2008. Those elements of FIG. 20 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 2000 can include additional cores up to and including additional core 2002N represented by the dashed lined boxes. Each of processor cores 2002A-2002N includes one or more internal cache units 2004A-2004N. In some embodiments, each processor core also has access to one or more shared cached units 2006.

[00217] The internal cache units 2004A-2004N and shared cache units 2006 represent a cache memory hierarchy within the processor 2000. 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),
5 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 2006 and 2004A-2004N.

[00218] In some embodiments, processor 2000 may also include a set of one or more bus controller units 216 and a system agent core 210. The one or more bus
10 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).

[00219] In some embodiments, one or more of the processor cores 2002A-2002N include support for simultaneous multi-threading. In such embodiment, the system agent
15 core 210 includes components for coordinating and operating cores 2002A-2002N 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
20 processor cores 2002A-2002N and graphics processor 2008.

[00220] In some embodiments, processor 2000 additionally includes graphics processor 2008 to execute graphics processing operations. In some embodiments, the graphics processor 2008 couples with the set of shared cache units 2006, and the system agent core 210, including the one or more integrated memory controllers 214. In some
25 embodiments, a display controller 211 is coupled with the graphics processor 2008 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 2008 or system agent core 210.

[00221] In some embodiments, a ring based interconnect unit 2012 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
5 embodiments, graphics processor 208 couples with the ring interconnect 2012 via an I/O link 2013.

[00222] The exemplary I/O link 2013 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
10 embedded memory module 218, such as an eDRAM module. In some embodiments, each of the processor cores 202A-202N and graphics processor 208 use embedded memory modules 218 as a shared Last Level Cache.

[00223] In some embodiments, processor cores 2002A-2002N are homogenous cores executing the same instruction set architecture. In another embodiment, processor
15 cores 2002A-2002N are heterogeneous in terms of instruction set architecture (ISA), where one or more of processor cores 2002A-2002N 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 2002A-2002N are heterogeneous in terms of microarchitecture, where one or more cores having a relatively higher power
20 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.

[00224] FIG. 21 is a block diagram of a graphics processor 2100, which may be a discrete graphics processing unit, or may be a graphics processor integrated with a
25 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 2114 to access memory. Memory interface 314 can be an interface to local memory, one or more internal caches, one or more shared
30 external caches, and/or to system memory.

[00225] In some embodiments, graphics processor 2100 also includes a display controller 2102 to drive display output data to a display device 2120. Display controller 2102 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 2100 includes a video codec engine 306 to encode, decode, or transcode media to, from, or between one or more media encoding formats, including, but not limited to Moving Picture Experts Group (MPEG) formats such as MPEG-2, Advanced Video Coding (AVC) formats such as H.264/MPEG-4 AVC, as well as the Society of Motion Picture & Television Engineers (SMPTE) 421M/VC-1, and Joint Photographic Experts Group (JPEG) formats such as JPEG, and Motion JPEG (MJPEG) formats.

[00226] In some embodiments, graphics processor 2100 includes a block image transfer (BLIT) engine 2104 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, GPE 2110 is a compute engine for performing graphics operations, including three-dimensional (3D) graphics operations and media operations.

[00227] In some embodiments, GPE 2110 includes a 3D pipeline 2112 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 2112 includes programmable and fixed function elements that perform various tasks within the element and/or spawn execution threads to a 3D/Media subsystem 315. While 3D pipeline 2112 can be used to perform media operations, an embodiment of GPE 310 also includes a media pipeline 2116 that is specifically used to perform media operations, such as video post-processing and image enhancement.

[00228] In some embodiments, media pipeline 2116 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 2106. In some embodiments, media pipeline 2116 additionally includes a thread spawning unit to spawn threads for execution on 3D/Media

sub-system 2115. The spawned threads perform computations for the media operations on one or more graphics execution units included in 3D/Media sub-system 2115.

[00229] In some embodiments, 3D/Media subsystem 2115 includes logic for executing threads spawned by 3D pipeline 2112 and media pipeline 2116. In one
5 embodiment, the pipelines send thread execution requests to 3D/Media subsystem 2115, 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 2115 includes one or more internal caches for thread instructions
10 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.

Graphics Processing Engine

[00230] FIG. 22 is a block diagram of a graphics processing engine 2210 of a graphics processor in accordance with some embodiments. In one embodiment, the
15 graphics processing engine (GPE) 2210 is a version of the GPE 2210 shown in FIG. 21. Elements of FIG. 22 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. For example, the 3D pipeline 2212 and media pipeline 2216 of FIG. 3 are illustrated. The media pipeline 2216 is optional in
20 some embodiments of the GPE 2210 and may not be explicitly included within the GPE 410. For example, and in at least one embodiment, a separate media and/or image processor is coupled to the GPE 2210.

[00231] In some embodiments, GPE 2210 couples with or includes a command streamer 2203, which provides a command stream to the 3D pipeline 2112 and/or media
25 pipelines 2116. In some embodiments, command streamer 2203 is coupled with memory, which can be system memory, or one or more of internal cache memory and shared cache memory. In some embodiments, command streamer 2203 receives commands from the memory and sends the commands to 3D pipeline 2112 and/or media pipeline 2116. The commands are directives fetched from a ring buffer, which stores commands for the 3D
30 pipeline 2112 and media pipeline 2116. In one embodiment, the ring buffer can

5 additionally include batch command buffers storing batches of multiple commands. The commands for the 3D pipeline 2112 can also include references to data stored in memory, such as but not limited to vertex and geometry data for the 3D pipeline 2112 and/or image data and memory objects for the media pipeline 2116. The 3D pipeline 2112 and media pipeline 2116 process the commands and data by performing operations via logic within the respective pipelines or by dispatching one or more execution threads to a graphics core array 2214.

10 **[00232]** In various embodiments, the 3D pipeline 2112 can execute one or more shader programs, such as vertex shaders, geometry shaders, pixel shaders, fragment shaders, compute shaders, or other shader programs, by processing the instructions and dispatching execution threads to the graphics core array 2214. The graphics core array 2214 provides a unified block of execution resources. Multi-purpose execution logic (e.g., execution units) within the graphic core array 2214 includes support for various 3D API shader languages and can execute multiple simultaneous execution threads associated with multiple shaders.

15 **[00233]** In some embodiments, the graphics core array 2214 also includes execution logic to perform media functions, such as video and/or image processing. In one embodiment, the execution units additionally include general-purpose logic that is programmable to perform parallel general purpose computational operations, in addition to graphics processing operations. The general-purpose logic can perform processing operations in parallel or in conjunction with general purpose logic within the processor core(s) 1907 of **FIG. 19** or core 2002A-2002N as in **FIG. 20**.

20 **[00234]** Output data generated by threads executing on the graphics core array 2214 can output data to memory in a unified return buffer (URB) 2218. The URB 2218 can store data for multiple threads. In some embodiments, the URB 2218 may be used to send data between different threads executing on the graphics core array 2214. In some embodiments, the URB 2218 may additionally be used for synchronization between threads on the graphics core array and fixed function logic within the shared function logic 2220.

[00235] In some embodiments, graphics core array 2214 is scalable, such that the array includes a variable number of graphics cores, each having a variable number of execution units based on the target power and performance level of GPE 2210. In one embodiment, the execution resources are dynamically scalable, such that execution
5 resources may be enabled or disabled as needed.

[00236] The graphics core array 2214 couples with shared function logic 2220 that includes multiple resources that are shared between the graphics cores in the graphics core array. The shared functions within the shared function logic 2220 are hardware logic units that provide specialized supplemental functionality to the graphics core array
10 2214. In various embodiments, shared function logic 2220 includes but is not limited to sampler 2221, math 2222, and inter-thread communication (ITC) 2223 logic. Additionally, some embodiments implement one or more cache(s) 2225 within the shared function logic 2220. A shared function is implemented where the demand for a given specialized function is insufficient for inclusion within the graphics core array 2214.
15 Instead a single instantiation of that specialized function is implemented as a stand-alone entity in the shared function logic 2220 and shared among the execution resources within the graphics core array 2214. The precise set of functions that are shared between the graphics core array 2214 and included within the graphics core array 2214 varies between embodiments.

20 [00237] **FIG. 23** is a block diagram of another exemplary embodiment of a graphics processor 500. Elements of **FIG. 23** 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.

[00238] In some embodiments, graphics processor 2300 includes a ring
25 interconnect 2302, a pipeline front-end 2304, a media engine 2337, and graphics cores 2380A-2380N. In some embodiments, ring interconnect 2302 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.

[00239] In some embodiments, graphics processor 2300 receives batches of commands via ring interconnect 2302. The incoming commands are interpreted by a command streamer 2303 in the pipeline front-end 2304. In some embodiments, graphics processor 2300 includes scalable execution logic to perform 3D geometry processing and media processing via the graphics core(s) 2380A-2380N. For 3D geometry processing commands, command streamer 2303 supplies commands to geometry pipeline 2336. For at least some media processing commands, command streamer 2303 supplies the commands to a video front end 2334, which couples with a media engine 2337. In some embodiments, media engine 2337 includes a Video Quality Engine (VQE) 2330 for video and image post-processing and a multi-format encode/decode (MFX) 2333 engine to provide hardware-accelerated media data encode and decode. In some embodiments, geometry pipeline 2336 and media engine 2337 each generate execution threads for the thread execution resources provided by at least one graphics core 2380A.

[00240] In some embodiments, graphics processor 2300 includes scalable thread execution resources featuring modular cores 2380A-2380N (sometimes referred to as core slices), each having multiple sub-cores 2350A-2350N, 2360A-2360N (sometimes referred to as core sub-slices). In some embodiments, graphics processor 2300 can have any number of graphics cores 2380A through 2380N. In some embodiments, graphics processor 2300 includes a graphics core 2380A having at least a first sub-core 2350A and a second sub-core 2360A. In other embodiments, the graphics processor is a low power processor with a single sub-core (e.g., 2350A). In some embodiments, graphics processor 2300 includes multiple graphics cores 2380A-2380N, each including a set of first sub-cores 2350A-2350N and a set of second sub-cores 2360A-2360N. Each sub-core in the set of first sub-cores 2350A-2350N includes at least a first set of execution units 2352A-2352N and media/texture samplers 2354A-2354N. Each sub-core in the set of second sub-cores 2360A-2360N includes at least a second set of execution units 2362A-2362N and samplers 2364A-2364N. In some embodiments, each sub-core 2350A-2350N, 2360A-2360N shares a set of shared resources 2370A-2370N. 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.

Execution Units

5 [00241] FIG. 24 illustrates thread execution logic 2400 including an array of processing elements employed in some exemplary embodiments of a GPE. Elements of FIG. 24 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.

10 [00242] In some embodiments, thread execution logic 2400 includes a shader processor 2402, a thread dispatcher 2404, instruction cache 2406, a scalable execution unit array including a plurality of execution units 2408A-2408N, a sampler 2410, a data cache 2412, and a data port 2414. In one embodiment, the scalable execution unit array can dynamically scale by enabling or disabling one or more execution units (e.g., any of execution unit 2408A, 2408B, 2408C, 2408D, through 2408N-1 and 2408N) based on the
15 computational requirements of a workload. 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 2400 includes one or more connections to memory, such as system memory or cache memory, through one or more of instruction cache 2406, data port 2414, sampler 2410, and execution units 2408A-
20 2408N. In some embodiments, each execution unit (e.g. 2408A) is a stand-alone programmable general purpose computational unit that is capable of executing multiple simultaneous hardware threads while processing multiple data elements in parallel for each thread. In various embodiments, the array of execution units 2408A-2408N is scalable to include any number individual execution units.

25 [00243] In some embodiments, the execution units 2408A-2408N are primarily used to execute shader programs. A shader processor 2402 can process the various shader programs and dispatch execution threads associated with the shader programs via a thread dispatcher 2404. In one embodiment, the thread dispatcher includes logic to arbitrate thread initiation requests from the graphics and media pipelines and instantiate
30 the requested threads on one or more execution unit in the execution units 2408A-2408N.

For example, the geometry pipeline (e.g., 2336 of FIG. 23) can dispatch vertex, tessellation, or geometry shaders to the thread execution logic 2400 (FIG. 24) for processing. In some embodiments, thread dispatcher 604 can also process runtime thread spawning requests from the executing shader programs.

5 [00244] In some embodiments, the execution units 2408A-2408N support 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),
10 pixel processing (e.g., pixel shaders, fragment shaders) and general-purpose processing (e.g., compute and media shaders). Each of the execution units 2408A-2408N is capable of multi-issue single instruction multiple data (SIMD) execution and multi-threaded operation enables an efficient execution environment in the face of higher latency memory accesses. Each hardware thread within each execution unit has a dedicated high-
15 bandwidth register file and associated independent thread-state. Execution is multi-issue per clock to pipelines capable of integer, single and double precision floating point operations, SIMD branch capability, logical operations, transcendental operations, and other miscellaneous operations. While waiting for data from memory or one of the shared functions, dependency logic within the execution units 2408A-2408N causes a
20 waiting thread to sleep until the requested data has been returned. While the waiting thread is sleeping, hardware resources may be devoted to processing other threads. For example, during a delay associated with a vertex shader operation, an execution unit can perform operations for a pixel shader, fragment shader, or another type of shader program, including a different vertex shader.

25 [00245] Each execution unit in execution units 2408A-2408N 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.

[00246] The execution unit instruction set includes 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.

[00247] One or more internal instruction caches (e.g., 2406) are included in the thread execution logic 2400 to cache thread instructions for the execution units. In some embodiments, one or more data caches (e.g., 2412) are included to cache thread data during thread execution. In some embodiments, a sampler 2410 is included to provide texture sampling for 3D operations and media sampling for media operations. In some embodiments, sampler 2410 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.

[00248] During execution, the graphics and media pipelines send thread initiation requests to thread execution logic 2400 via thread spawning and dispatch logic. Once a group of geometric objects has been processed and rasterized into pixel data, pixel processor logic (e.g., pixel shader logic, fragment shader logic, etc.) within the shader processor 2402 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, a pixel shader or fragment shader calculates the values of the various vertex attributes that are to be interpolated across the rasterized object. In some embodiments, pixel processor logic within the shader processor 2402 then executes an application programming interface (API)-supplied pixel or fragment shader program. To execute the shader program, the shader processor 2402 dispatches threads to an execution

unit (e.g., 2408A) via thread dispatcher 2404. In some embodiments, pixel shader 2402 uses texture sampling logic in the sampler 2410 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
5 from further processing.

[00249] In some embodiments, the data port 2414 provides a memory access mechanism for the thread execution logic 2400 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 2412) to cache
10 data for memory access via the data port.

[00250] FIG. 25 is a block diagram illustrating a graphics processor instruction formats 2500 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
15 that are only included in a sub-set of the instructions. In some embodiments, instruction format 2500 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.

[00251] In some embodiments, the graphics processor execution units natively support instructions in a 128-bit instruction format 2510. A 64-bit compacted instruction format 2530 is available for some instructions based on the selected instruction, instruction options, and number of operands. The native 128-bit instruction format 2510 provides access to all instruction options, while some options and operations are
20 restricted in the 64-bit instruction format 2530. The native instructions available in the 64-bit instruction format 2530 vary by embodiment. In some embodiments, the instruction is compacted in part using a set of index values in an index field 2513. 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
25 instruction format 2510.
30

[00252] For each format, instruction opcode 2512 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 2514 enables control over certain execution options, such as channels selection (e.g., predication) and data channel order (e.g., swizzle). For instructions in the 128-bit instruction format 2510 an exec-size field 2516 limits the number of data channels that will be executed in parallel. In some embodiments, exec-size field 2516 is not available for use in the 64-bit compact instruction format 2530.

[00253] Some execution unit instructions have up to three operands including two source operands, src0 2520, src1 2522, and one destination 2518. 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 2524), where the instruction opcode 2512 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.

[00254] In some embodiments, the 128-bit instruction format 2510 includes an access/address mode field 2526 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.

[00255] In some embodiments, the 128-bit instruction format 2510 includes an access/address mode field 2526, which specifies an address mode and/or an access mode for the instruction. In one embodiment, the access mode is used to define a data access alignment for the instruction. Some embodiments support access modes including a 16-byte aligned access mode and a 1-byte aligned access mode, where the byte alignment of the access mode determines the access alignment of the instruction operands. For example, when in a first mode, the instruction may use byte-aligned addressing for source

and destination operands and when in a second mode, the instruction may use 16-byte-aligned addressing for all source and destination operands.

[00256] 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 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.

[00257] In some embodiments instructions are grouped based on opcode 2512 bit-fields to simplify Opcode decode 2540. 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 2542 includes data movement and logic instructions (e.g., move (mov), compare (cmp)). In some embodiments, move and logic group 2542 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 2544 (e.g., call, jump (jmp)) includes instructions in the form of 0010xxxxb (e.g., 0x20). A miscellaneous instruction group 2546 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 2548 includes component-wise arithmetic instructions (e.g., add, multiply (mul)) in the form of 0100xxxxb (e.g., 0x40). The parallel math group 2548 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.

Graphics Pipeline

[00258] FIG. 26 is a block diagram of another embodiment of a graphics processor 800. Elements of FIG. 26 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.

[00259] In some embodiments, graphics processor 2600 includes a graphics pipeline 2620, a media pipeline 2630, a display engine 2640, thread execution logic 2650, and a render output pipeline 2670. In some embodiments, graphics processor 2600 is a graphics processor within a multi-core processing system that includes one or more
5 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 2600 via a ring interconnect 2602. In some embodiments, ring interconnect 802 couples graphics processor 2600 to other processing components, such as other graphics processors or general-purpose processors. Commands from ring interconnect
10 802 are interpreted by a command streamer 2603, which supplies instructions to individual components of graphics pipeline 2620 or media pipeline 2630.

[00260] In some embodiments, command streamer 2603 directs the operation of a vertex fetcher 2605 that reads vertex data from memory and executes vertex-processing commands provided by command streamer 2603. In some embodiments, vertex fetcher
15 805 provides vertex data to a vertex shader 2607, which performs coordinate space transformation and lighting operations to each vertex. In some embodiments, vertex fetcher 805 and vertex shader 2607 execute vertex-processing instructions by dispatching execution threads to execution units 2652A-2652B via a thread dispatcher 2631.

[00261] In some embodiments, execution units 2652A-2652B are an array of
20 vector processors having an instruction set for performing graphics and media operations. In some embodiments, execution units 2652A-2652B have an attached L1 cache 2651 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.

[00262] In some embodiments, graphics pipeline 2620 includes tessellation
25 components to perform hardware-accelerated tessellation of 3D objects. In some embodiments, a programmable hull shader 2611 configures the tessellation operations. A programmable domain shader 2617 provides back-end evaluation of tessellation output. A tessellator 2613 operates at the direction of hull shader 2611 and contains special
30 purpose logic to generate a set of detailed geometric objects based on a coarse geometric

model that is provided as input to graphics pipeline 2620. In some embodiments, if tessellation is not used, tessellation components (e.g., hull shader 2611, tessellator 2613, and domain shader 2617) can be bypassed.

[00263] In some embodiments, complete geometric objects can be processed by a geometry shader 2619 via one or more threads dispatched to execution units 2652A-2652B, or can proceed directly to the clipper 2629. 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 2619 receives input from the vertex shader 2607. In some embodiments, geometry shader 2619 is programmable by a geometry shader program to perform geometry tessellation if the tessellation units are disabled.

[00264] Before rasterization, a clipper 2629 processes vertex data. The clipper 2629 may be a fixed function clipper or a programmable clipper having clipping and geometry shader functions. In some embodiments, a rasterizer and depth test component 2673 in the render output pipeline 2670 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 2650. In some embodiments, an application can bypass the rasterizer and depth test component 2673 and access un-rasterized vertex data via a stream out unit 2623.

[00265] The graphics processor 2600 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 2652A-2652B and associated cache(s) 2651, texture and media sampler 2654, and texture/sampler cache 2658 interconnect via a data port 2656 to perform memory access and communicate with render output pipeline components of the processor. In some embodiments, sampler 2654, caches 2651, 2658 and execution units 2652A-2652B each have separate memory access paths.

[00266] In some embodiments, render output pipeline 2670 contains a rasterizer and depth test component 2673 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 2678 and depth cache 2679 are also available in some embodiments. A pixel operations component 2677 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 2641, or substituted at display time by the display controller 2643 using overlay display planes. In some embodiments, a shared L3 cache 2675 is available to all graphics components, allowing the sharing of data without the use of main system memory.

[00267] In some embodiments, graphics processor media pipeline 2630 includes a media engine 2637 and a video front end 2634. In some embodiments, video front end 2634 receives pipeline commands from the command streamer 2603. In some embodiments, media pipeline 2630 includes a separate command streamer. In some embodiments, video front-end 2634 processes media commands before sending the command to the media engine 2637. In some embodiments, media engine 2637 includes thread spawning functionality to spawn threads for dispatch to thread execution logic 2650 via thread dispatcher 2631.

[00268] In some embodiments, graphics processor 2600 includes a display engine 840. In some embodiments, display engine 2640 is external to processor 2600 and couples with the graphics processor via the ring interconnect 2602, or some other interconnect bus or fabric. In some embodiments, display engine 2640 includes a 2D engine 2641 and a display controller 2643. In some embodiments, display engine 2640 contains special purpose logic capable of operating independently of the 3D pipeline. In some embodiments, display controller 2643 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.

[00269] In some embodiments, graphics pipeline 2620 and media pipeline 2630 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), Open Computing Language (OpenCL), and/or Vulkan graphics and compute API, all from the Khronos Group. In some embodiments, support may also be provided for the Direct3D library from the Microsoft Corporation. In some embodiments, a combination of these libraries may be supported. 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.

10

Graphics Pipeline Programming

[00270] FIG. 27A is a block diagram illustrating a graphics processor command format 2700 according to some embodiments. FIG. 27B is a block diagram illustrating a graphics processor command sequence 2710 according to an embodiment. The solid lined boxes in FIG. 27A 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 2700 of FIG. 27A includes data fields to identify a target client 2702 of the command, a command operation code (opcode) 2704, and the relevant data 2706 for the command. A sub-opcode 2705 and a command size 2708 are also included in some commands.

20

[00271] In some embodiments, client 2702 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 2704 and, if present, sub-opcode 2705 to determine the operation to perform. The client unit performs the command using information in data field 2706. For some commands an explicit command size 908 is

30

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.

5 [00272] The flow diagram in FIG. 27B shows an exemplary graphics processor command sequence 2710. 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
10 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.

[00273] In some embodiments, the graphics processor command sequence 910
15 may begin with a pipeline flush command 2712 to cause any active graphics pipeline to complete the currently pending commands for the pipeline. In some embodiments, the 3D pipeline 2722 and the media pipeline 2724 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
20 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 2712 can be used for pipeline synchronization or before placing the graphics processor into a low power state.

25 [00274] In some embodiments, a pipeline select command 2713 is used when a command sequence requires the graphics processor to explicitly switch between pipelines. In some embodiments, a pipeline select command 2713 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 2712 is
30 required immediately before a pipeline switch via the pipeline select command 2713.

[00275] In some embodiments, a pipeline control command 2714 configures a graphics pipeline for operation and is used to program the 3D pipeline 2722 and the media pipeline 2724. In some embodiments, pipeline control command 2714 configures the pipeline state for the active pipeline. In one embodiment, the pipeline control
5 command 2714 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.

[00276] In some embodiments, commands for the return buffer state 2716 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
10 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, configuring the return buffer state 2716 includes selecting the size and number of return buffers to use for a set of pipeline operations.

[00277] The remaining commands in the command sequence differ based on the active pipeline for operations. Based on a pipeline determination 2720, the command
15 sequence is tailored to the 3D pipeline 2722 beginning with the 3D pipeline state 2730 or the media pipeline 2724 beginning at the media pipeline state 2740.

[00278] The commands to configure the 3D pipeline state 930 include 3D state
20 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 on the particular 3D API in use. In some embodiments, 3D pipeline state 2730 commands are also able to selectively disable or bypass certain pipeline elements if those
25 elements will not be used.

[00279] In some embodiments, 3D primitive 2732 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 2732 command are forwarded to the vertex fetch function in the graphics pipeline. The vertex fetch function uses the 3D
30 primitive 2732 command data to generate vertex data structures. The vertex data

structures are stored in one or more return buffers. In some embodiments, 3D primitive 2732 command is used to perform vertex operations on 3D primitives via vertex shaders. To process vertex shaders, 3D pipeline 2722 dispatches shader execution threads to graphics processor execution units.

5 [00280] In some embodiments, 3D pipeline 2722 is triggered via an execute 2734 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
10 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.

[00281] In some embodiments, the graphics processor command sequence 2710
15 follows the media pipeline 2724 path when performing media operations. In general, the specific use and manner of programming for the media pipeline 2724 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
20 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.

25 [00282] In some embodiments, media pipeline 2724 is configured in a similar manner as the 3D pipeline 2722. A set of commands to configure the media pipeline state 2740 are dispatched or placed into a command queue before the media object commands 2742. In some embodiments, commands for the media pipeline state 2740 include data to configure the media pipeline elements that will be used to process the
30 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, commands for the media pipeline state 2740 also support the use of one or more pointers to “indirect” state elements that contain a batch of state settings.

[00283] In some embodiments, media object commands 2742 supply pointers to
5 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 2742. Once the pipeline state is configured and media object commands 2742 are queued, the media pipeline 924 is triggered via an execute command 2744 or an equivalent execute event (e.g., register
10 write). Output from media pipeline 2724 may then be post processed by operations provided by the 3D pipeline 2722 or the media pipeline 2724. In some embodiments, GPGPU operations are configured and executed in a similar manner as media operations.

15 Graphics Software Architecture

[00284] FIG. 28 illustrates exemplary graphics software architecture for a data processing system 2800 according to some embodiments. In some embodiments, software architecture includes a 3D graphics application 2810, an operating system 2820, and at least one processor 2830. In some embodiments, processor 2830 includes a
20 graphics processor 2832 and one or more general-purpose processor core(s) 2834. The graphics application 2810 and operating system 2820 each execute in the system memory 1050 of the data processing system.

[00285] In some embodiments, 3D graphics application 2810 contains one or more
25 shader programs including shader instructions 2812. 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 2814 in a machine language suitable for execution by the general-purpose processor core 2834. The application also includes graphics objects 1016 defined by vertex data.

[00286] In some embodiments, operating system 2820 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. The operating system 1020 can support a graphics API 2822 such as the
5 Direct3D API, the OpenGL API, or the Vulkan API. When the Direct3D API is in use, the operating system 2820 uses a front-end shader compiler 2824 to compile any shader instructions 2812 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
10 compilation of the 3D graphics application 2810. In some embodiments, the shader instructions 2812 are provided in an intermediate form, such as a version of the Standard Portable Intermediate Representation (SPIR) used by the Vulkan API.

[00287] In some embodiments, user mode graphics driver 2826 contains a back-end shader compiler 2827 to convert the shader instructions 2812 into a hardware specific
15 representation. When the OpenGL API is in use, shader instructions 2812 in the GLSL high-level language are passed to a user mode graphics driver 2826 for compilation. In some embodiments, user mode graphics driver 2826 uses operating system kernel mode functions 2828 to communicate with a kernel mode graphics driver 2829. In some embodiments, kernel mode graphics driver 1029 communicates with graphics processor
20 2832 to dispatch commands and instructions.

IP Core Implementations

[00288] 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
25 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
30 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.

5 [00289] FIG. 29 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
10 2930 can generate a software simulation 2910 of an IP core design in a high level programming language (e.g., C/C++). The software simulation 2910 can be used to design, test, and verify the behavior of the IP core using a simulation model 2912. The simulation model 2912 may include functional, behavioral, and/or timing simulations. A register transfer level (RTL) design 2915 can then be created or synthesized from the
15 simulation model 2912. The RTL design 2915 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 2915, 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
20 simulation may vary.

[00290] The RTL design 2915 or equivalent may be further synthesized by the design facility into a hardware model 2920, 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
25 for delivery to a 3rd party fabrication facility 2965 using non-volatile memory 2940 (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 2950 or wireless connection 2960. The fabrication facility 2965 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.

Exemplary System on a Chip Integrated Circuit

[00291] FIGS. 30-32 illustrate exemplary integrated circuits and associated
5 graphics processors that may be fabricated using one or more IP cores, according to various embodiments described herein. In addition to what is illustrated, other logic and circuits may be included, including additional graphics processors/cores, peripheral interface controllers, or general purpose processor cores.

[00292] FIG. 30 is a block diagram illustrating an exemplary system on a chip
10 integrated circuit 3000 that may be fabricated using one or more IP cores, according to an embodiment. Exemplary integrated circuit 1200 includes one or more application processor(s) 3005 (e.g., CPUs), at least one graphics processor 3010, and may additionally include an image processor 3015 and/or a video processor 3020, any of which may be a modular IP core from the same or multiple different design facilities.
15 Integrated circuit 3000 includes peripheral or bus logic including a USB controller 1225, UART controller 3030, an SPI/SDIO controller 3035, and an I²S/I²C controller 3040. Additionally, the integrated circuit can include a display device 3045 coupled to one or more of a high-definition multimedia interface (HDMI) controller 1250 and a mobile industry processor interface (MIPI) display interface 3055. Storage may be provided by a
20 flash memory subsystem 3060 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 3070.

[00293] FIG. 31 is a block diagram illustrating an exemplary graphics processor
25 3110 of a system on a chip integrated circuit that may be fabricated using one or more IP cores, according to an embodiment. Graphics processor 3110 can be a variant of the graphics processor 3010 of FIG. 30. Graphics processor 3110 includes a vertex processor 3105 and one or more fragment processor(s) 3115A-3115N (e.g., 3115A, 3115B, 3115C, 3115D, through 3115N-1, and 3115N). Graphics processor 3110 can
30 execute different shader programs via separate logic, such that the vertex processor 3105

is optimized to execute operations for vertex shader programs, while the one or more fragment processor(s) 3115A-3115N execute fragment (e.g., pixel) shading operations for fragment or pixel shader programs. The vertex processor 3105 performs the vertex processing stage of the 3D graphics pipeline and generates primitives and vertex data.

5 The fragment processor(s) 3115A-3115N use the primitive and vertex data generated by the vertex processor 3105 to produce a framebuffer that is displayed on a display device. In one embodiment, the fragment processor(s) 3115A-3115N are optimized to execute fragment shader programs as provided for in the OpenGL API, which may be used to perform similar operations as a pixel shader program as provided for in the Direct 3D
10 API.

[00294] Graphics processor 3110 additionally includes one or more memory management units (MMUs) 3120A-3120B, cache(s) 3125A-3125B, and circuit interconnect(s) 3130A-3130B. The one or more MMU(s) 3120A-3120B provide for virtual to physical address mapping for graphics processor 3110, including for the vertex
15 processor 3105 and/or fragment processor(s) 3115A-3115N, which may reference vertex or image/texture data stored in memory, in addition to vertex or image/texture data stored in the one or more cache(s) 3125A-3125B. In one embodiment, the one or more MMU(s) 3120A-3120B may be synchronized with other MMUs within the system, including one or more MMUs associated with the one or more application processor(s) 3005, image
20 processor 3015, and/or video processor 3020 of FIG. 30, such that each processor 3005-3020 can participate in a shared or unified virtual memory system. The one or more circuit interconnect(s) 3130A-3130B enable graphics processor 3110 to interface with other IP cores within the SoC, either via an internal bus of the SoC or via a direct connection, according to embodiments.

25 [00295] FIG. 32 is a block diagram illustrating an additional exemplary graphics processor 3210 of a system on a chip integrated circuit that may be fabricated using one or more IP cores, according to an embodiment. Graphics processor 3210 can be a variant of the graphics processor 3010 of FIG. 30. Graphics processor 3210 includes the one or more MMU(s) 3120A-3120B, cache(s) 3125A-3125B, and circuit interconnect(s) 3130A-
30 1330B of the integrated circuit 3100 of FIG. 31.

[00296] Graphics processor 3210 includes one or more shader core(s) 3215A-3215N (e.g., 3215A, 3215B, 3215C, 3215D, 3215E, 3215F, through 3215N-1, and 3115N), which provides for a unified shader core architecture in which a single core or type or core can execute all types of programmable shader code, including shader
5 program code to implement vertex shaders, fragment shaders, and/or compute shaders. The exact number of shader cores present can vary among embodiments and implementations. Additionally, graphics processor 3210 includes an inter-core task manager 3205, which acts as a thread dispatcher to dispatch execution threads to one or more shader core(s) 3215A-3215N and a tiling unit 3218 to accelerate tiling operations
10 for tile-based rendering, in which rendering operations for a scene are subdivided in image space, for example to exploit local spatial coherence within a scene or to optimize use of internal caches.

[00297] Embodiments of the present invention include methods and systems for budgeted and simplified training of deep neural networks (DNN).

15 [00298] In one example, a method for a deep neural network (DNN) includes sub-sampling a training image into a plurality of training sub-images. A number of sub-images are randomly selected. A DNN is trained with the randomly selected number of training sub-images to obtain a training result.

[00299] In one example, a method includes sub-sampling a testing image into a
20 plurality of testing sub-images. A first testing sub-image is randomly selected. The trained DNN is tested with the randomly selected first testing image to obtain a first testing result.

[00300] In one example, a method includes randomly selecting a second testing sub-image. The trained DNN is tested with the randomly selected second testing image
25 to obtain a second testing result. The first testing result and second testing result are combined in selecting one of the first testing result and second testing result as a final result.

[00301] In one example, for a method, each training sub-image has a smaller resolution than the training image.

[00302] In one example, for a method, each testing sub-image has a smaller resolution than the testing image.

[00303] In one example, a system for a server includes a processing core, an I/O controller hub, and a graphics processor. The processing core has a deep neural network (DNN). The I/O controller hub is coupled to the processing core and provides network, data storage, and DNN access for the processing core. The graphics processor is coupled to the I/O controller hub and trains the DNN using a plurality of training sub-images derived from a down-sampled training image. The graphics processor also tests the trained DNN using a plurality of testing sub-images derived from a down-sampled testing image.

[00304] In one example, the graphics processor sub-samples a testing image into a plurality of testing sub-images. The graphics processor also randomly selects a first testing sub-image. The graphics processor also tests the trained DNN with the randomly selected first testing image to obtain a first testing result.

[00305] In one example, the graphics processor randomly selects a second testing sub-image. The graphics processor also tests the trained DNN with the randomly selected second testing image to obtain a second testing result. The graphics processor also combines the first testing result and second testing result in selecting one of the first testing result and second testing result as a final result.

[00306] In one example, for the system, each training sub-image has a smaller resolution than the training image.

[00307] In one example, for the system, each testing sub-image has a smaller resolution than the testing image.

[00308] In one example, a method for a recurrent deep Q- network (RDQN) having a local attention mechanism located between a convolutional neural network (CNN) and a long-short time memory (LSTM) includes generating a plurality of feature maps by the CNN from an input image. Hard-attention is applied by the local attention mechanism to the generated plurality of feature maps by selecting a subset of the generated feature maps. Soft-attention is applied by the local attention mechanism to the selected subset of generated feature maps by providing weights to the selected subset of

generated feature maps in obtaining weighted feature maps. The weighted feature maps are stored in the LSTM. A Q value is calculated for different actions based on the weighted feature maps stored in the LSTM.

5 [00309] In one example, a method includes determining if an iteration condition is satisfied. A final deep Q-network is output if the iteration condition is satisfied.

[00310] In one example, a method includes writing generated feature maps into a hierarchical memory.

[00311] In one example, a method includes reading the selected subset of generated feature maps in the hierarchical memory.

10 [00312] In one example, a system for a server includes a processing core, an I/O hub controller, and a graphics processor. The processing core has a recurrent deep Q-network (RDQN) including a convolutional neural network (CNN), and long-short time memory (LSTM), wherein a local attention mechanism is located between the CNN and LSTM, and wherein the CNN and LSTM are input and output modules for the local
15 attention mechanism having soft-attention and hard-attention. The I/O hub controller is coupled to the processing core and provides network, data storage, LSTM, RDQN, CNN, and local attention mechanism access for the processing core. The graphics processor is coupled to the I/O hub controller and processes an input image and a plurality of feature maps based on the input image. The graphics processor also applies hard-attention to
20 generated plurality of feature maps by selecting a subset of the generated feature maps. The graphics processor also applies soft-attention to the selected subset of generated feature maps by providing weights to the selected subset of generated feature maps in obtaining weighted feature maps. The graphics processor also stores the weighted feature maps in the LSTM. The graphics processor also calculates a Q value for different actions
25 based on the weighted feature maps stored in the LSTM.

[00313] In one example, the graphics processor determines if an iteration condition is satisfied, and outputs a final deep Q-network if the iteration condition is satisfied.

[00314] In one example, the graphics processor writes generated feature maps into a hierarchical memory.

[00315] In one example, the graphics processor reads the selected subset of generated feature maps in the hierarchical memory.

[00316] The foregoing description and drawings are to be regarded in an illustrative rather than a restrictive sense. Persons skilled in the art will understand that
5 various modifications and changes may be made to the embodiments described herein without departing from the broader spirit and scope of the invention as set forth in the appended claims.

CLAIMS

What is claimed is:

1. A method for a deep neural network (DNN), comprising:
sub-sampling a training image into a plurality of training sub-images;
5 randomly selecting a number of sub-images; and
training a DNN with the randomly selected number of training sub-images to
obtain a training result.
2. The method of claim 1, further comprising:
10 sub-sampling a testing image into a plurality of testing sub-images;
randomly selecting a first testing sub-image; and
testing the trained DNN with the randomly selected first testing image to obtain a
first testing result.
3. The method of claim 2, further comprising:
15 randomly selecting a second testing sub-image;
testing the trained DNN with the randomly selected second testing image to
obtain a second testing result; and
combining the first testing result and second testing result in selecting one of the
20 first testing result and second testing result as a final result.
4. The method of claim 1, wherein each training sub-image has a smaller
resolution than the training image.
5. The method of claim 2, wherein each testing sub-image has a smaller
25 resolution than the testing image.
6. A system for a server comprising:
a processing core having a deep neural network (DNN);

an I/O controller hub coupled to the processing core and to provide network, data storage, and DNN access for the processing core; and
a graphics processor coupled to the I/O controller hub and to
train the DNN using a plurality of training sub-images derived
5 from a down-sampled training image, and
test the trained DNN using a plurality of testing sub-images
derived from a down-sampled testing image.

7. The system of claim 6, wherein the graphics processor is to
10 sub-sample a testing image into a plurality of testing sub-images,
randomly select a first testing sub-image, and
test the trained DNN with the randomly selected first testing image to obtain a
first testing result.

8. The system of claim 7, wherein the graphics processor is to
15 randomly select a second testing sub-image,
test the trained DNN with the randomly selected second testing image to obtain a
second testing result, and
combine the first testing result and second testing result in selecting one of the
20 first testing result and second testing result as a final result.

9. The system of claim 6, wherein each training sub-image has a smaller
resolution than the training image.

10. The system of claim 7, wherein each testing sub-image has a smaller
25 resolution than the testing image.

11. In a recurrent deep Q- network (RDQN) having a local attention
mechanism located between a convolutional neural network (CNN) and a long-short time
30 memory (LSTM), a method comprising:

generating a plurality of feature maps by the CNN from an input image;
applying hard-attention by the local attention mechanism to the generated
plurality of feature maps by selecting a subset of the generated feature maps;
applying soft-attention by the local attention mechanism to the selected subset of
5 generated feature maps by providing weights to the selected subset of generated feature
maps in obtaining weighted feature maps;
storing the weighted feature maps in the LSTM; and
calculating a Q value for different actions based on the weighted feature maps
stored in the LSTM.

10

12. The method of claim 11, further comprising:
determining if an iteration condition is satisfied; and
outputting a final deep Q-network if the iteration condition is satisfied.

15

13. The method of claim 11, wherein applying hard-attention includes writing
generated feature maps into a hierarchical memory.

14. The method of claim 13, wherein applying soft-attention includes reading
the selected subset of generated feature maps in the hierarchical memory.

20

15. A system for a server comprising:
a processing core having a recurrent deep Q-network (RDQN) including a
convolutional neural network (CNN), and long-short time memory (LSTM), wherein a
local attention mechanism is located between the CNN and LSTM, and wherein the CNN
25 and LSTM are input and output modules for the local attention mechanism having soft-
attention and hard-attention;

an I/O hub controller coupled to the processing core and to provide network, data
storage, LSTM, RDQN, CNN, and local attention mechanism access for the processing
core; and

30

a graphics processor coupled to the I/O hub controller and to

process an input image and a plurality of feature maps based on the input image,
apply hard-attention to generated plurality of feature maps by selecting a subset of the generated feature maps,
5 apply soft-attention to the selected subset of generated feature maps by providing weights to the selected subset of generated feature maps in obtaining weighted feature maps.
store the weighted feature maps in the LSTM, and
calculate a Q value for different actions based on the weighted feature maps stored in the LSTM.
10

16. The system of claim 15, wherein the graphics processor is to determine if an iteration condition is satisfied, and output a final deep Q-network if the iteration condition is satisfied.

15

17. The system of claim 15, wherein the graphics processor is to write generated feature maps into a hierarchical memory.

18. The system of claim 17, wherein the graphics processor is to read the selected subset of generated feature maps in the hierarchical memory.
20

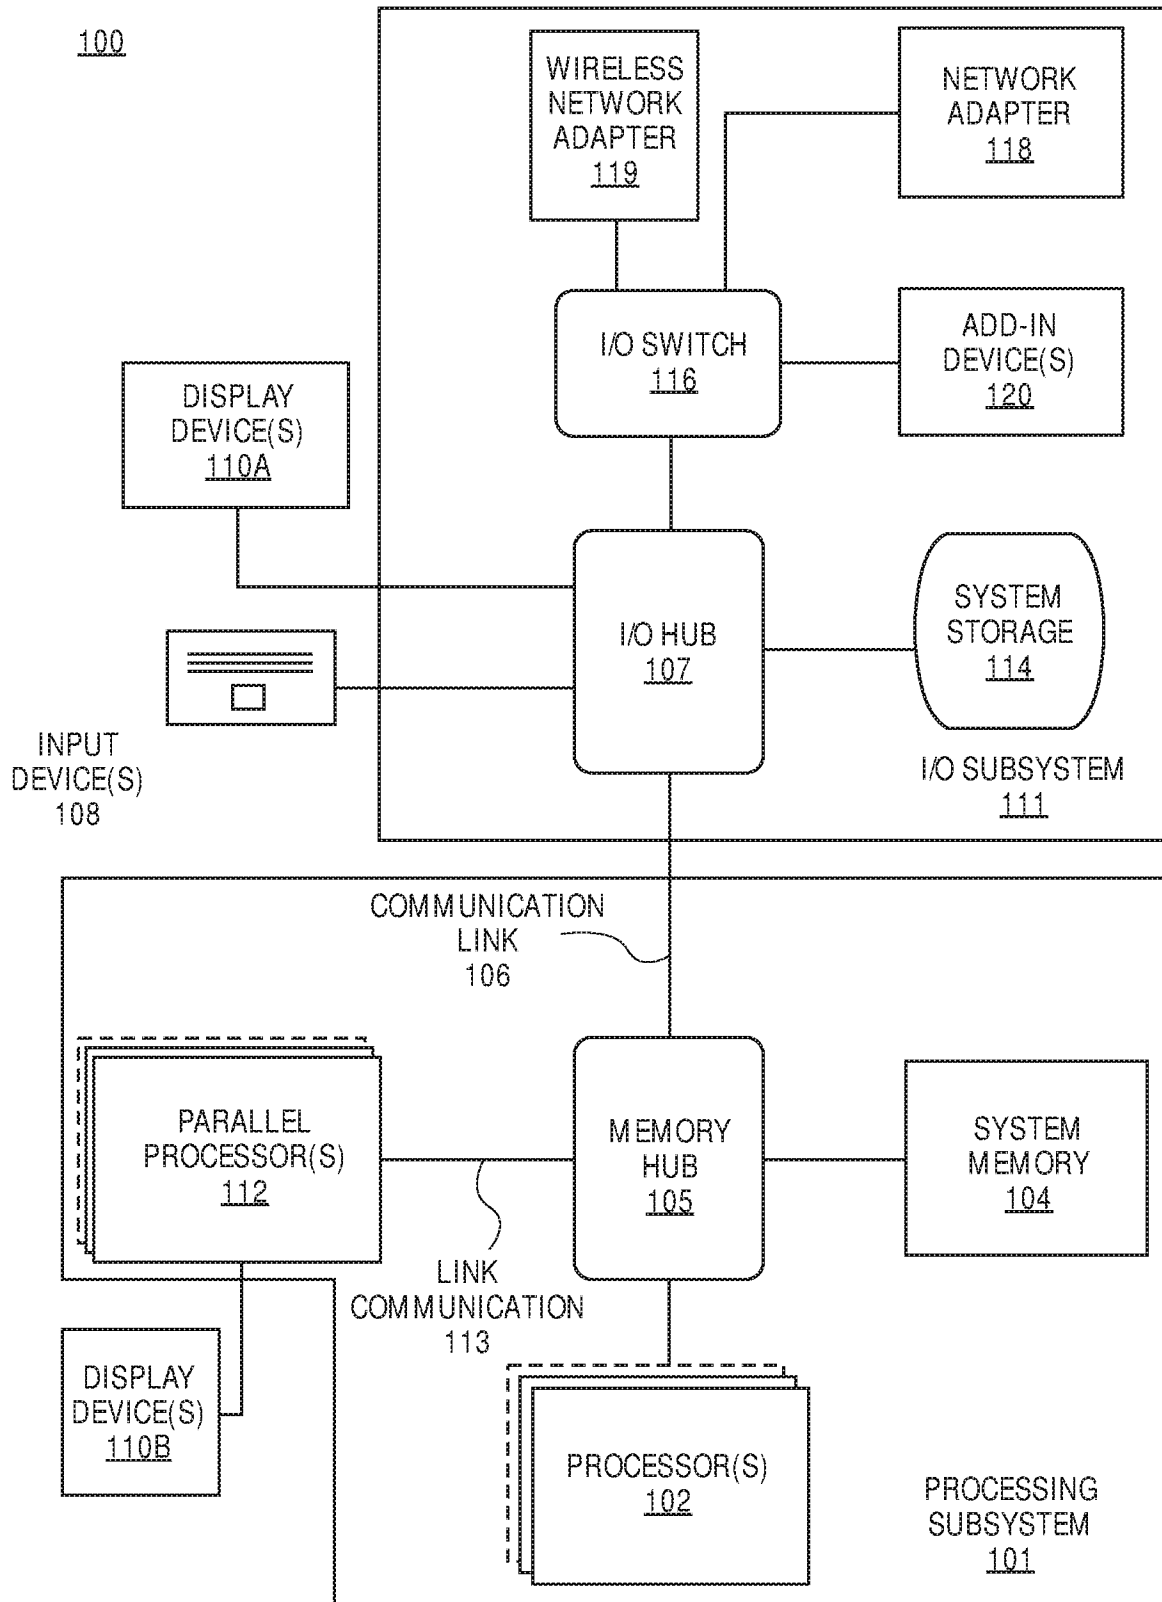


FIG. 1




FIG. 2A

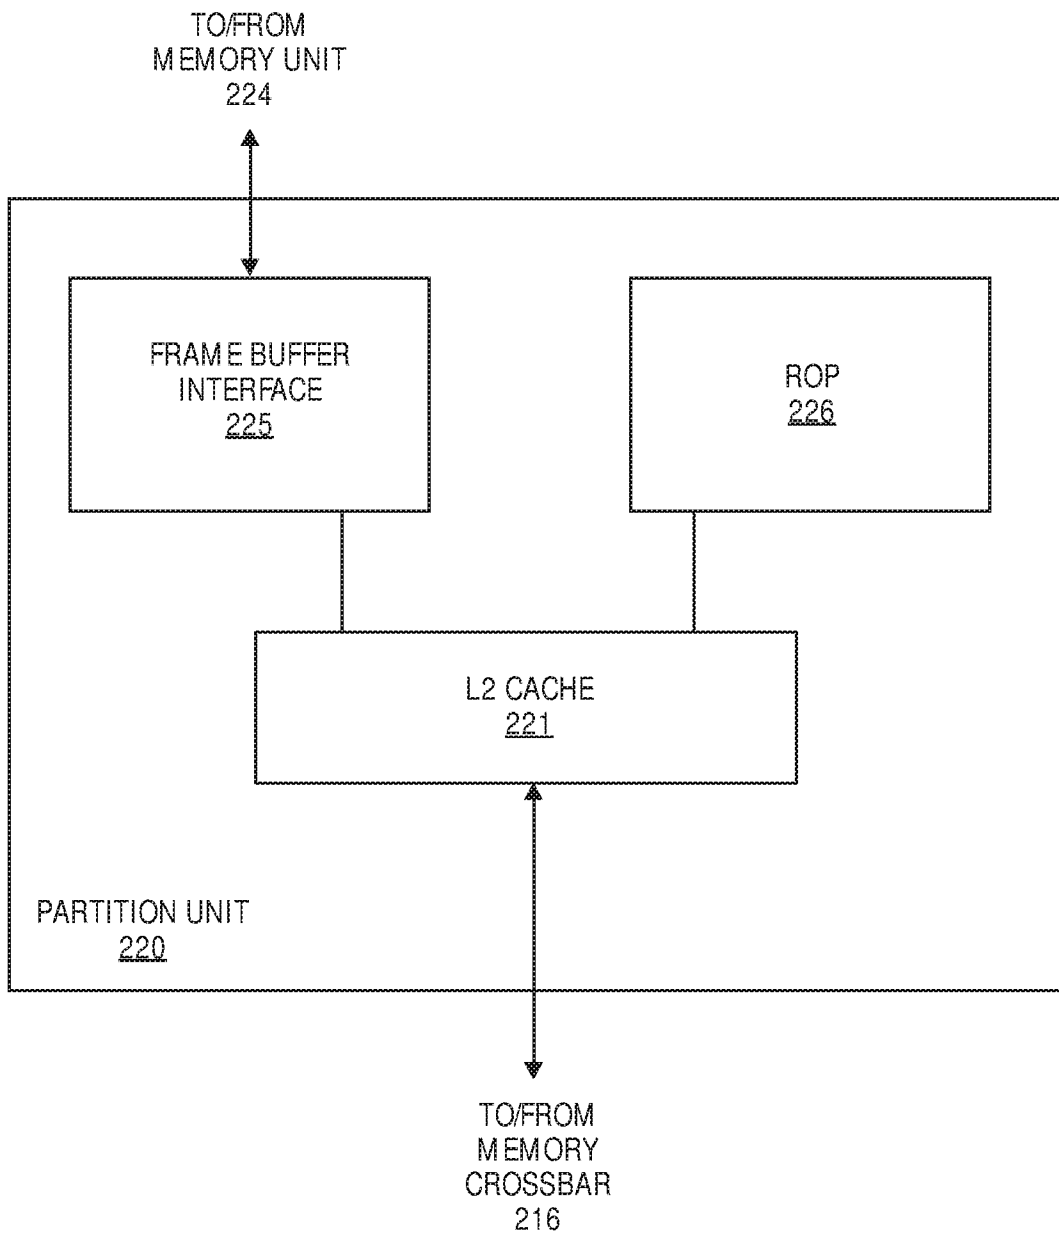


FIG. 2B

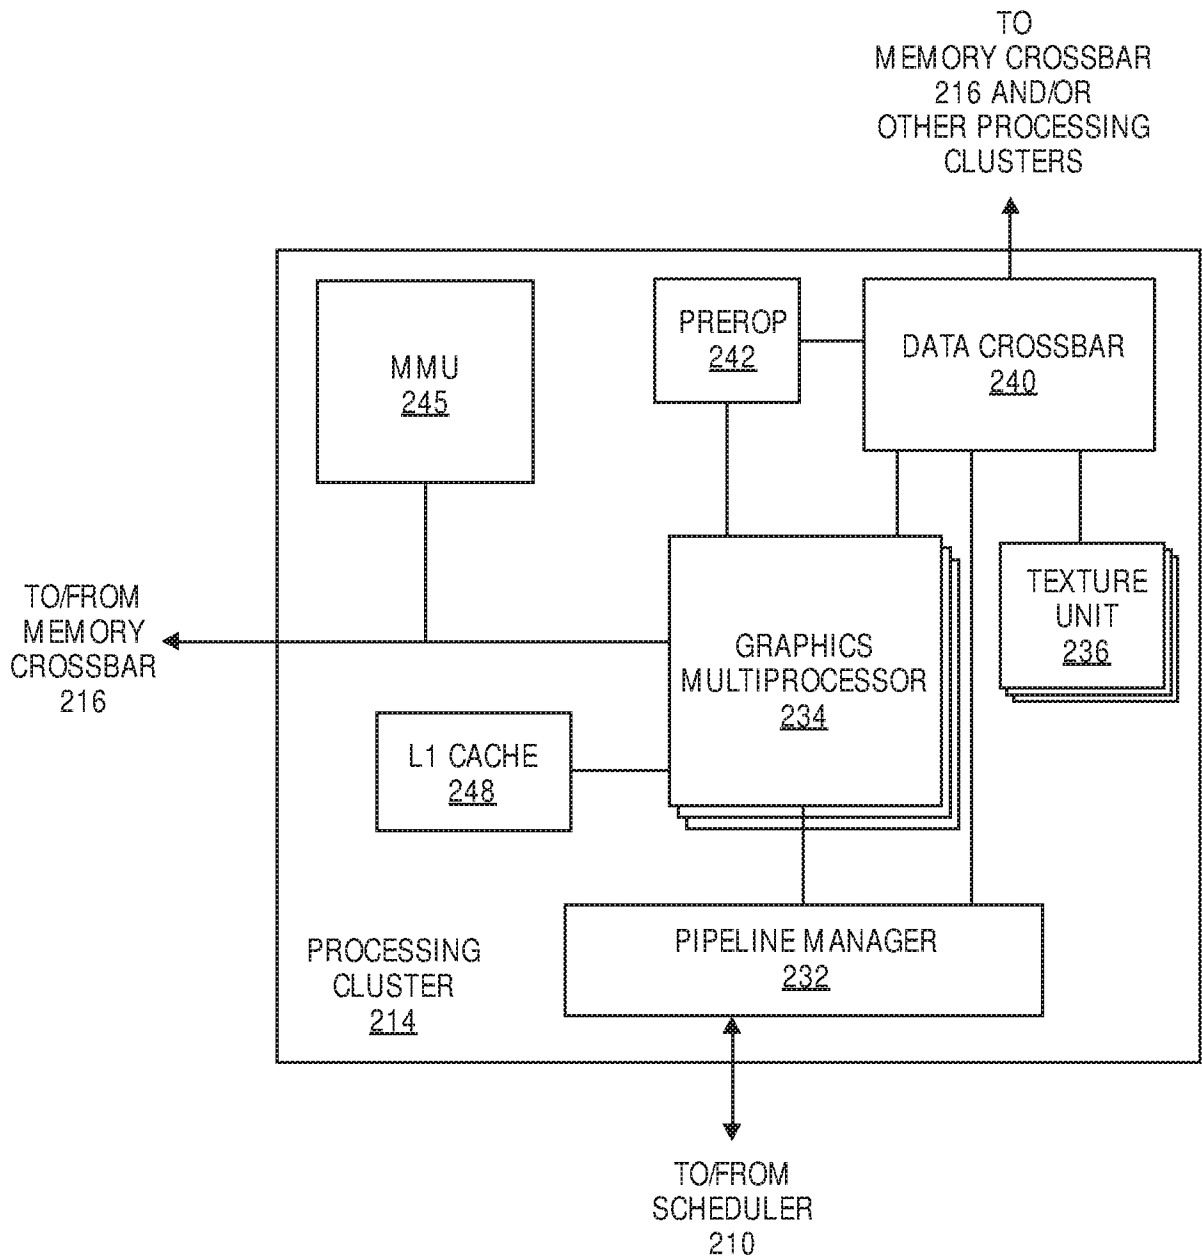


FIG. 2C

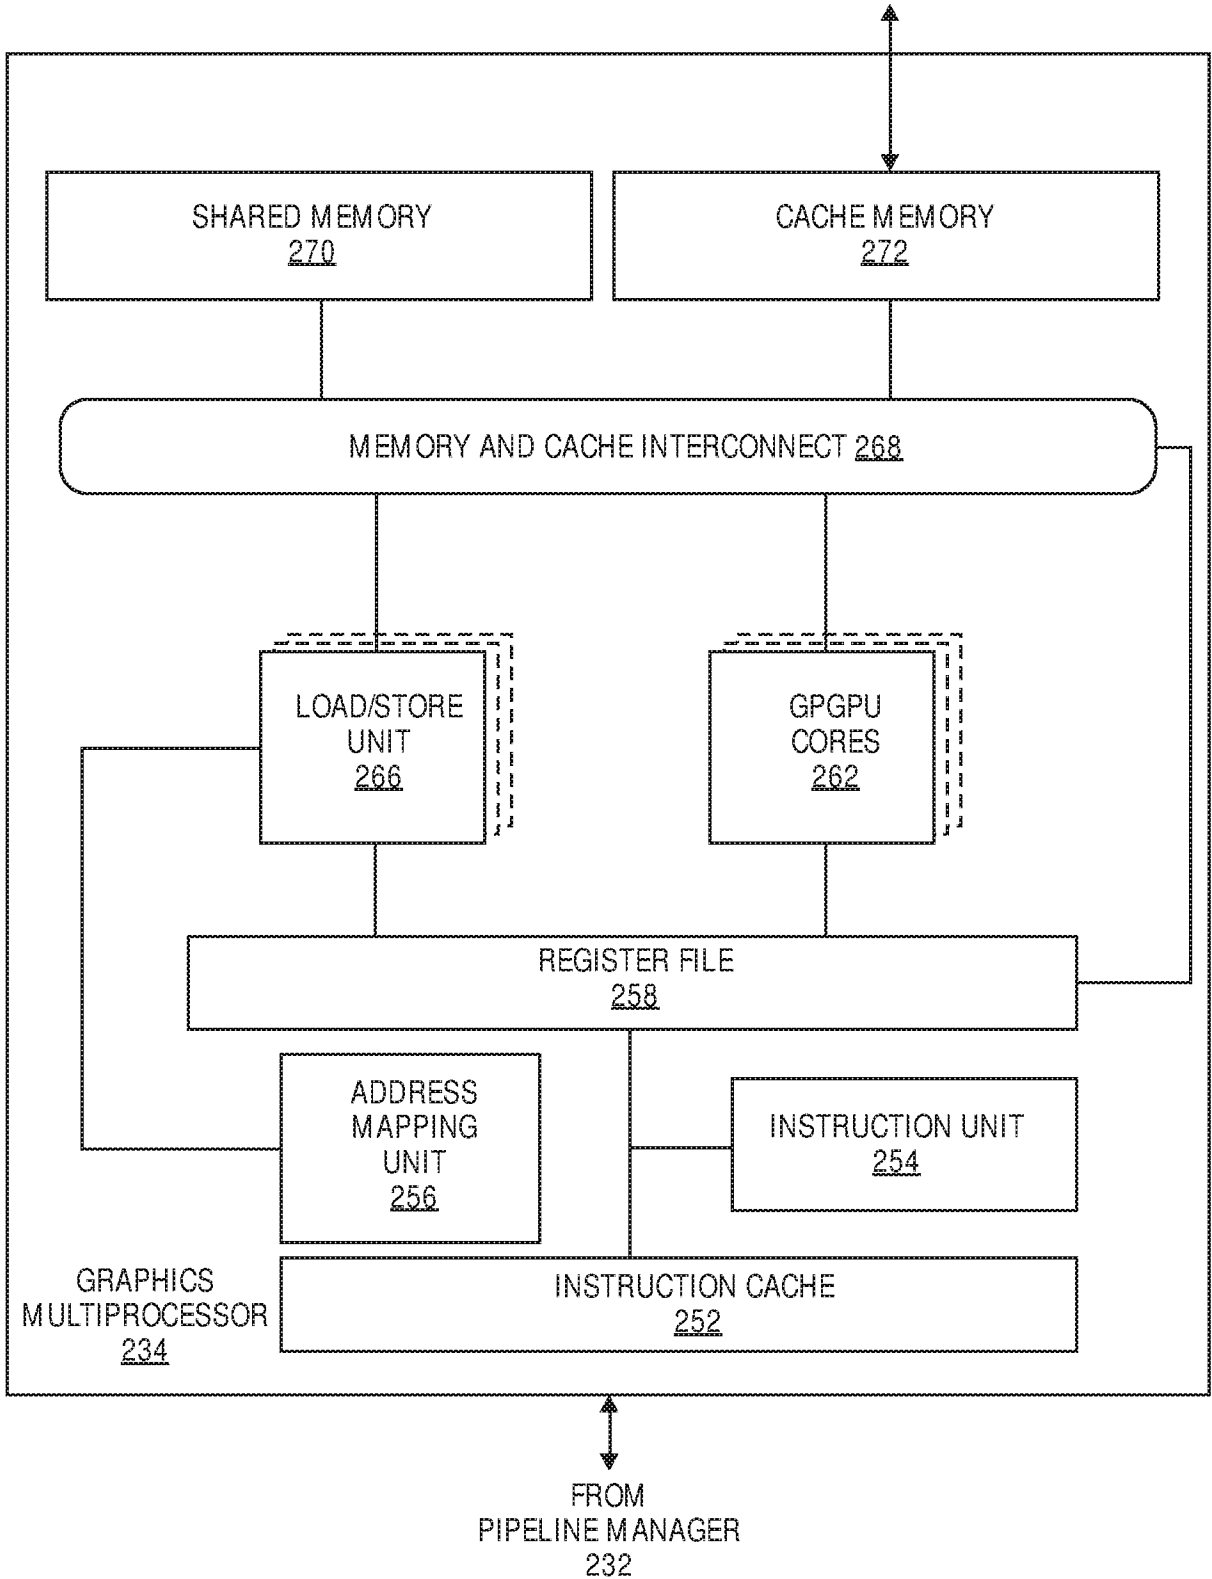


FIG. 2D

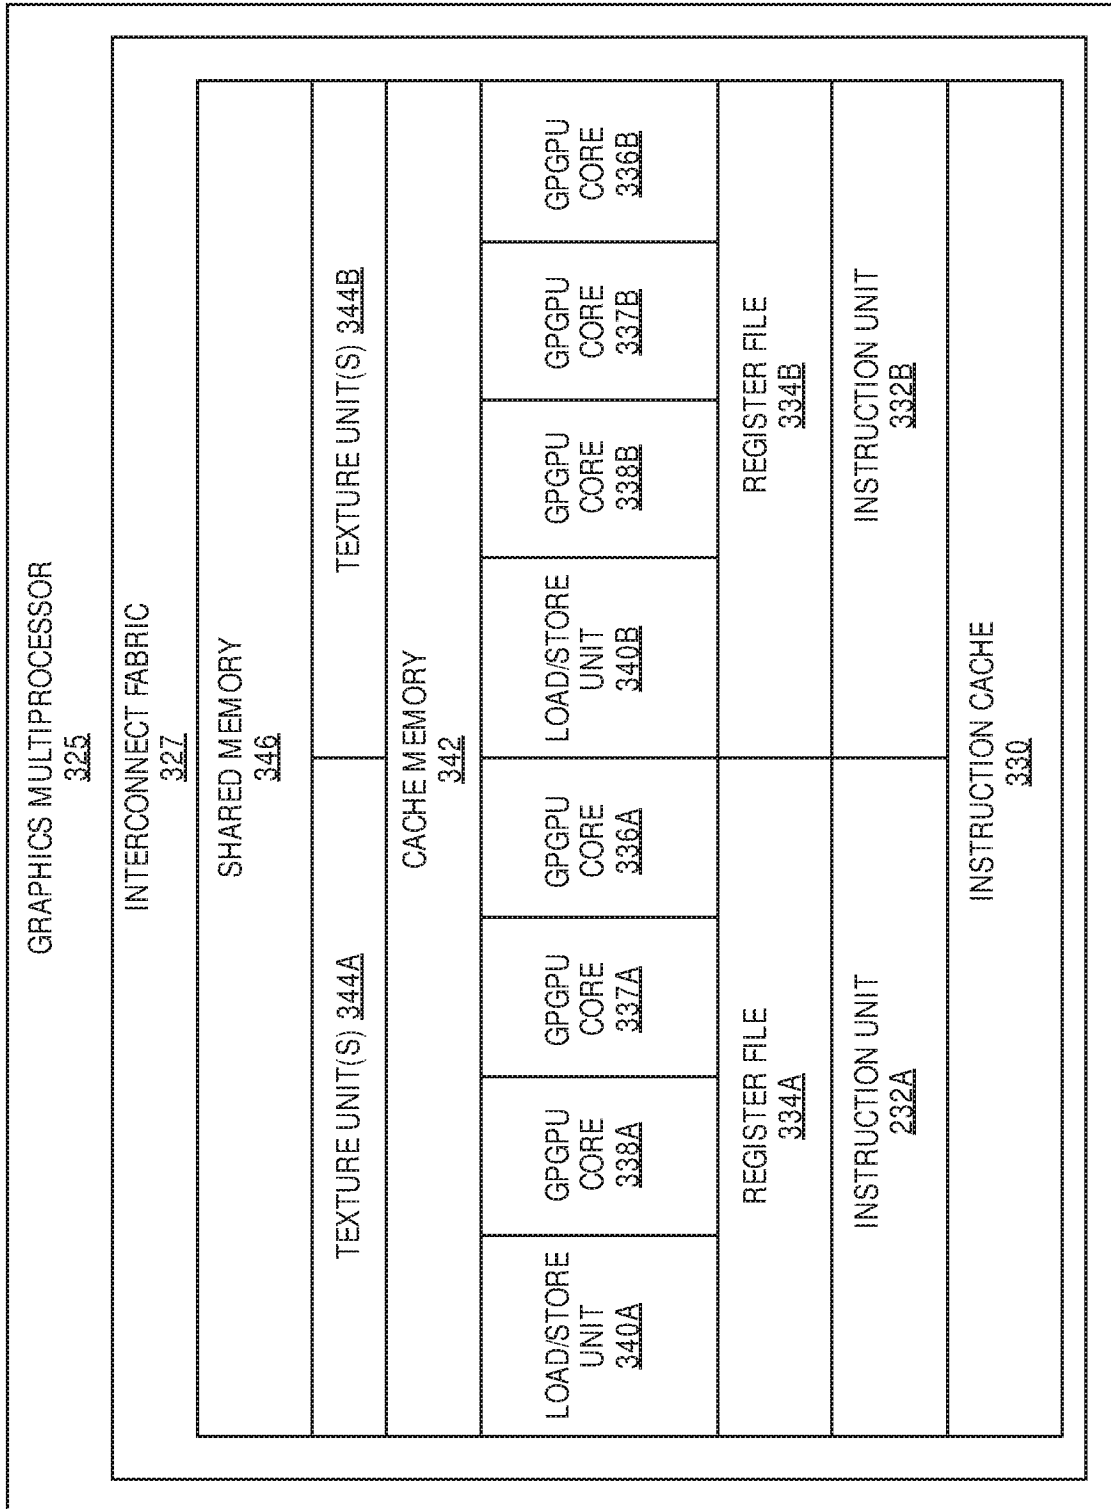


FIG. 3A

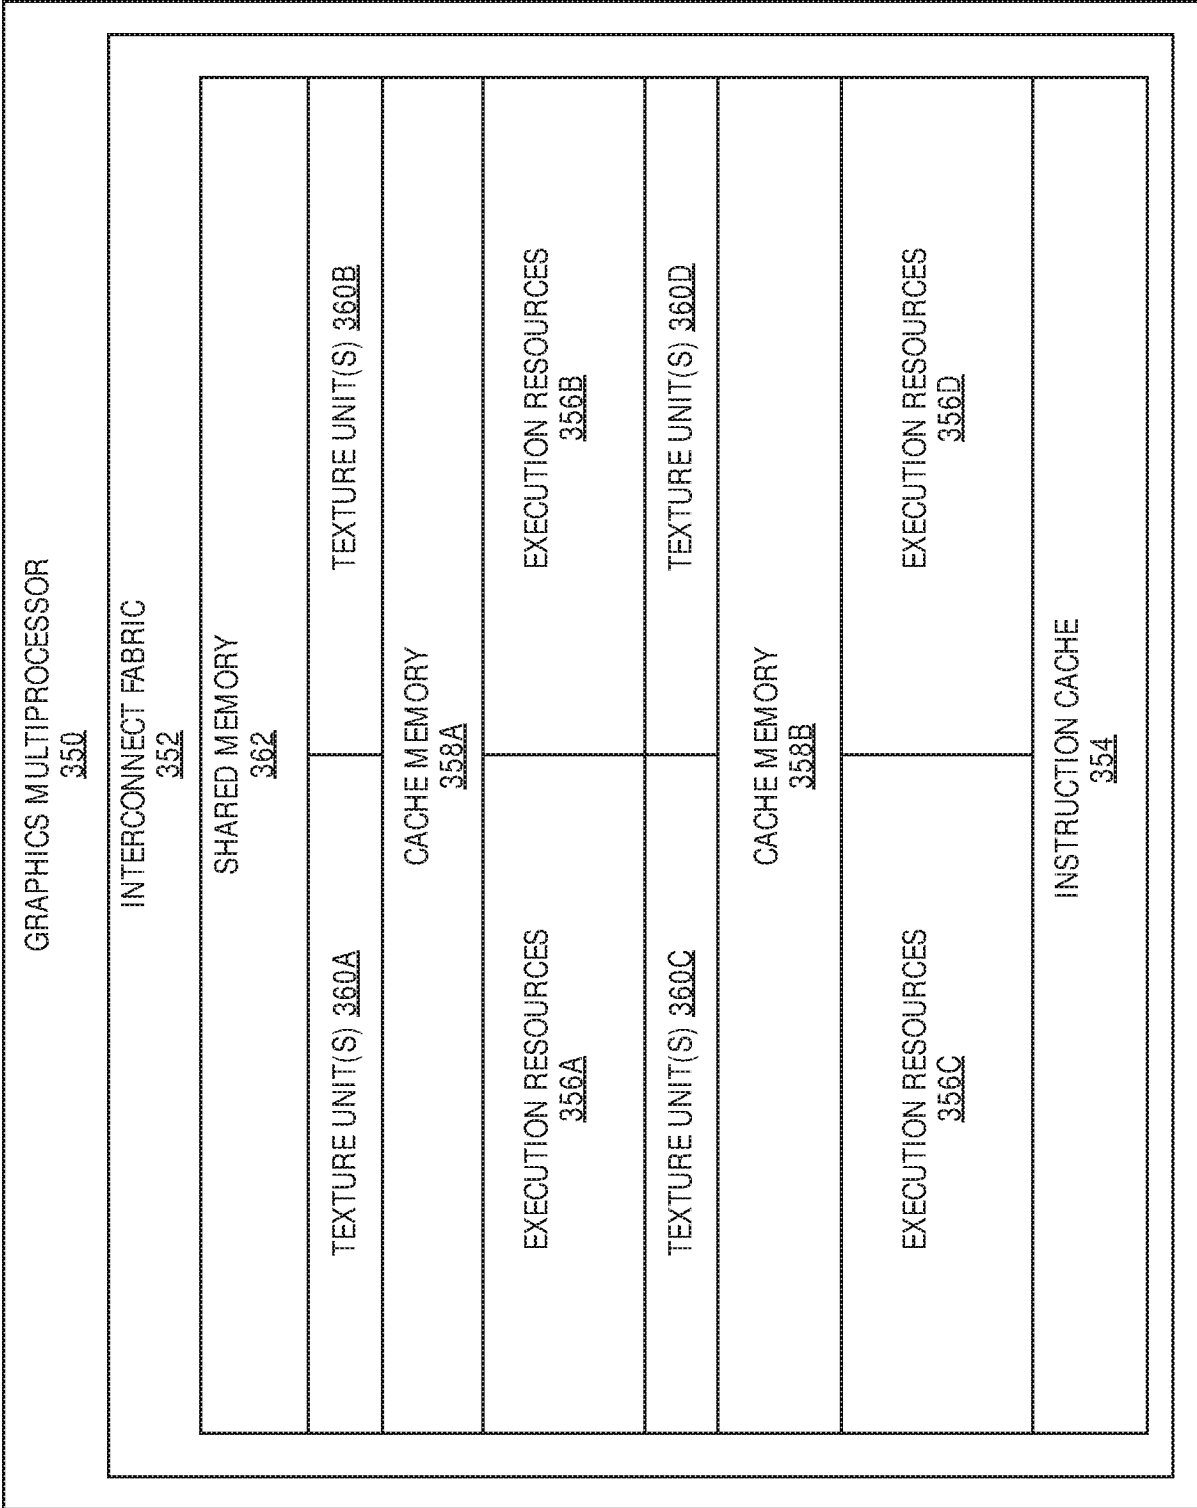


FIG. 3B




FIG. 4A




FIG. 4B




FIG. 4C

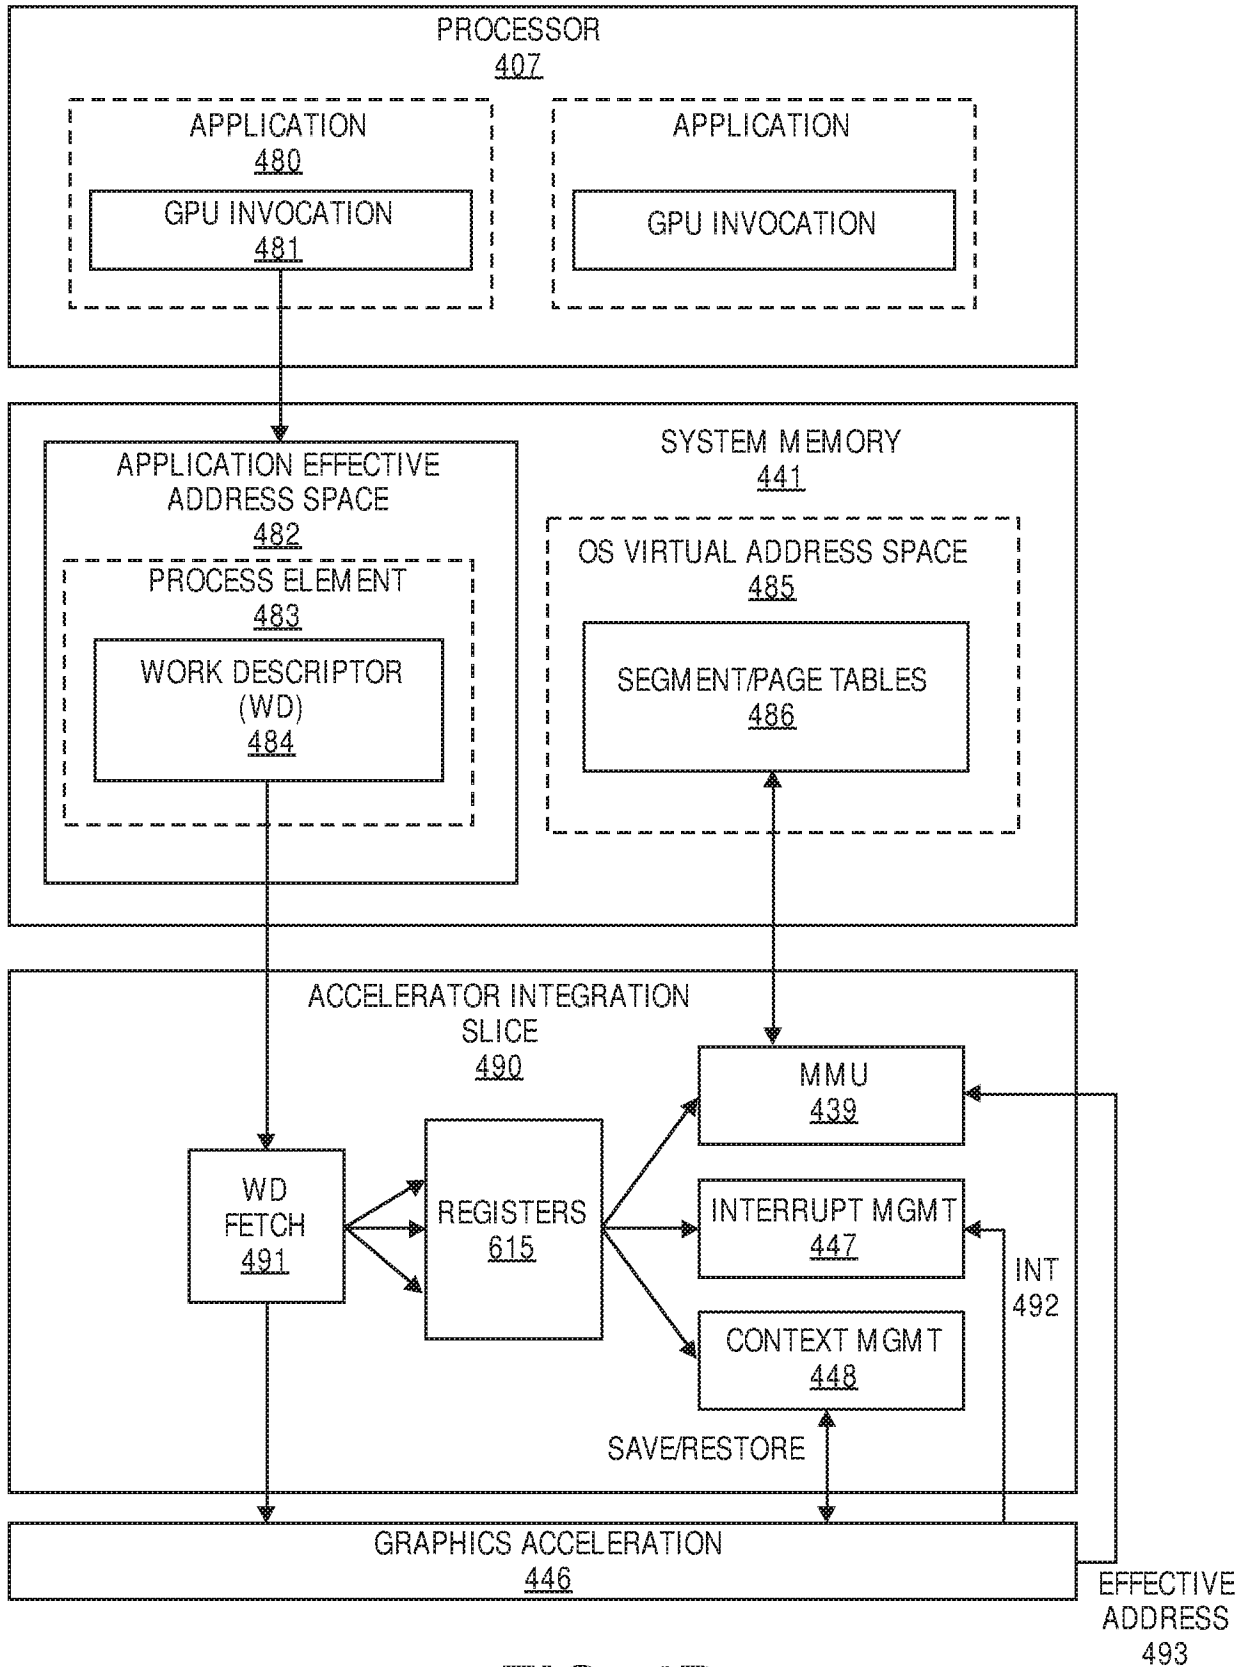


FIG. 4D

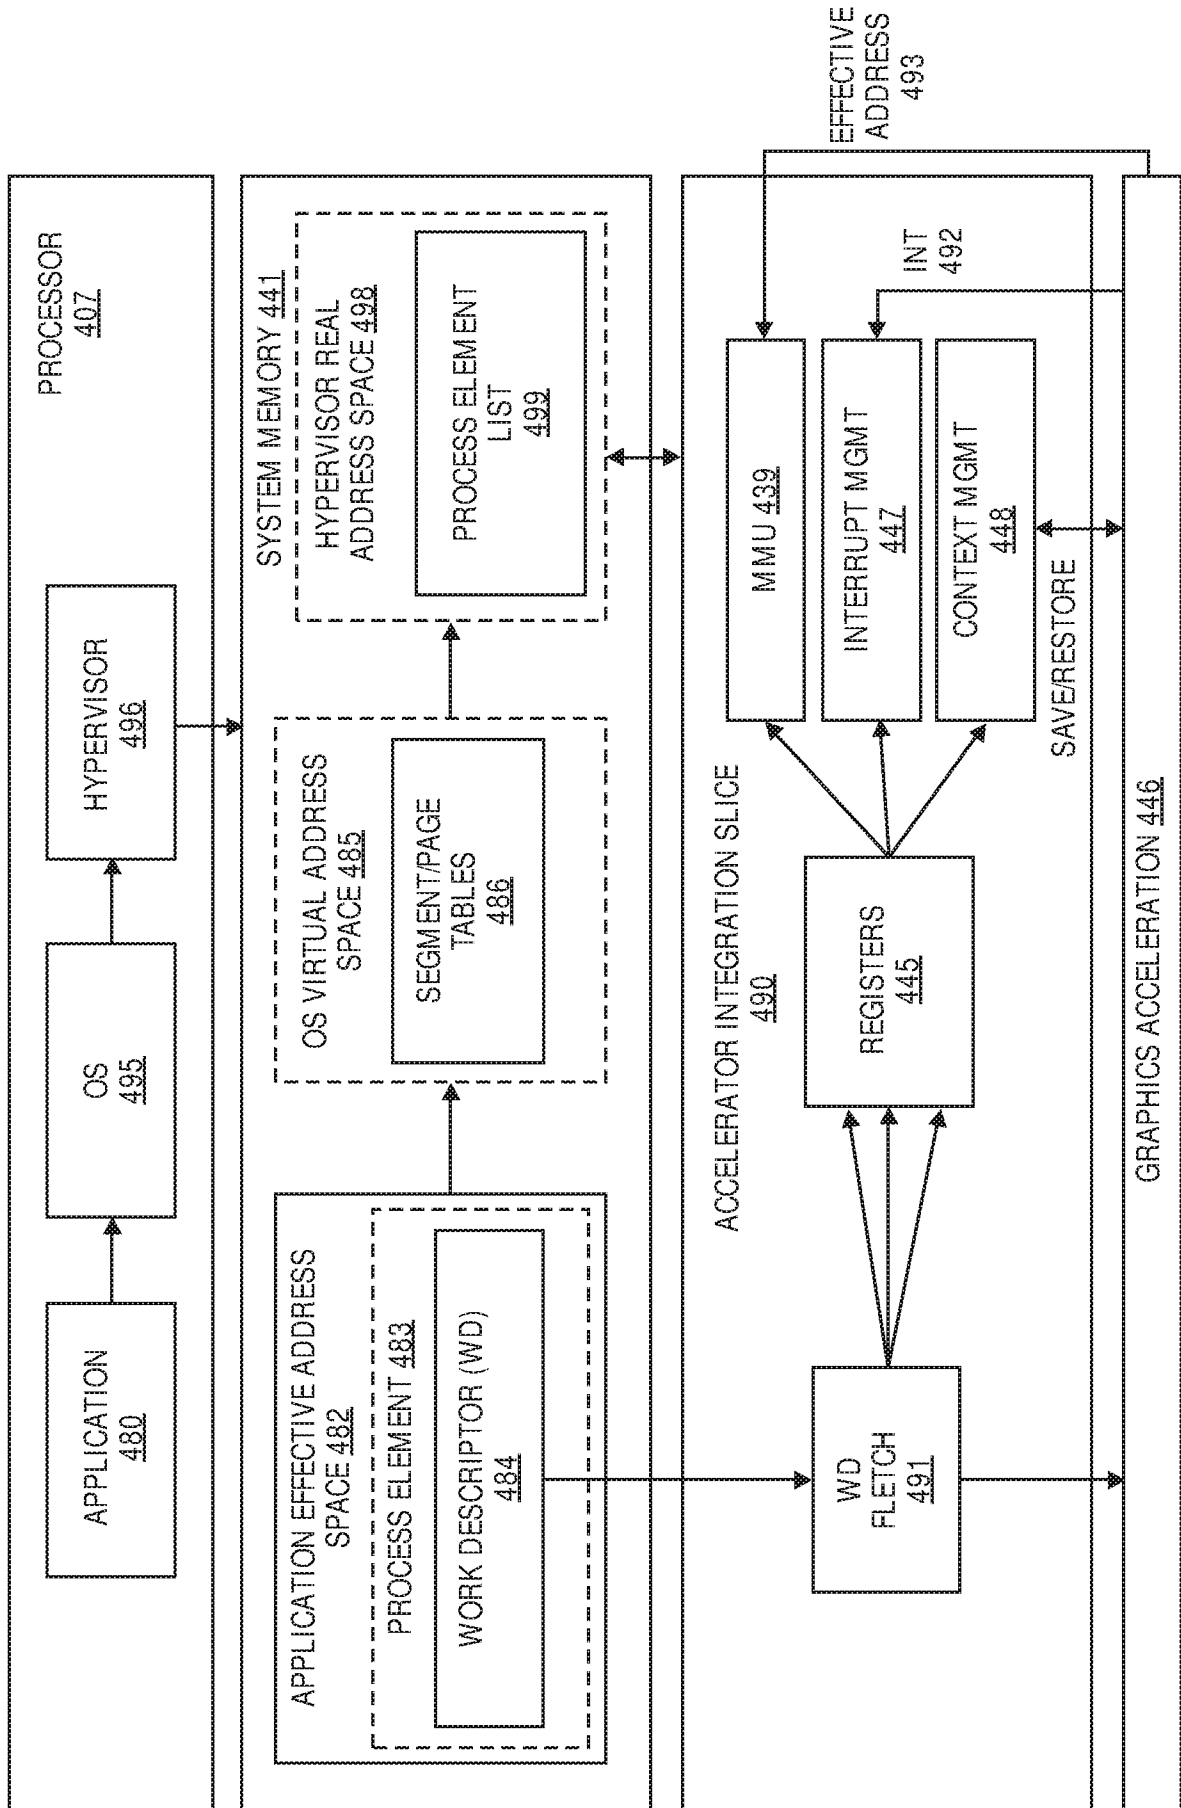


FIG. 4E




FIG. 4F

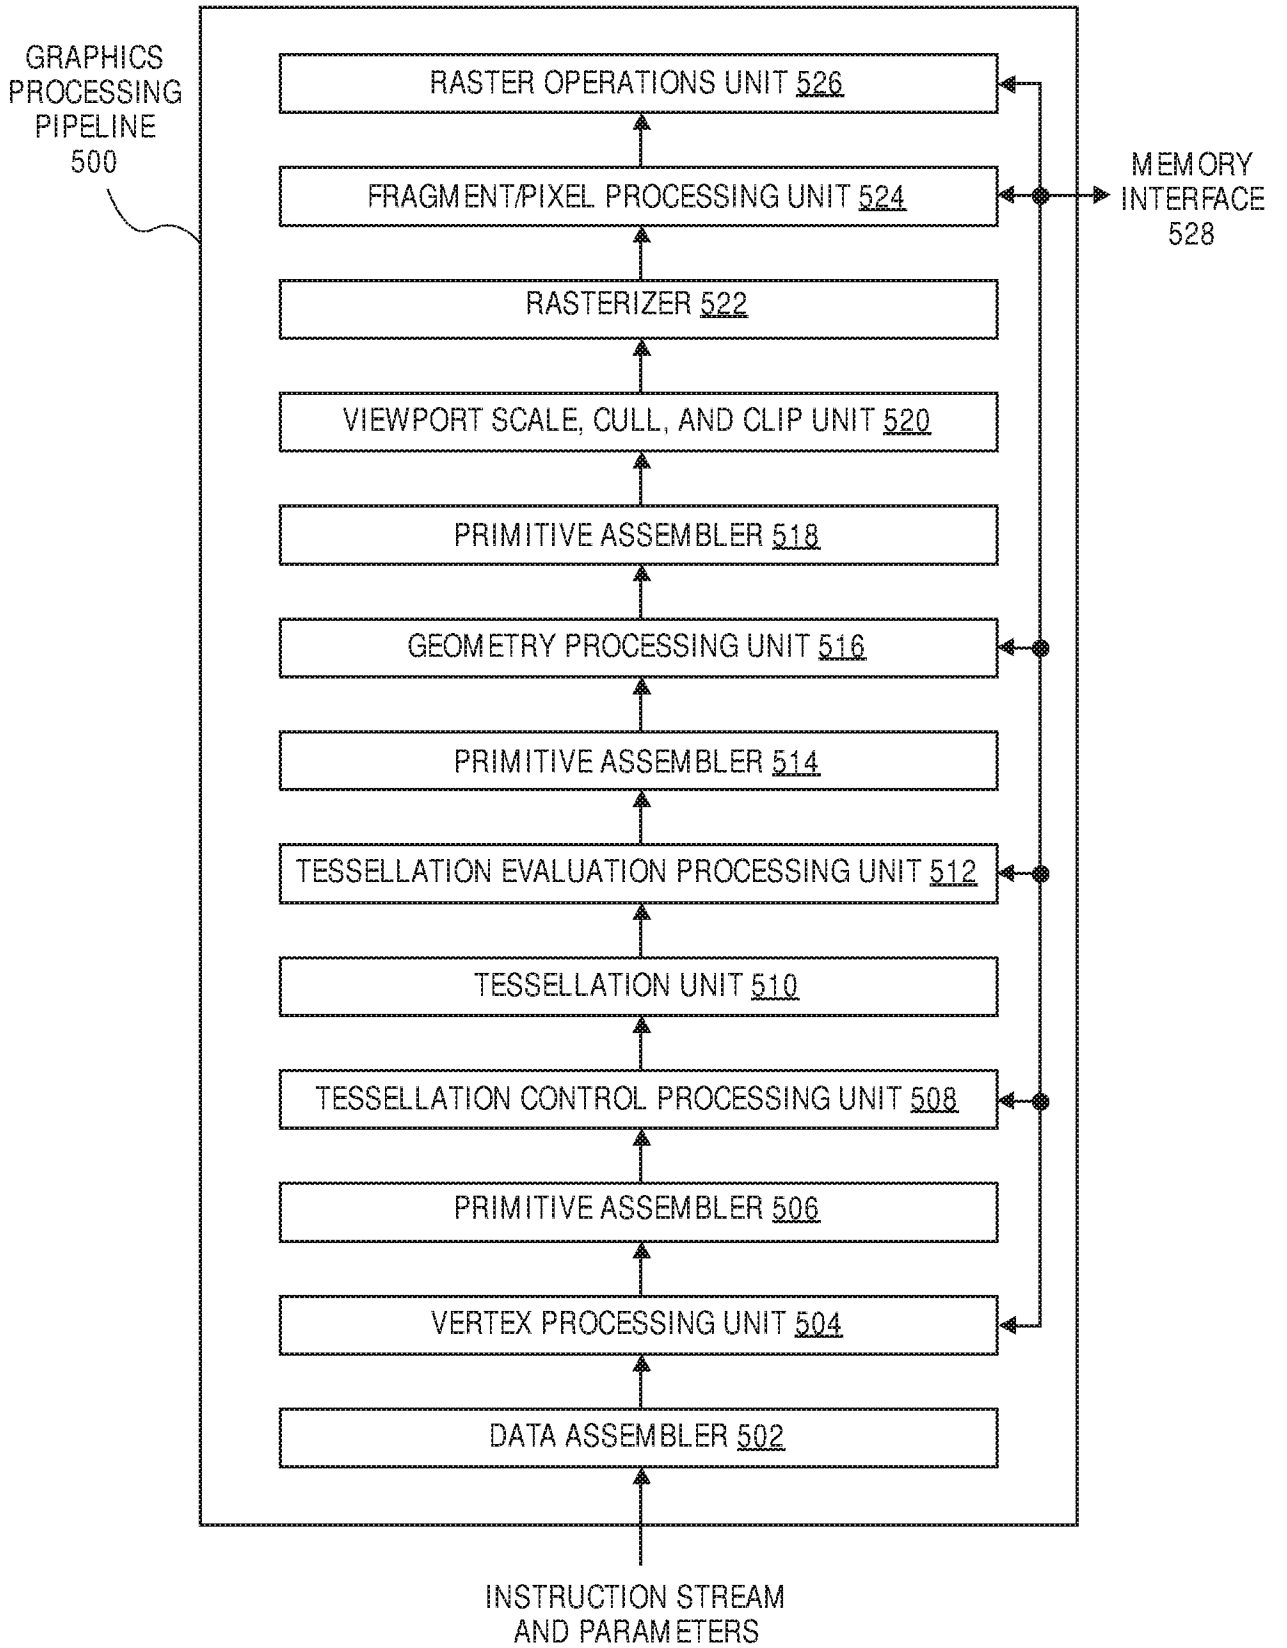


FIG. 5

600

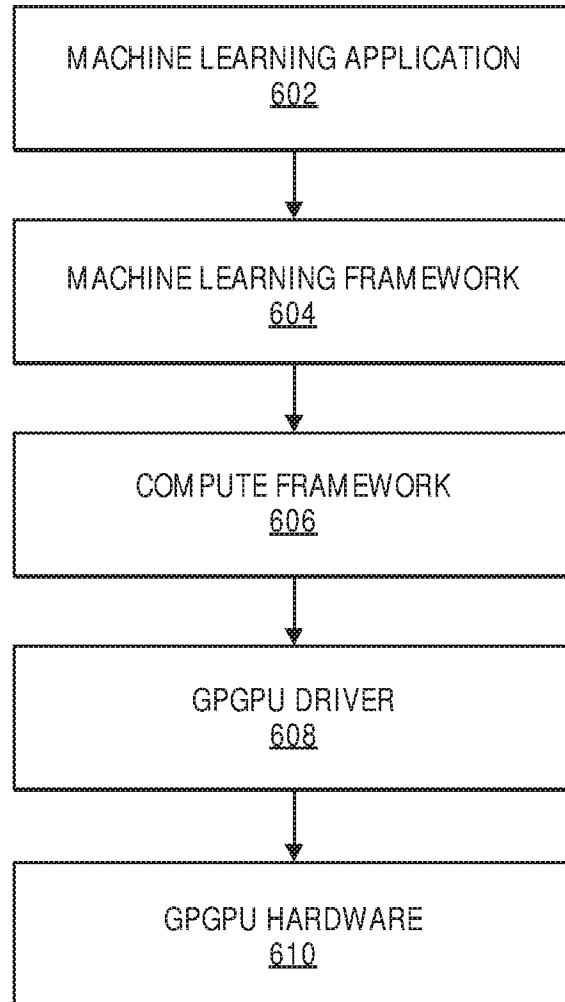


FIG. 6

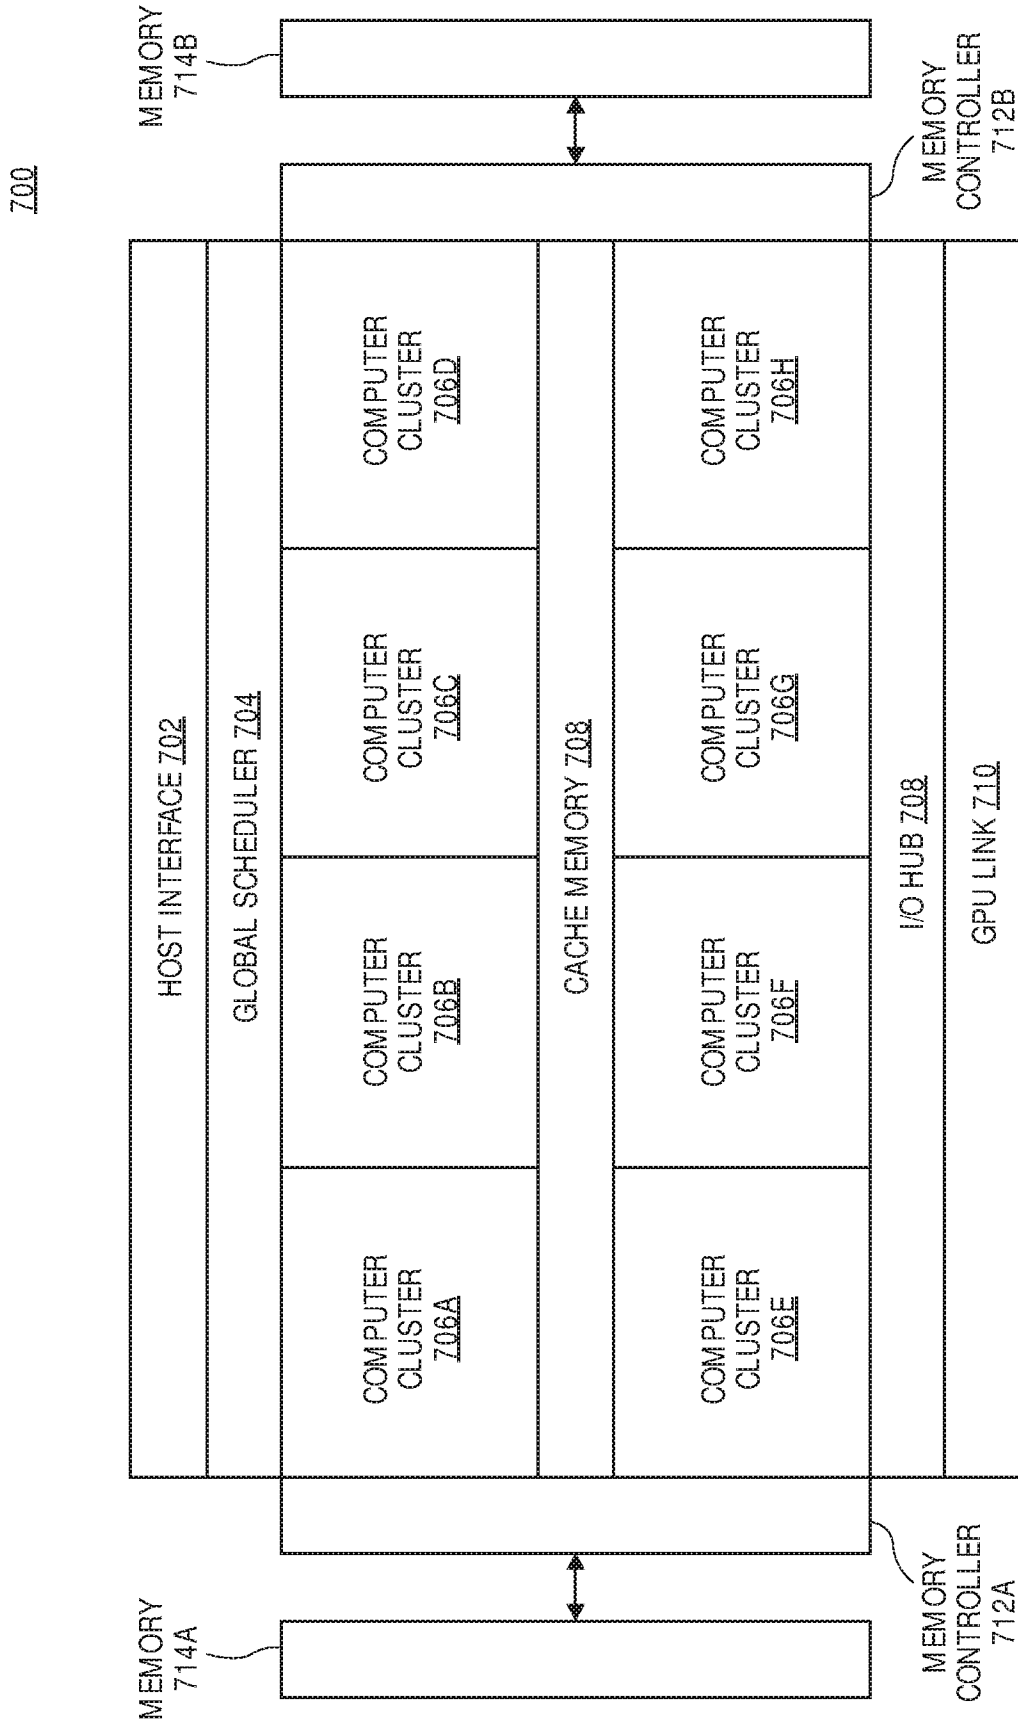


FIG. 7

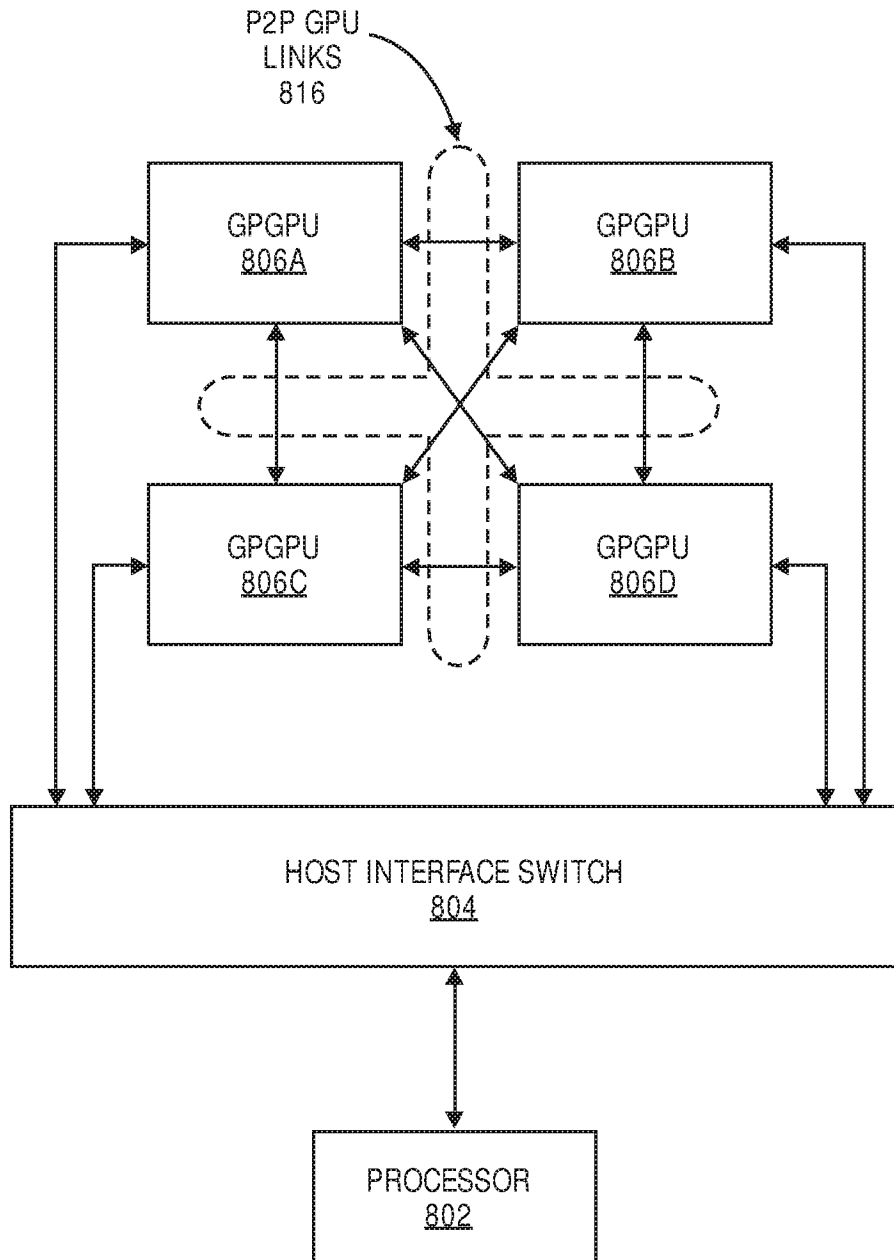


FIG. 8




FIG. 9A

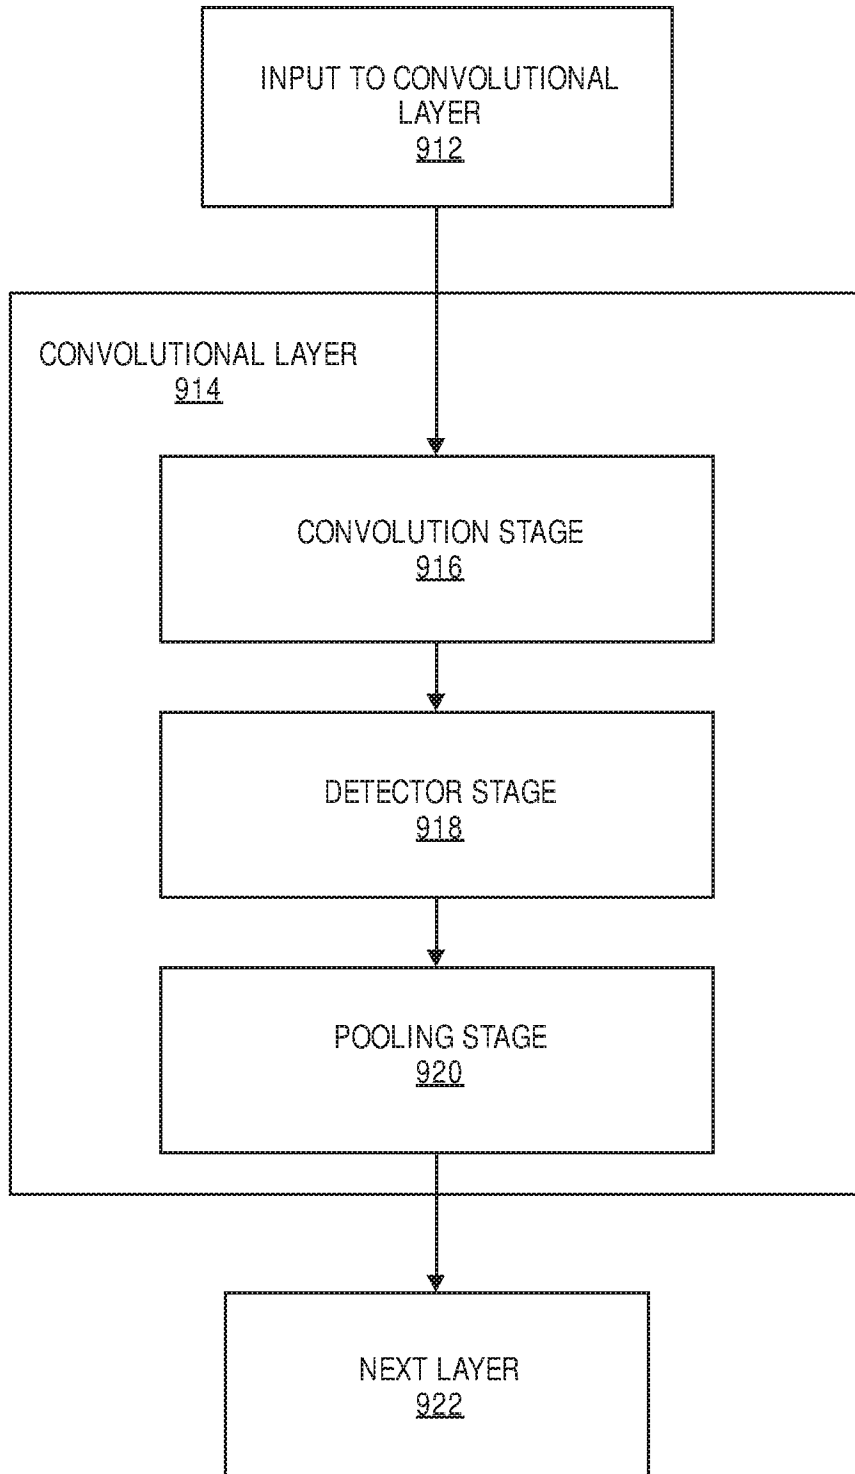


FIG. 9B




FIG. 10

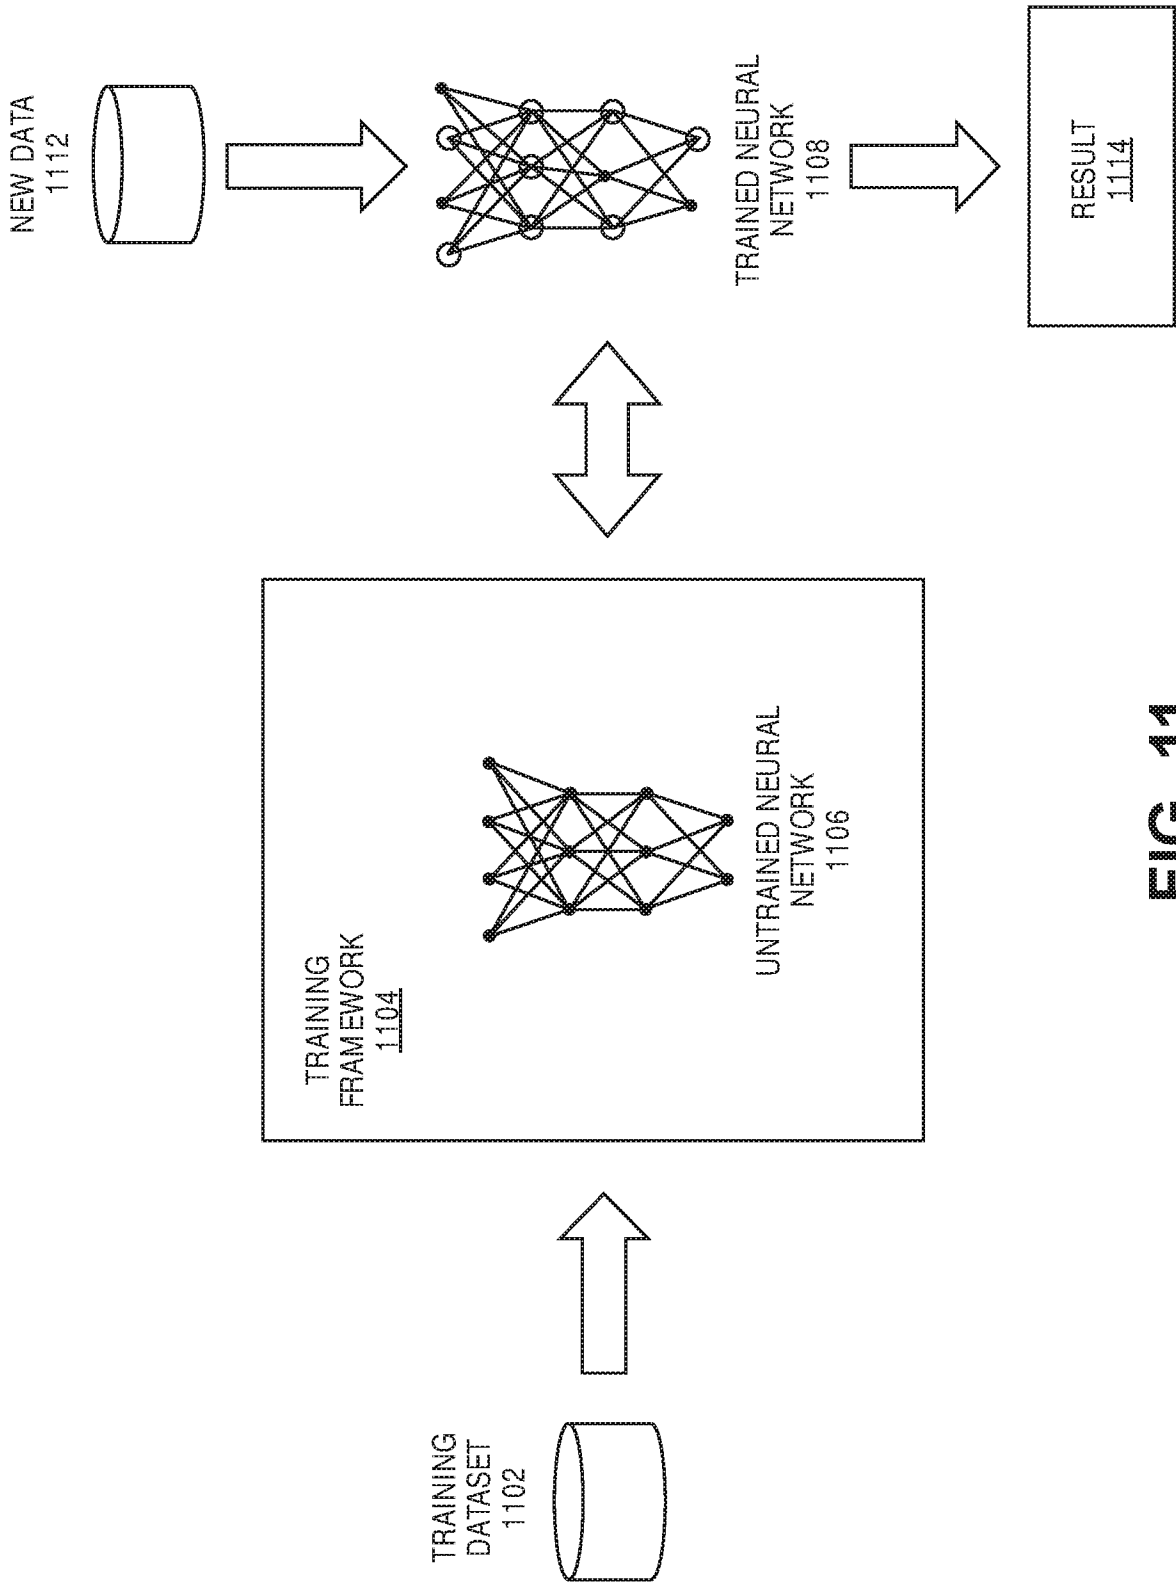


FIG. 11

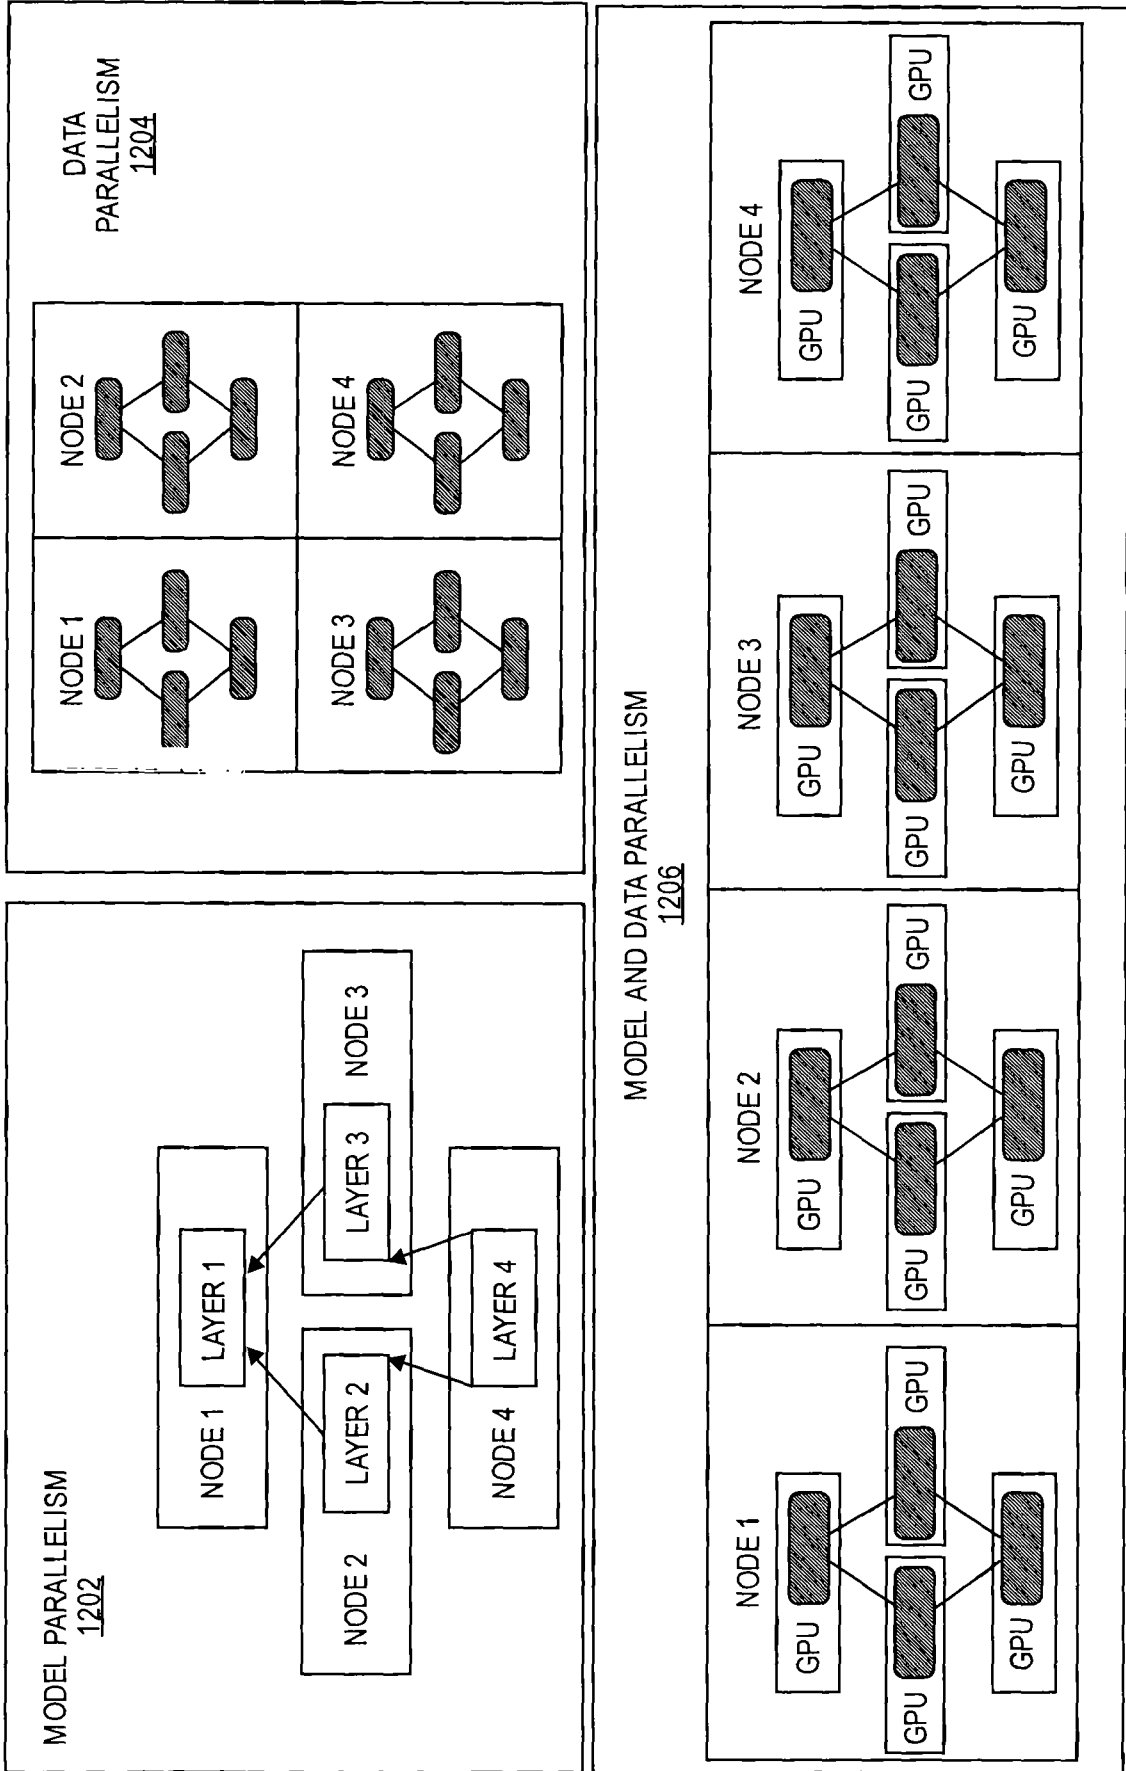


FIG. 12

1300

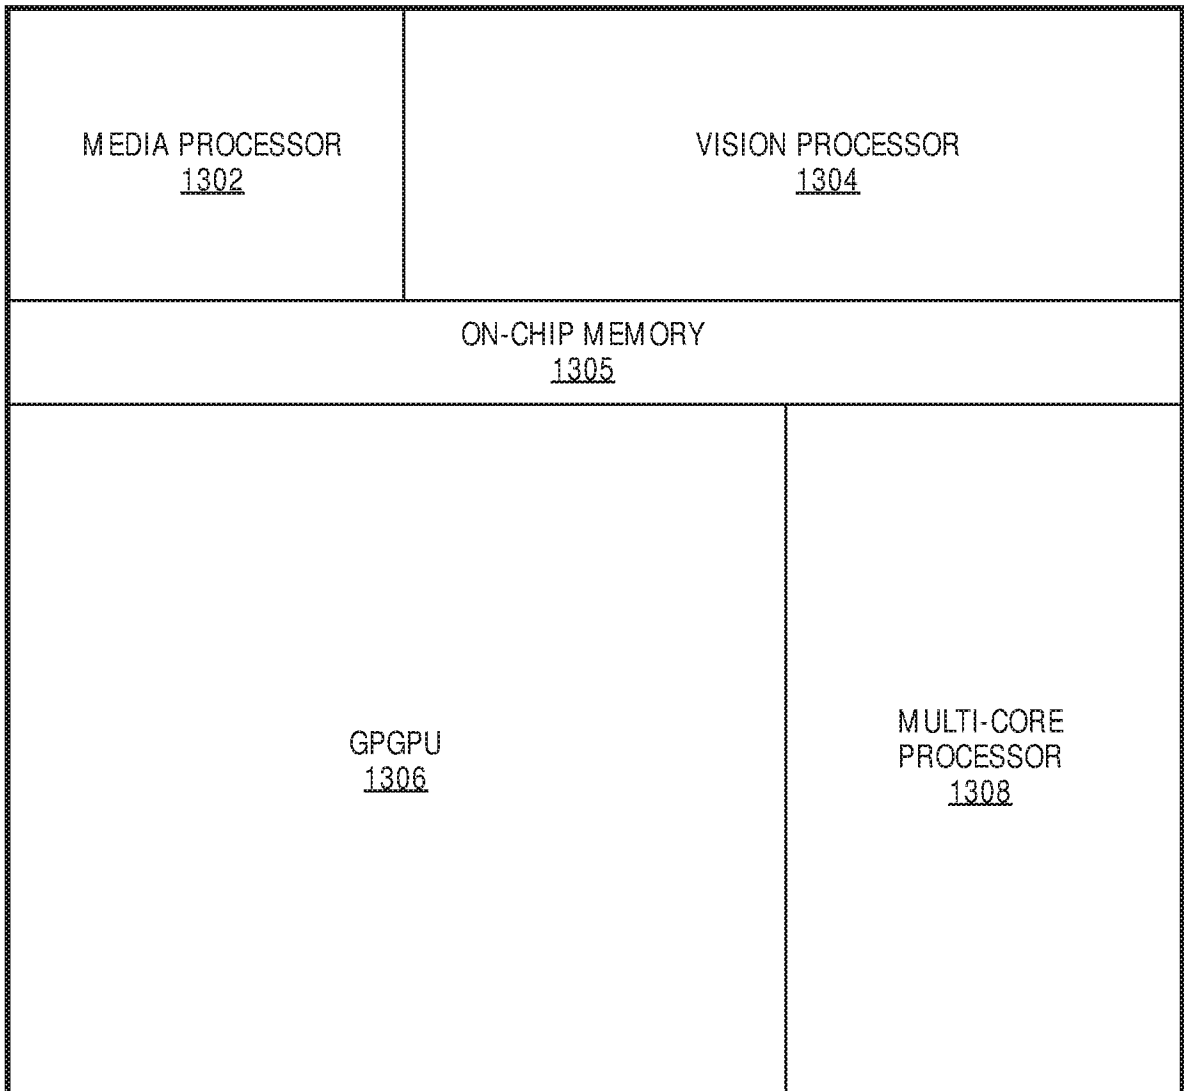


FIG. 13

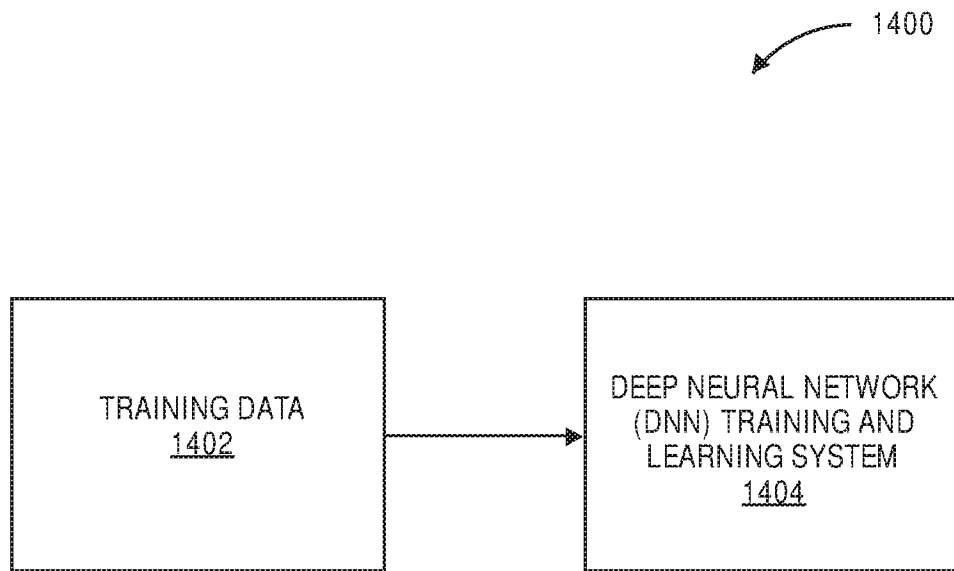


FIG. 14




FIG. 15




FIG. 16A

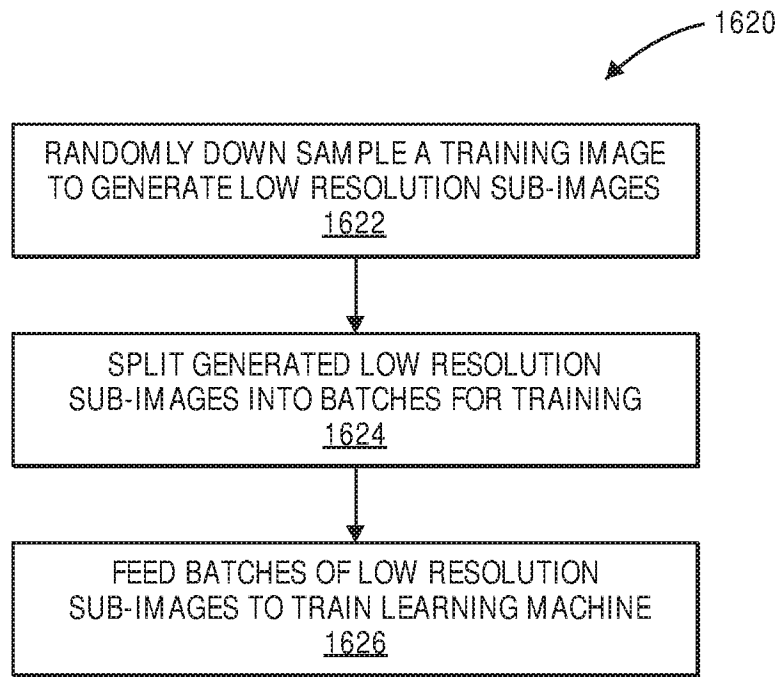


FIG. 16B

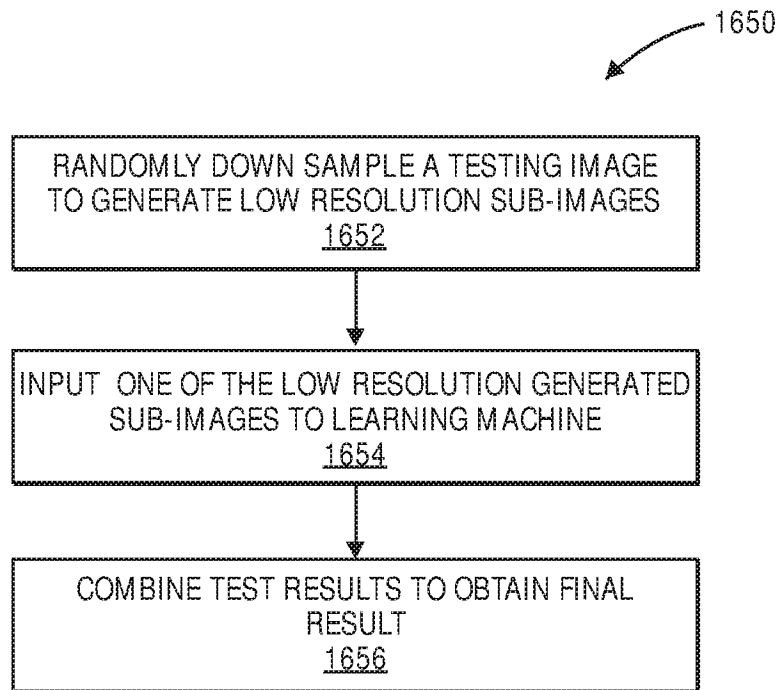


FIG. 16C




FIG. 17A

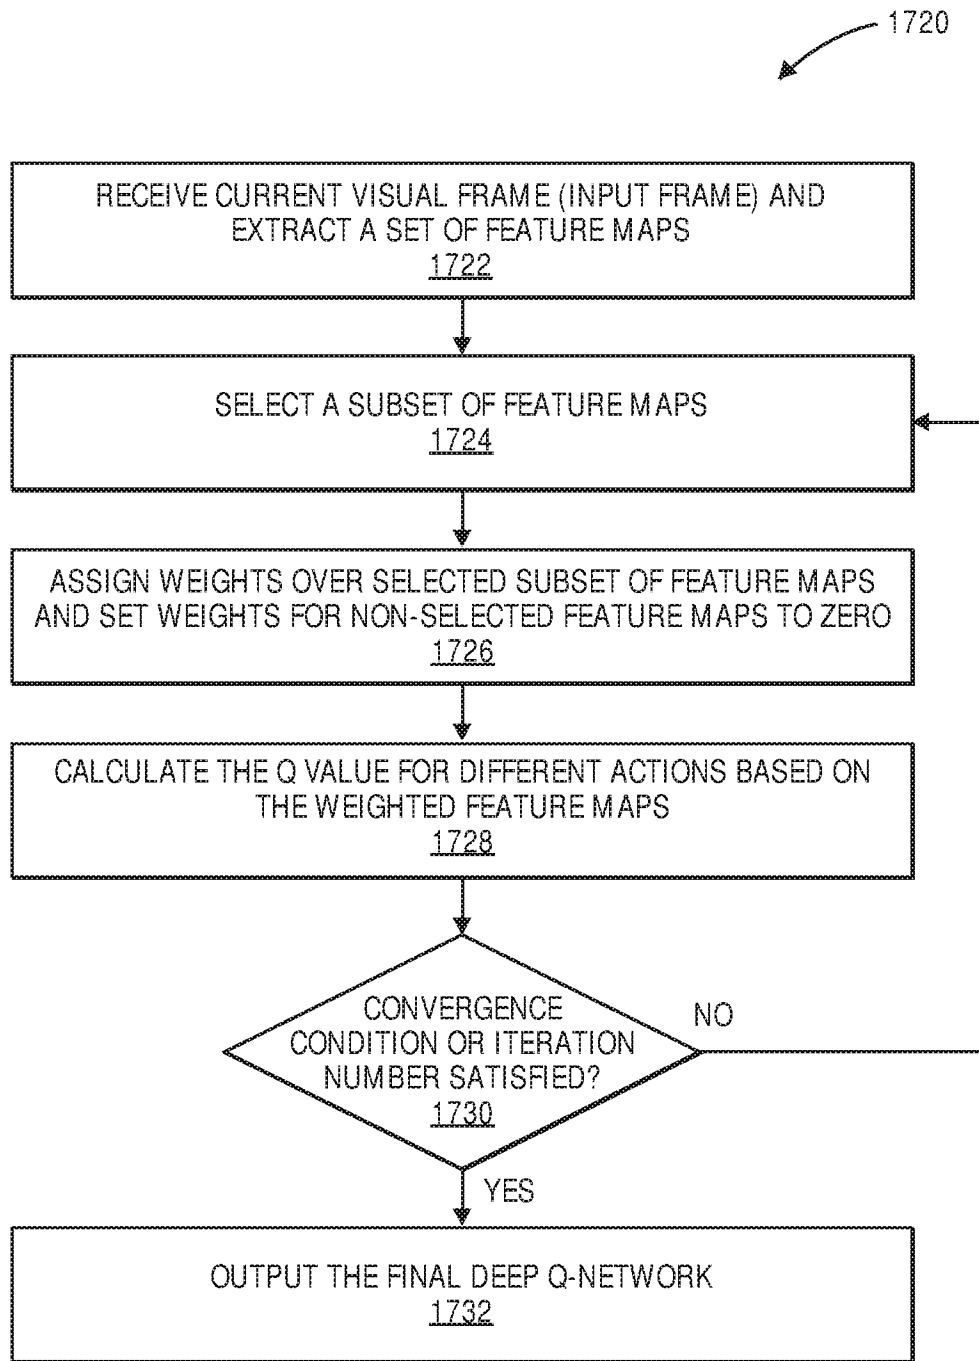


FIG. 17B




FIG. 18A

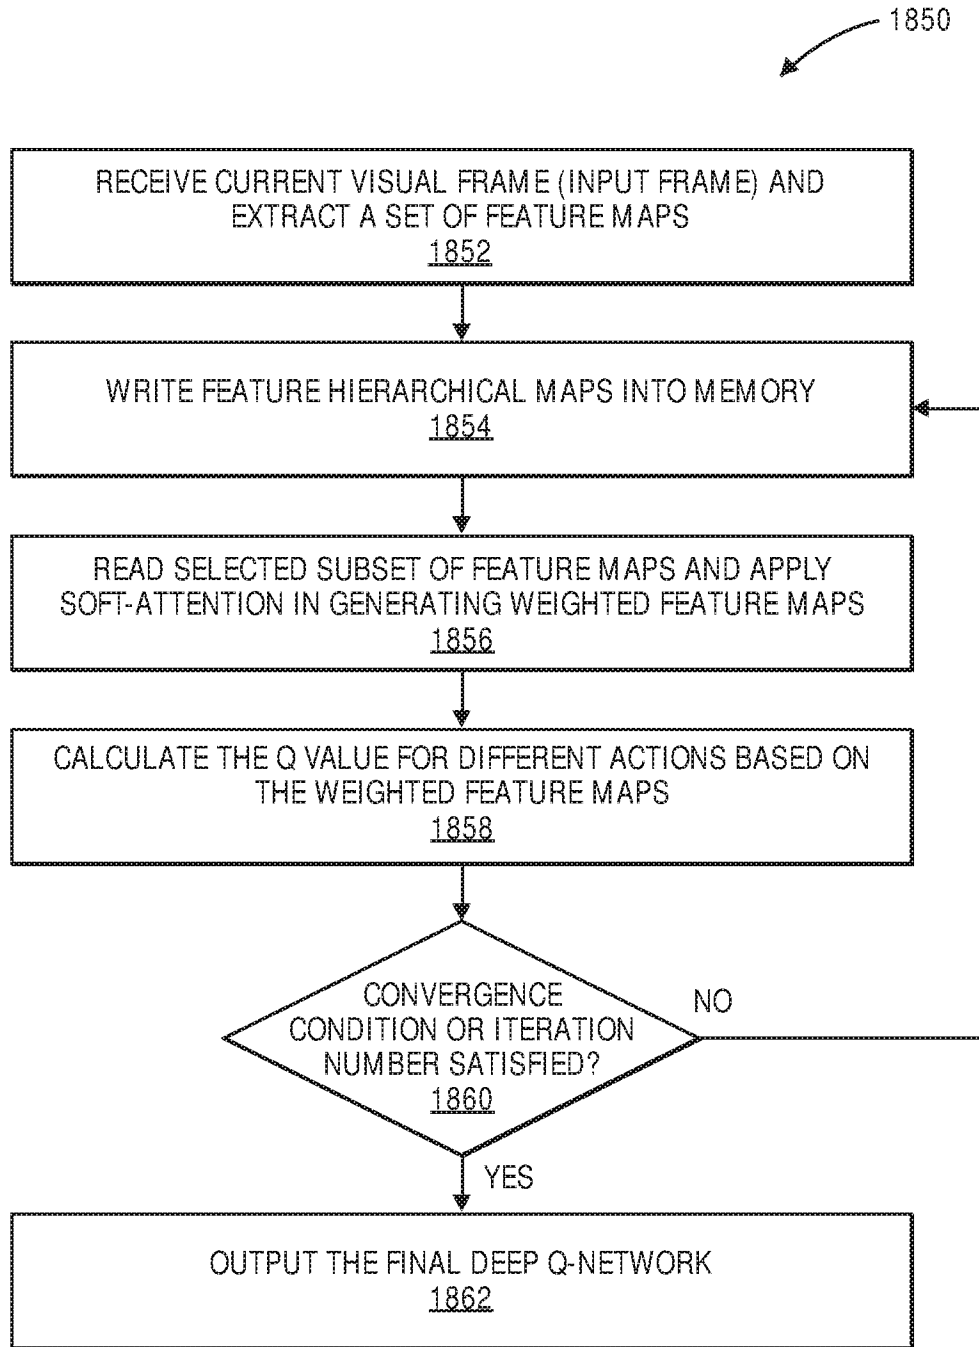


FIG. 18B

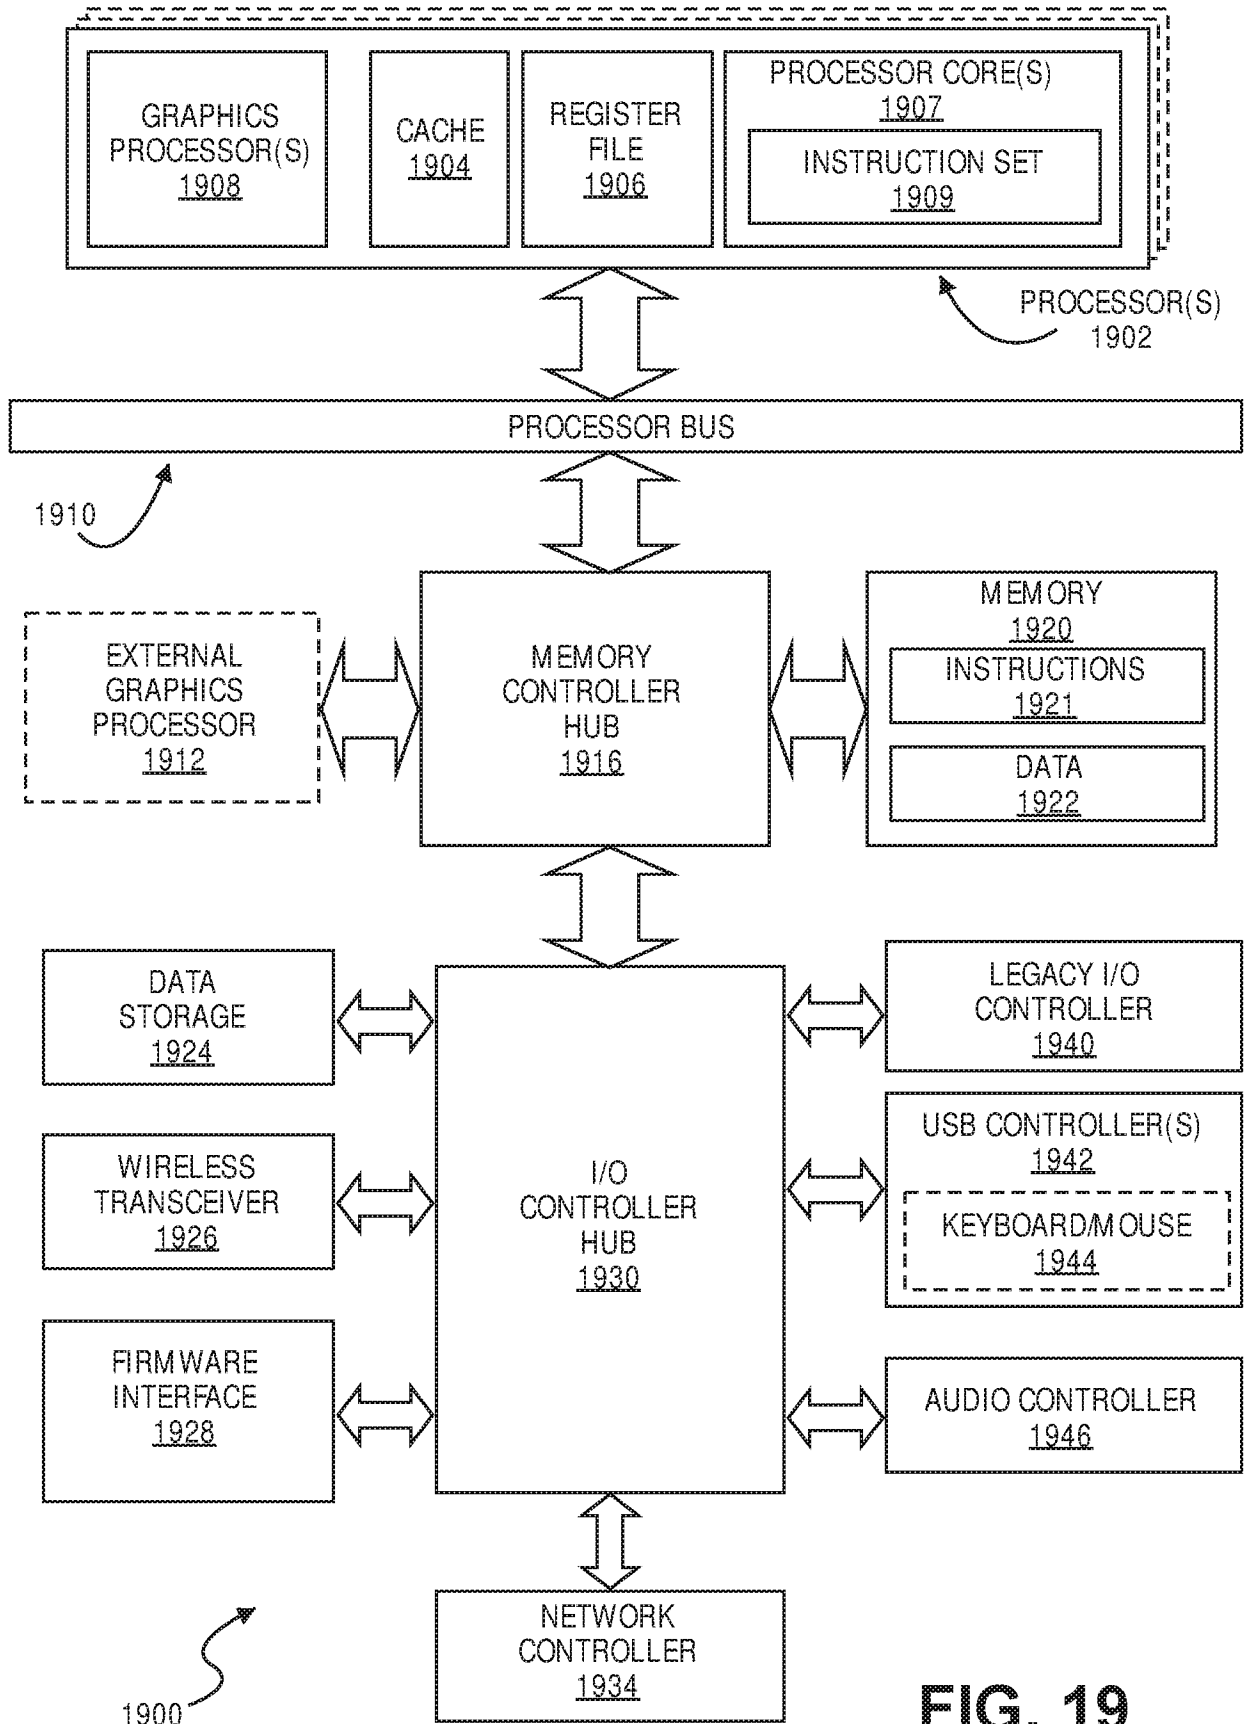


FIG. 19

PROCESSOR
2000




FIG. 20




FIG. 21




FIG. 22




FIG. 23






FIG. 24

GRAPHICS CORE INSTRUCTION FORMATS

2500



64-BIT COMPACT INSTRUCTION 2530



OPCODE DECODE 2540




FIG. 25




FIG. 26

FIG. 27A

SAMPLE COMMAND FORMAT
2700


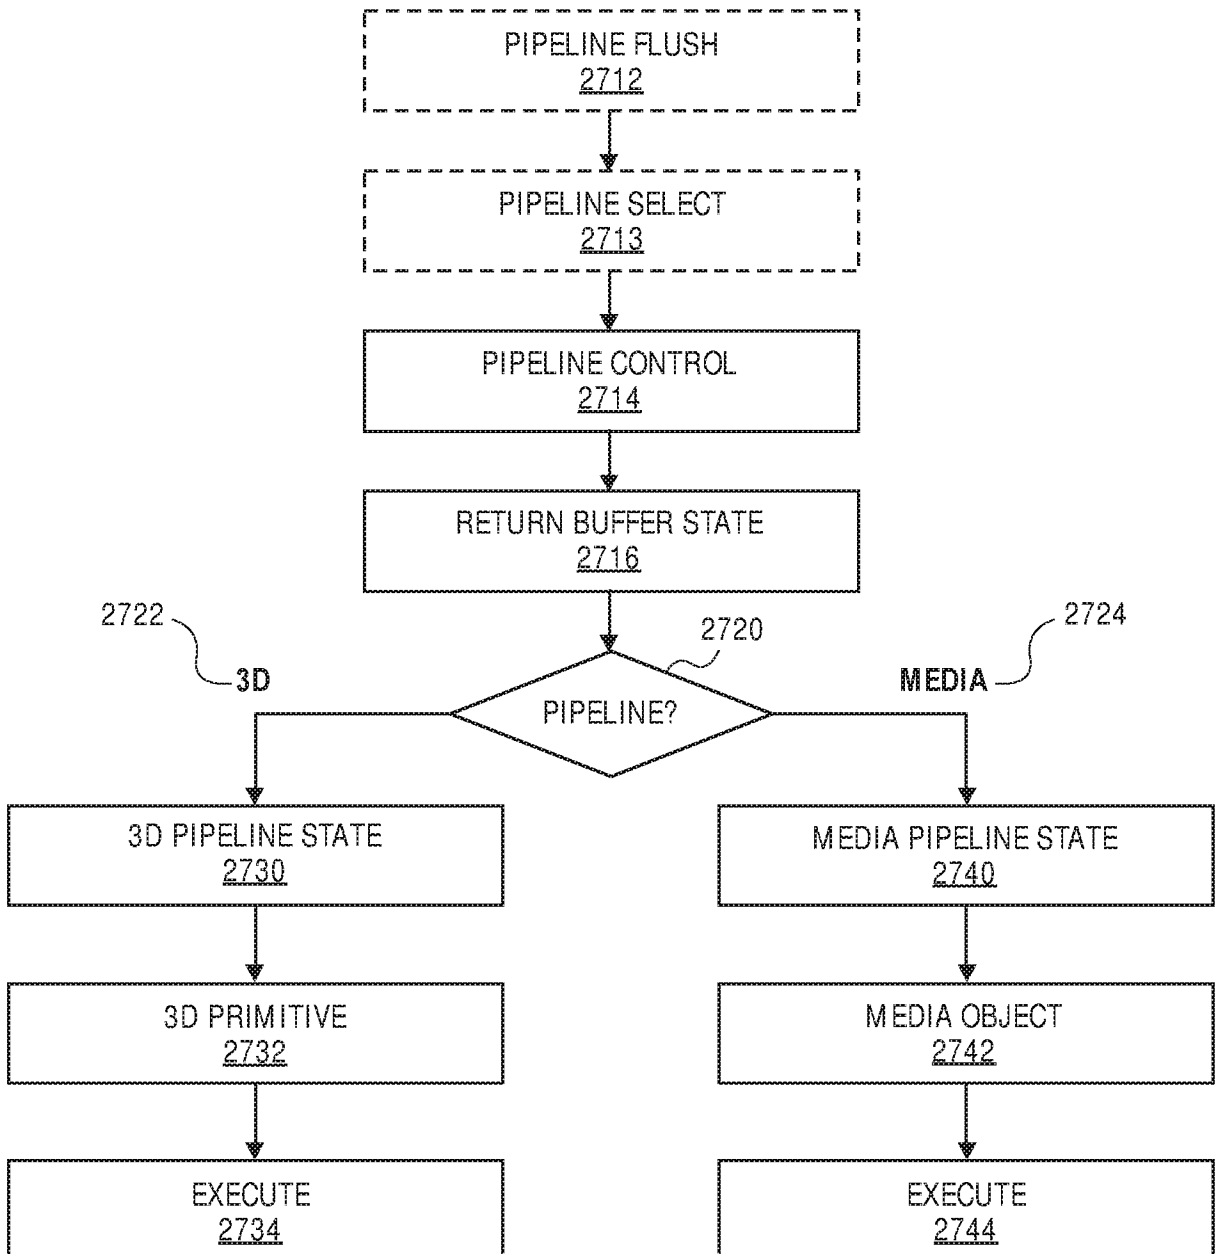


FIG. 27B

SAMPLE COMMAND SEQUENCE
2710



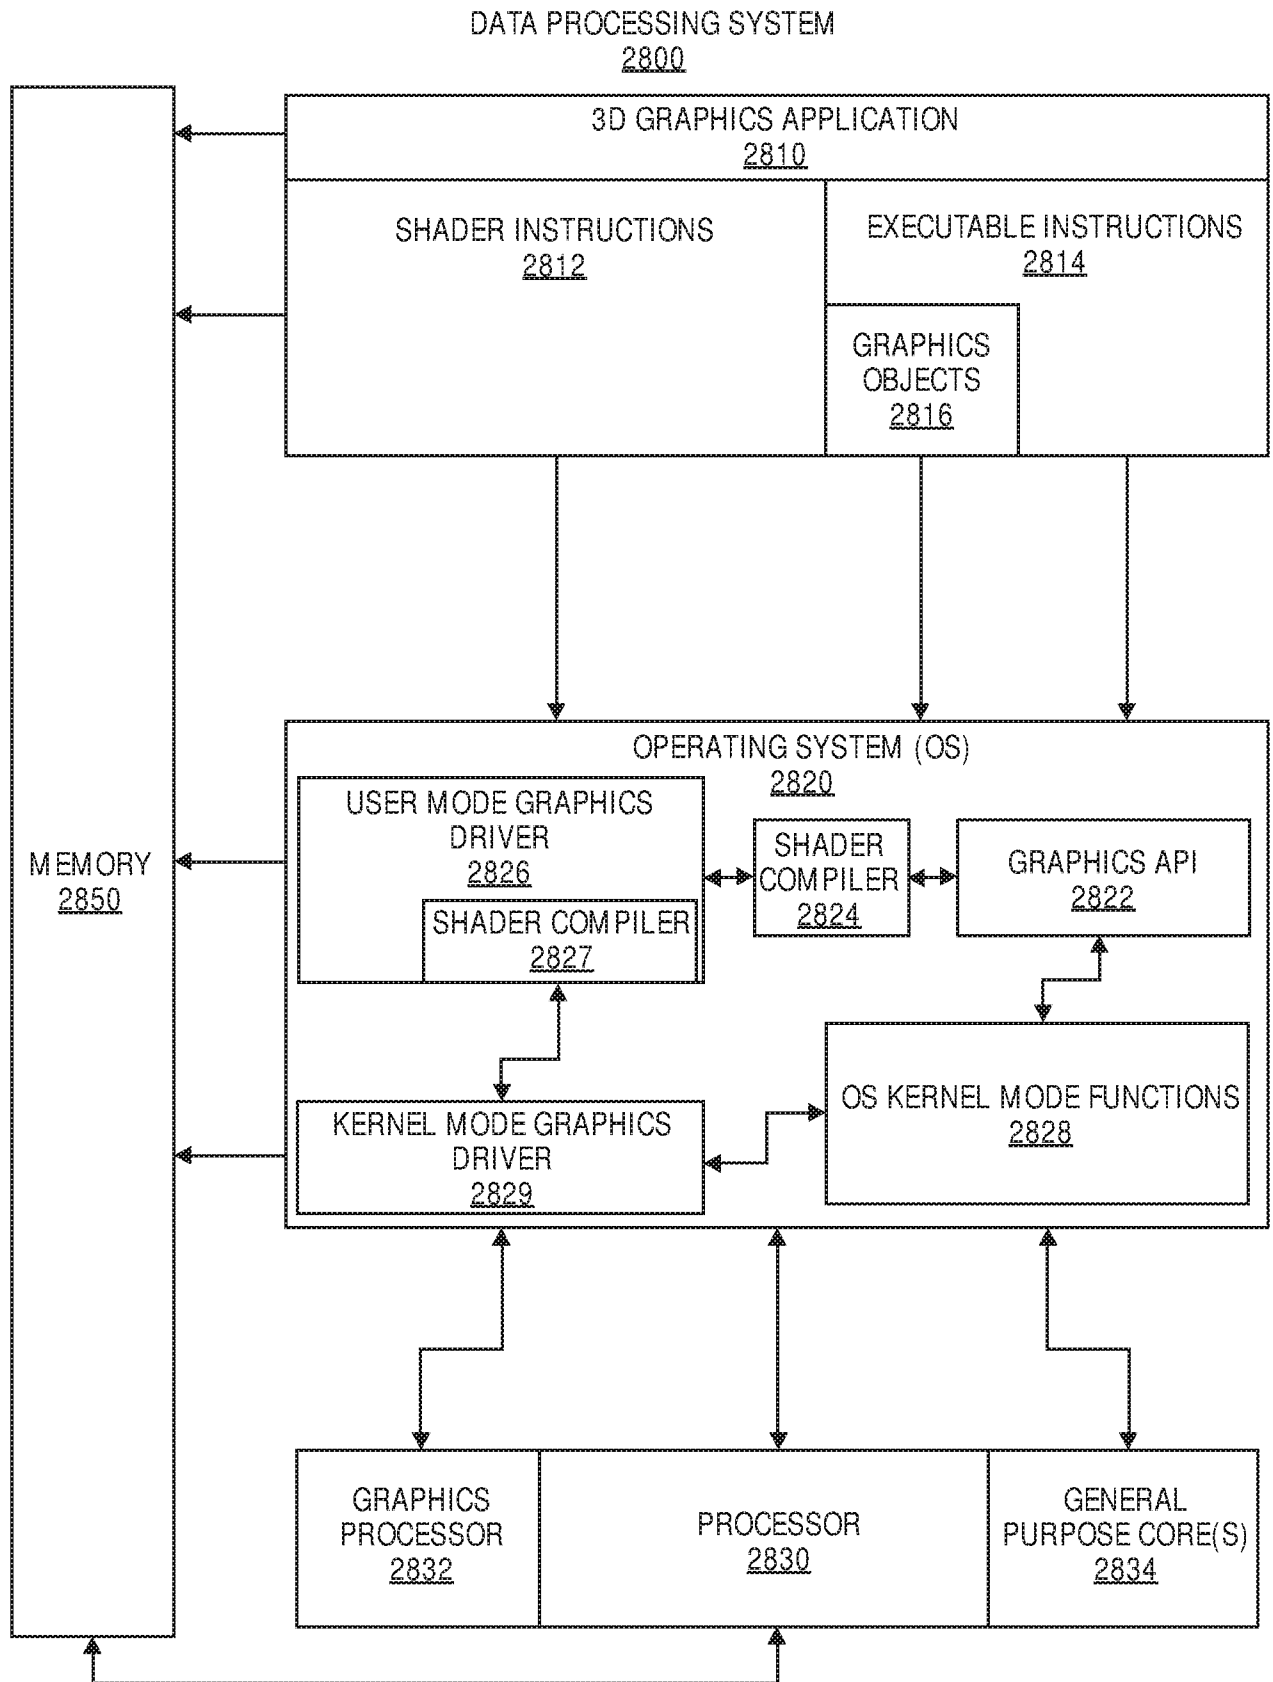


FIG. 28

IP CORE DEVELOPMENT
2900




FIG. 29




FIG. 30

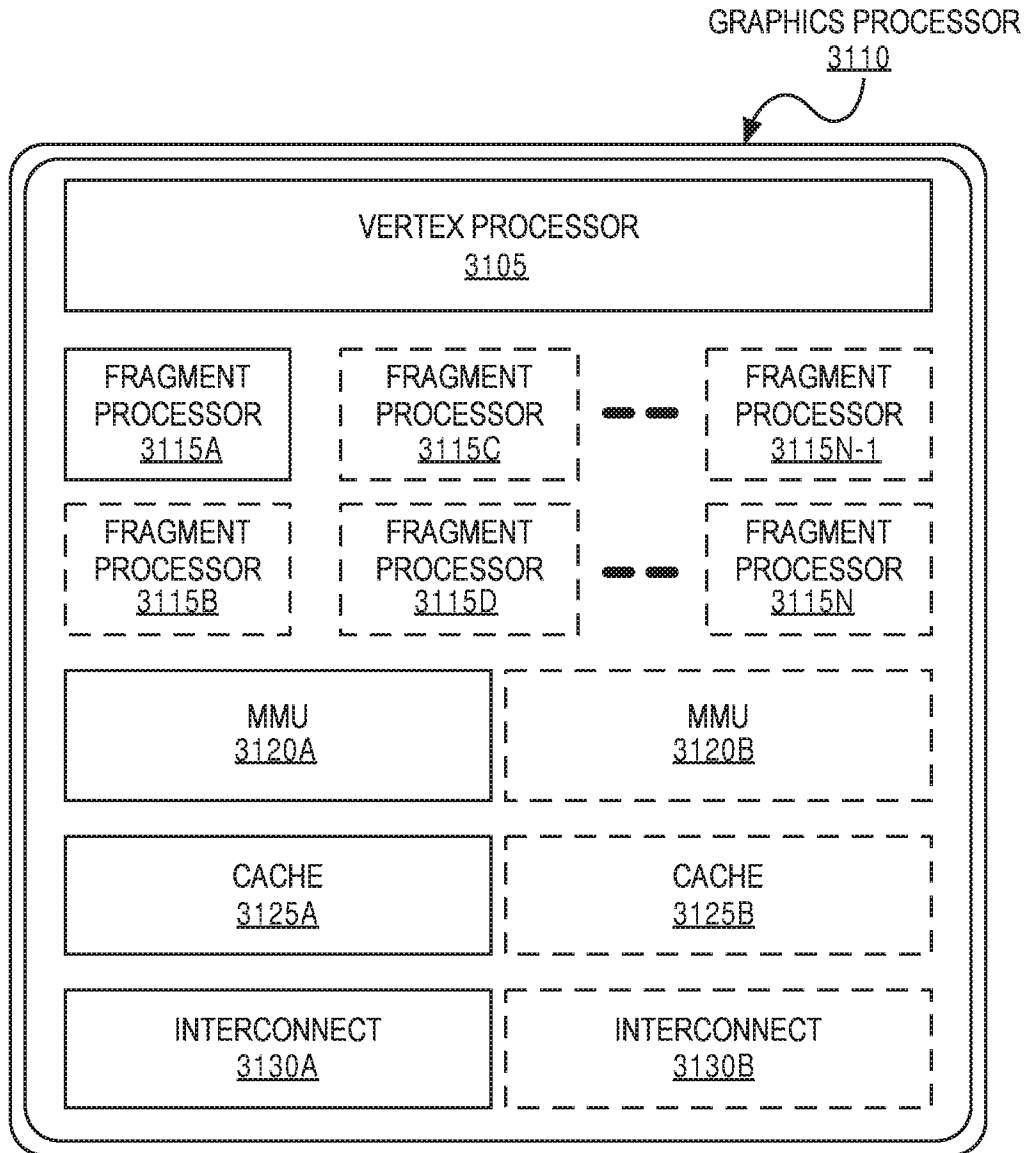


FIG. 31

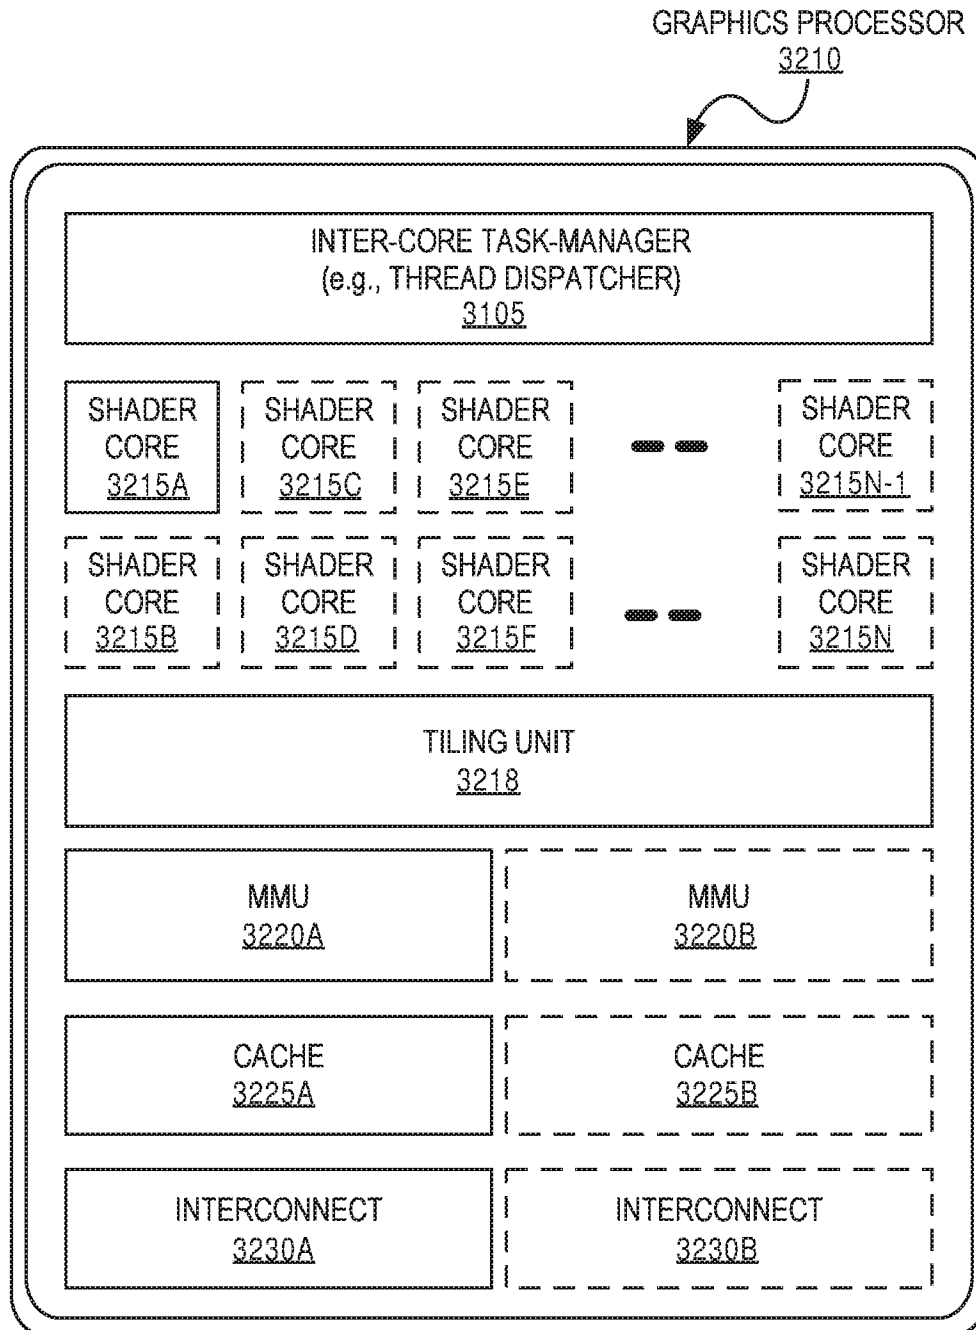


FIG. 32

INTERNATIONAL SEARCH REPORT

International application No.

PCT/CN2017/079719

A. CLASSIFICATION OF SUBJECT MATTER		
G06K 9/66(2006.01)i; G06N 3/08(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)		
G06K; G06N		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)		
CNPAT,EPODOC,WPI,CNKI,IEEE: deep neural network, DNN, random+, train+, test+, image?, graphics, process+, I/O, recurrent, Q-network?, CNN, LSTM, +attention, RDQN, DRQN, DQN, long short time memory, Q value		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	CN 104408483 A (XI DIAN UNIVERSITY) 11 March 2015 (2015-03-11) description, paragraphs [0036]-[0077] and figures 1-4	1-10
X	SOROKIN, Ivan et al. "Deep Attention Recurrent Q-Network" <i>Proceedings of the NIPS Workshop on Deep Reinforcement Learning</i> , 05 December 2015 (2015-12-05), pages 1-7	11-18
A	CN 106446930 A (UNIVERSITY SHENYANG TECHNOLOGY) 22 February 2017 (2017-02-22) the whole document	1-18
A	WO 2017044842 A1 (GOOGLE INC.) 16 March 2017 (2017-03-16) the whole document	1-18
A	US 2013266214 A1 (BRIGHAM YOUNG UNIVERSITY) 10 October 2013 (2013-10-10) the whole document	1-18
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input checked="" type="checkbox"/> 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		Date of mailing of the international search report
27 December 2017		18 January 2018
Name and mailing address of the ISA/CN		Authorized officer
STATE INTELLECTUAL PROPERTY OFFICE OF THE P.R.CHINA 6, Xitucheng Rd., Jimen Bridge, Haidian District, Beijing 100088 China		MA,Zeyu
Facsimile No. (86-10)62019451		Telephone No. (86-10)62413958

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.

PCT/CN2017/079719

Patent document cited in search report			Publication date (day/month/year)	Patent family member(s)			Publication date (day/month/year)
CN	104408483	A	11 March 2015	None			
CN	106446930	A	22 February 2017	None			
WO	2017044842	A1	16 March 2017	US	2017076201	A1	16 March 2017
US	2013266214	A1	10 October 2013	None			