



- (51) **International Patent Classification:**
G06F 9/38 (2006.01) G06T 1/20 (2006.01)
G06F 9/48 (2006.01) G06F 12/1018 (2016.01)
- (21) **International Application Number:**
PCT/US2020/035203
- (22) **International Filing Date:**
29 May 2020 (29.05.2020)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
16/426,613 30 May 2019 (30.05.2019) US
- (71) **Applicant: ADVANCED MICRO DEVICES, INC.**
[US/US]; 2485 Augustine Drive, Santa Clara, California 95054 (US).
- (72) **Inventors: MCCRARY, Rex Eldon;** 3501 Quadrangle Blvd., Suite 375, Orlando, Florida 32817 (US). **LUO, Yi;** 3501 Quadrangle Blvd., Suite 375, Orlando, Florida 32817 (US). **WISE, Harry J.;** 3501 Quadrangle Blvd., Suite 375, Orlando, Florida 32817 (US). **ASHKAR, Alexander Fuad;** 3501 Quadrangle Blvd., Suite 375, Orlando, Florida 32817 (US). **MANTOR, Michael;** 3501 Quadrangle Blvd., Suite 375, Orlando, Florida 32817 (US).

(74) **Agent: SHEEHAN, Adam D.;** Davidson Sheehan LLP, 6836 Austin Center Blvd., Suite 320, Austin, Texas 78731 (US).

(81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, 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, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, 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).

(54) **Title:** GRAPHICS CONTEXT BOUNCING

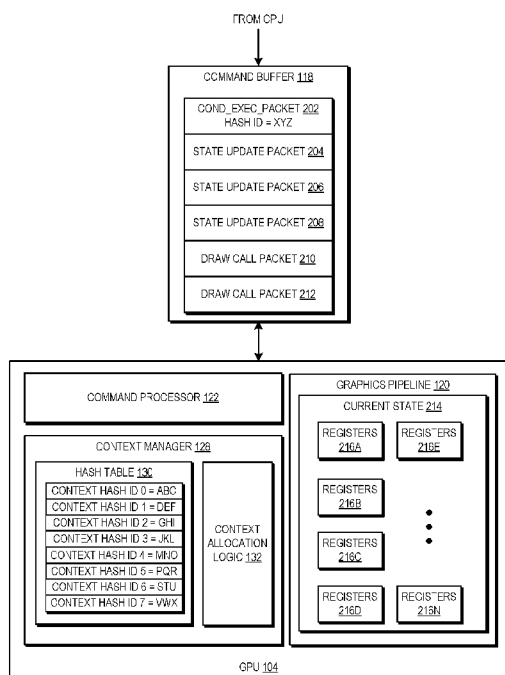


FIG. 2

(57) **Abstract:** A method of context bouncing includes receiving, at a command processor [122] of a graphics processing unit (GPU) [104], a conditional execute packet [202] providing a hash identifier corresponding to an encapsulated state. The encapsulated state includes one or more context state packets [204, 206] following the conditional execute packet. A command packet [210] following the encapsulated state is executed based at least in part on determining whether the hash identifier of the encapsulated state matches one of a plurality of hash identifiers of active context states currently stored at the GPU.

WO 2020/243482 A1

Published:

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

GRAPHICS CONTEXT BOUNCING

BACKGROUND

A graphics processing unit (GPU) is a processing unit that is specially designed to perform graphics processing tasks. A GPU may, for example, execute graphics processing tasks required by an end-user application, such as a video game application. Typically, there are several layers of software between the end-user application and the GPU. For example, in some cases the end-user application communicates with the GPU via an application programming interface (API). The API allows the end-user application to output graphics data and commands in a standardized format, rather than in a format that is dependent on the GPU.

Many GPUs include graphics pipelines for executing instructions of graphics applications. A graphics pipeline includes a plurality of processing blocks that work on different steps of an instruction at the same time. Pipelining enables a GPU to take advantage of parallelism that exists among the steps needed to execute the instruction. As a result, a GPU can execute more instructions in a shorter period of time. The output of the graphics pipeline is dependent on the state of the graphics pipeline. The state of a graphics pipeline is updated based on state packages (e.g., context-specific constants including texture handlers, shader constants, transform matrices, and the like) that are locally stored by the graphics pipeline. Because the context-specific constants are locally maintained, they can be quickly accessed by the graphics pipeline.

BRIEF DESCRIPTION OF THE DRAWINGS

The present disclosure may be better understood, and its numerous features and advantages made apparent to those skilled in the art by referencing the accompanying drawings. The use of the same reference symbols in different drawings indicates similar or identical items.

FIG. 1 is a block diagram illustrating a system including a graphics processing unit configured to employ graphics context bouncing in accordance with some embodiments.

FIG. 2 is a block diagram illustrating processing of command stream instructions at the GPU of FIG. 1 in accordance with some embodiments.

FIG. 3 is a block diagram illustrating draining of encapsulated state at the GPU of FIG. 1 in accordance with some embodiments.

5 FIG. 4 is a flow diagram illustrating a method of graphics context bouncing in accordance with some embodiments.

DETAILED DESCRIPTION

To perform graphics processing, a central processing unit (CPU) of a system often issues to a GPU a call, such as a draw call, which includes a series of
10 commands instructing the GPU to draw an object according to the CPU's instructions. As the draw call is processed through the GPU graphics pipeline, the draw call uses various configurable settings to decide how meshes and textures are rendered. A common GPU workflow involves updating the values of constants in a memory array and then performing a draw operation using the constants as data. A GPU whose
15 memory array contains a given set of constants may be considered to be in a particular state. These constants and settings, referred to as context state (also referred to as "rendering state", "GPU state", or simply "state"), affect various aspects of rendering and include information the GPU needs to render an object. The context state provides a definition of how meshes are rendered and includes information
20 such as the current vertex/index buffers, the current vertex/pixel shader programs, shader inputs, texture, material, lighting, transparency, and the like. The context state contains information unique to the draw or set of draws being rendered at the graphics pipeline. Context therefore refers to the required GPU pipeline state to correctly draw something.

25 Many GPUs use a technique known as pipelining to execute instructions. Pipelining enables a GPU to work on different steps of an instruction at the same time, thereby taking advantage of parallelism that exists among the steps needed to execute the instruction. As a result, the GPU can execute more instructions in a shorter period of time. The video data output by the graphics pipeline are dependent
30 on state packages (e.g., context-specific constants) that are locally stored by the

graphics pipeline. In GPUs, it is common to set up the state of the GPU, perform a draw operation, and then make only a small number of changes to the state before the next draw operation. The state settings (e.g. values of constants in memory) often remain the same from one draw operation to the next.

5 The context-specific constants are locally maintained for quick access by the graphics pipeline. However, GPU hardware is generally memory constrained and only locally stores (and therefore operates on) a limited number of sets of context state. Accordingly, the GPU will often change context state in order to start working on a new set of context registers because the graphics pipeline state needs to be
10 changed to draw something else. The GPU performs a context roll to a newly supplied context to release the current context by copying the current registers into a newly allocated context before applying any new state updates by programming fresh register values. Due to the limited number of context state sets stored locally, the GPU sometimes runs out of context and the graphics pipeline is stalled while waiting
15 for a context to be freed so that a new context may be allocated. In various embodiments, the CPU reuses a context state that was previously active, and without context bouncing would release the current context and allocate a new one. If no context is available, the GPU would have to wait for one to be freed, thus introducing a stall.

20 To improve GPU system performance, FIGs. 1-4 illustrate systems and methods for context state management by assigning hash-based identifiers to context states and switching to matching context (referred to as context bouncing) without allocating a new context. In various embodiments, a method of context bouncing includes receiving, at a command processor of a graphics processing unit (GPU), a conditional
25 execute packet providing a hash identifier corresponding to an encapsulated state. The encapsulated state includes one or more context state packets following the conditional execute packet. A command packet following the encapsulated state is executed based at least in part on determining whether the hash identifier of the encapsulated state matches one of a plurality of hash identifiers of active context
30 states currently stored at the GPU. In various embodiments, determining whether the hash identifier of the encapsulated state matches one of a plurality of hash identifiers

includes querying a context manager to search a hash table storing the plurality of hash identifiers of active context states currently stored at the GPU.

In some embodiments, context bouncing also includes draining the encapsulated state in response to determining the hash identifier matches one of the plurality of hash identifiers of active context states. In other embodiments, a new context state set is allocated in response to determining the hash identifier does not match any of the plurality of hash identifiers of active context states. Additionally, one or more context state packets of the encapsulated state are executed prior to executing the command packet. Bouncing back to an existing context (i.e., without allocating a new context) reduces the likelihood of running out of context and stalling the graphics pipeline.

FIG. 1 is a block diagram illustrating a system 100 including a graphics processing unit configured to employ graphics context bouncing in accordance with some embodiments. In the depicted example, the system 100 includes a central processing unit (CPU) 102, a graphics processing unit (GPU) 104, a device memory 106 utilized by the GPU 104, and a system memory 108 shared by the CPU 102 and the GPU 104. In various embodiments, the GPU 104 may be a dedicated device, several devices, or integrated into a larger device (e.g., a north bridge device). The memories 106, 108 include any of a variety of random access memories (RAMs) or combinations thereof, such as a double-data-rate dynamic random access memory (DDR DRAM), a graphics DDR DRAM (GDDR DRAM), and the like. The GPU 104 communicates with the CPU 102, the device memory 106, and the system memory 108 via a bus 110. The bus 110 includes any type of bus used in computing systems, including, but not limited to, a peripheral component interface (PCI) bus, an accelerated graphics port (AGP) bus, a PCI Express (PCIE) bus, and the like.

The CPU 102 sends instructions intended for processing at the GPU 104 to command buffers. In various embodiments, the command buffer may be located, for example, in system memory 108 in a separate memory coupled to the bus 110 (e.g., device memory 106).

As illustrated, the CPU 102 includes a number of processes, such as executing one or more application(s) 112 to generate graphic commands and a user mode

driver 116 (or other drivers, such as a kernel mode driver). In various embodiments, the one or more applications 112 include applications that utilizes the functionality of GPU 104. An application 112 may include one or more graphics instructions that instruct GPU 104 to render a graphical user interface (GUI) and/or a graphics scene.

5 For example, the graphics instructions may include instructions that define a set of one or more graphics primitives to be rendered by GPU 104.

In some embodiments, the application 112 utilizes a graphics application programming interface (API) 114 to invoke a user mode driver 116 (or a similar GPU driver. User mode driver 116 issues one or more commands to GPU 104 for rendering one or more graphics primitives into displayable graphics images. Based on the graphics instructions issued by application 112 to the user mode driver 116, the user mode driver 116 formulates one or more graphics commands that specify one or more operations for GPU 104 to perform for rendering graphics. In some embodiments, the user mode driver 116 is a part of the application 112 running on the CPU 102. For example, the user mode driver 116 may be part of a gaming application running on the CPU 102. Similarly, a kernel mode driver (not shown) may be part of an operating system running on the CPU 102. The graphics commands generated by the user mode driver 116 include graphics commands intended to generate an image or a frame for display. The user mode driver 116 translates standard code received from the API 114 into a native format of instructions understood by the GPU 104. The user-mode driver 116 is typically written by the manufacturer of the GPU 104. Graphics commands generated by the user mode driver 116 are sent to GPU 104 for execution. The GPU 104 can then execute the graphics commands and use the results to control what is displayed on a display screen.

In various embodiments, the CPU 102 sends graphics commands intended for the GPU 104 to a command buffer 118. Although depicted in FIG. 1 as a separate module for ease of illustration, the command buffer 118 may be located, for example, in device memory 106, system memory 108, or a separate memory communicably coupled to the bus 110. The command buffer 118 temporarily stores a stream of graphics commands that include input to the GPU 104. The stream of graphics commands include, for example, one or more command packets and/or one or more

state update packets. In some embodiments, a command packet includes a draw command (also interchangeably referred to as a "draw call") instructing the GPU 104 to execute processes on image data to be output for display. For example, a draw command may instruct the GPU 104 to render defined by a group of one or more
5 vertices (e.g., defined in a vertex buffer) stored in memory. The geometry defined by the group of one or more vertices correspond, in some embodiments, to a plurality of primitives to be rendered.

In various embodiments, a state command (also interchangeably referred to as a "context state packet" or "state update packet") instructs the GPU 104 to change
10 one or more state variables (e.g., a draw color). In some embodiments, a context state packet includes a constant or a set of constants that update the state of graphics pipeline 120 at the GPU 104. A state update packet may, for example, update colors that are to be drawn or blended during execution of a draw call. Although described here in the context of draw calls and context state packets, those
15 skilled in the art will recognize that other commands may be utilized without departing from the scope of this disclosure.

The GPU 104 includes a command processor 122 that receives commands in a command stream to be executed from the CPU 102 (e.g., via the command buffer
20 118 and bus 110) and coordinates execution of those commands at the graphics pipeline 120. In various embodiments, the command processor 122 is implemented in hardware, firmware, software, or a combination thereof. The command stream includes one or more draw calls and context state packets, as described above. The command processor 122 also manages the context states written to registers of the
25 graphics pipeline 120. In various embodiments, in response to receiving a context state packet, the command processor sets one or more state registers in the GPU 104 to particular values based on the context state packet and/or configures one or more of fixed-function processing units based on the context state packet.

Although illustrated in FIG. 1 as having one command processor 122, those
30 skilled in the art will recognize that the GPU 104 may include any number and type of command processors for retrieving and executing packets from hardware queues. In various embodiments, a "packet" refers to a memory buffer encoding a single command. Different types of packets may be stored in hardware queues, memory

buffers, and the like. Additionally, as used herein, the term “block” refers to a processing module included in an ASIC, an execution pipeline of a CPU, and/or a graphics pipeline of a GPU. Such a processing module may include, but is not limited to, an arithmetic logic unit, a multiply/divide unit, a floating point unit, a color
5 buffer, a vertex shader, a pixel shader, a clipping unit, or some other processing module as would be apparent to a person skilled in the art.

The graphics pipeline 120 includes a number of stages 124, including stage A 124a, stage B 124b, and through stage N 124n. In various embodiments, the various stage 124 each represent a stage of the graphics pipeline 120 that execute various
10 aspects of a draw call. The command processor 122 writes context state updates to local banks of context registers for storing and updating operating state. As illustrated in FIG. 1, the command processor writes state updates to current state registers 126.

Conventional context management techniques dispatch information or
15 command packets to buffers (sometimes referred to as queue buffers). Such queue buffers are created in a first-in-first-out (FIFO)-like sequence to which running programs submits commands. However, due to memory capacity at the GPUs, only a limited number of unique rendering state sets may be stored and therefore operated on in the graphics pipeline at any given moment. Accordingly, sequential
20 processes may be run on one state set but if the driver changes context state after performing a draw operation, the CP must execute a sequence of register writes followed by a read to ensure completion once the GPU changes to a different state set. This context rolling process for a new context includes the command processor 122 sending an event down the graphics pipeline to release the current context.
25 Once completed, the command processor is notified and allows the context to be reused (or a new context to be allocated). However, these context rolling processes increases the likelihood of running out of contexts to use and stalling the graphics pipeline until the draw call associated with the current context completes processing through the various stages of the graphics pipeline. In general, GPU operations are
30 pipelined with various amounts of work in flight at any given moment; stalling the graphics pipeline such that in-flight tasks must retire before changing state increases likelihood of starving the GPU backend.

To more efficiently manage context states, the GPU 104 includes a graphics context manager 128. The graphics context manager 128 maintains a hash table 130 storing a set of hash-based identifiers corresponding to sets of context state currently stored at registers of the GPU 104. The graphics context manager 128
5 further includes a context allocation logic module 132 for performing the hash identifier search and context allocation operations described herein. As described further herein, the user mode driver 116 provides a unique hash identifier to identify a new context state that the user mode driver 116 programs into a graphics command. In various embodiments, the user mode driver 116 indicates to the command
10 processor 122, via a new state packet (or other token method), to scan for all active states at the GPU 104 and determine whether the unique hash identifier of the new context matches any one of the plurality of hash identifiers of currently active context states (i.e., hash identifiers stored at the hash table 130).

In some embodiments, the command processor 122 queries the graphics
15 context manager 128 to search the hash table 130 for determining whether the hash table 130 stores the requested unique hash identifier corresponding to the new context (i.e., indicating that the context is available for use). If the hash table 130 does not contain the requested unique hash identifier, then the graphics context manager 128 allocates a new context using the requested unique hash identifier.
20 However, if the hash table 130 does contain the requested unique hash identifier (thereby informing that the requested unique hash identifier corresponds to a state that is already active at the GPU 104), then the graphics context manager 128 returns that context.

Instead of allocating a new context or potentially waiting for a context to free up
25 (thereby resulting in a graphics pipeline stall), the GPU 104 performs context bouncing by switching back to the existing, matching context that is already running at the GPU 104 without allocating a new context (e.g., in a manner similar to how a cache memory shares cache lines). Accordingly, the GPU 104 can run non-sequential processes on multiple state sets provided that the state sets are still
30 available (i.e., loaded at the GPU 104). This approach more efficiently utilizes the onboard state of the GPU 104 and reduces a number of stalls associated with running out of contexts to use, especially when the application 112 is quickly toggling

between a small number of states (which would involve context rolling and acquiring new states as the state changes after each draw in conventional context management techniques, and quickly running out of context).

FIG. 2 is a block diagram illustrating an embodiment in which the CPU 102
5 sends instructions intended for the GPU 104 in accordance with some embodiments. As illustrated, the CPU 102 sends instructions to the command buffer 118. The command buffer 118 temporarily stores commands of the command stream from the CPU 102. The stream of commands includes, for example, conditional execution packets, state update packets, and draw call packets. In various embodiments, other
10 commands may be included in the stream of commands as would be apparent to a person skilled in the art.

A conditional execute packet is a command inserted by the user mode driver 116 of FIG. 1 (or other software driver) which provides the hash identifier corresponding to one or more context state packets following the conditional execute
15 packet. In some embodiments, the conditional execute packet also includes dword data corresponding to the size of the next group of context state packets in dwords (which is short for "double word"). As the context state packets vary in size from a single dword to many dwords, the dword data defines a number N of context state packages associated with the hash identifier of the conditional execute packet and
20 further represents all the context register updates required to setup the context state for the next draw(s) following that hash identifier. Additionally, as further described herein, the dword data also defines a number of dwords to discard if the graphics context manager 128 returns and indicates that context associated with that hash identifier is already active at the GPU 104. Based on that determination, the
25 command processor 122 pulls from the command buffer 118 based on that given dword data size until all the state associated with that particular hash ID (as provided by the conditional execute packet) is discarded.

A state update packet is a constant or a collection of constants that update the state of graphics pipeline 120. A state update packet may, for example, update
30 colors that are to be blended during execution of a draw command. In some embodiments, a state update packet includes a set context packet that programs multi-context registers of the GPU 104. The set context packet includes all data

required to program the state in the packet. In other embodiments, a state update packet also includes a load context packet which requires fetching context information from memory before the state is written to context registers of the GPU 104. A draw call packet is a command that causes graphics pipeline 120 to execute processes on data to be output for display.

The execution of a draw call is dependent on all the state updates that were retrieved since a previous draw call. For example, FIG. 2 illustrates six commands that are included in the command stream—(1) a conditional execute packet 202, (2) a first state update packet 204, (3) a second state update packet 206, (4) a third state update packet 208, (5) a first draw call packet 210, and (6) a second draw call packet 212. The draw call packet 210 is dependent on the first, second, and third state update packets 204-208 because these are the state updates that are encapsulated by the conditional execute packet 202 prior to the next draw command (i.e., draw call packet 210). As used herein, the term “encapsulated” refers to a number of packets used to define a given hash identifier. Accordingly, the term “encapsulated state” refers to a number of context state packets that can be conditionally executed or discarded depending on whether that given hash identifier is available (e.g., stored at registers of the GPU 104) to context bounce / switch to. In other words, an encapsulated state is a collection of context state packets used to define a given hash identifier.

In one embodiment, the output of operations executed by the GPU 104 is dependent on a current context state 214. In various embodiments, the current context state 214 is based on the values of various context-specific constants such as texture handlers, shader constants, transform matrices, and the like which are stored in state registers. As illustrated in FIG. 2, the current context state 214 includes registers 216A through 216N, which are representative of any number and type (e.g., general purpose register) of state registers and instruction pointer. It is noted that current context state 214 can also include other variables and other values which define the current architectural state of the GPU 104. It is also noted that a “context state” of a processing unit can also be referred to as a “context” or as a “state” of the processing unit. In various embodiments, the processing unit is a GPU (e.g., GPU 104 of FIGs. 1-2). In other embodiments, the processing unit includes

other types of processing units. It is noted that the terms “processing unit” and “processor” are utilized interchangeably herein.

In operation, the command processor 122 receives the conditional execute packet 202. The conditional execute packet 202 informs, before the next draw command (e.g., draw call packet 210), the command processor 122 of the unique hash identifier associated with the upcoming state update packets 204-208. The command processor 122 queries the context manager 128 to search the hash table 130 to determine whether the unique hash identifier provided by the conditional execute packet 202 is active at the GPU 104 (and therefore stored at the hash table 130).

As illustrated in FIG. 2, the conditional execute packet 202 provides a unique hash identifier of XYZ. The command processor 122 queries the context manager 128 and determines that hash identifier XYZ is not present in the hash table 130. Accordingly, based on that query miss, the context manager 128 retires one of the existing hash identifiers in hardware (as tracked at the hash table 130) and allocates a new context. In various embodiments, the context manager 128 retires the least recently used (LRU) inactive hash identifier in hash table 130. The retired hash identifier is replaced with the hash identifier supplied for the query (e.g., hash identifier XYZ in this discussion) for that new context. Additionally, based on that determination of no match for the hash identifier XYZ at hash table 130, the context manager 128 indicates the query miss back to the command processor 122 and returns the new context. The command processor 122 then processes the encapsulated state by executing the state update packets 204-208. In other words, a query miss causes the GPU 104 to perform context rolling to allocate a new context, execute the upcoming state update packets 204-208, and set up the next state with the new hash identifier (i.e., hash identifier XYZ).

FIG. 3 is a block diagram illustrating draining of encapsulated state in accordance with some embodiments. In another embodiment, but with continued reference to FIG. 2, a command stream 300 at command buffer 118 includes five commands—(1) a conditional execute packet 302, (2) a first state update packet 304, (3) a second state update packet 306, (4) a third state update packet 308, and (5) a draw call packet 310. The draw call packet 310 is dependent on the first, second,

and third state update packets 304-308 because these are the state updates that are encapsulated by the conditional execute packet 302 prior to the next draw command (i.e., draw call packet 310). The first, second, and third state update packets 304-308 are collectively referred to as the encapsulated state 312. In this embodiment, the
5 encapsulated state 312 includes three of the context state packets (i.e., as defined by conditional execute packet 302) that can be conditionally executed or discarded depending on whether that given hash identifier is available (e.g., stored at registers of the GPU 104) to context bounce / switch to.

As illustrated in FIG. 3, the conditional execute packet 302 provides a unique
10 hash identifier of ABC. Referring back to FIG. 2, the command processor 122 queries the context manager 128 and determines that hash identifier ABC is present in the hash table 130. Accordingly, based on that query hit, the context manager 128 returns the context associated with the queried hash identifier (i.e., hash identifier ABC). The command processor 122 then drains the encapsulated state 312 (e.g.,
15 skipping and discarding state update packets 304-308) and assigns the context matching the query hit for execution of the draw call packet 310 using an existing context hash identifier. By determining that a desired hash identifier is available for use, switching to the context for that hash identifier, and discarding any command packets that would have redefined that context, the GPU 104 improves operating
20 efficiency by realizing that the requested hash identifier is currently available at GPU hardware and switching to it rather than forcing the processor to stall from lack of available contexts.

FIG. 4 is a flow diagram illustrating a method 400 of graphics context bouncing in accordance with some embodiments. At block 402, a command processor of a
25 graphics processing unit (GPU) receives a conditional execute packet providing a hash identifier corresponding to an encapsulated state. For example, with reference to FIG. 2, the command processor 122 receives conditional execute packet 202 from the command buffer 118. The conditional execute packet 202 defines an encapsulated state corresponding to one or more context state packets following the
30 conditional execute packet (e.g., state update packets 204-208 of FIG. 2). In various embodiments, the conditional execute packet 202 defines a unique hash identifier associated with the encapsulated state (e.g., hash identifier XYZ at conditional

execute packet 202 of FIG. 2 or hash identifier ABC at conditional execute packet 302 of FIG. 3). Additionally, in some embodiments, the conditional execute packet 202 also includes dword data corresponding to the size of the group of context state packets for the encapsulated state.

5 At block 404, the command processor determines whether the hash identifier of the encapsulated state matches one of a plurality of hash identifiers of active context states currently stored at the GPU. For example, with reference to FIG. 2, the command processor 122 queries the context manager 128 to search hash table 130 storing a plurality of hash identifiers of active context states currently stored at the
10 GPU.

 If the command processor determines that the hash identifier of the encapsulated state does not matches any of the plurality of hash identifiers of active context states currently stored at the GPU, the method 400 proceeds to block 406. However, if the command processor determines that the hash identifier of the
15 encapsulated state matches one of the plurality of hash identifiers of active context states currently stored at the GPU, the method 400 proceeds to block 408.

 At block 406, and with continued reference to FIG. 2, the context manger 128 retires one of the inactive existing hash identifiers in hardware (as tracked at the hash table 130) and allocates a new context based on the query miss from context
20 manager 128 determining that hash identifier XYZ is not present in the hash table 130. In some embodiments, the context manager 128 retires the least recently used (LRU) hash identifier in hash table 130. At block 410, the context manger 128 replaces the retired hash identifier with the hash identifier supplied for the query (e.g., hash identifier XYZ in this discussion) for that new context. Additionally, in various
25 embodiments, the context manager 128 indicates the query miss back to the command processor 122 and returns the new context.

 At block 412, the command processor 122 processes the encapsulated state by executing the state update packets 204-208. In other words, a query miss causes the GPU 104 to perform context rolling to allocate a new context state set, execute
30 the upcoming state update packets 204-208, and set up the next state with the new hash identifier (i.e., hash identifier XYZ). Subsequently, at block 414, the command

processor 122 executes the command packet following the encapsulated state (e.g., state draw call packet 210 of FIG. 2).

Returning now to block 408, and with reference to FIG. 3, the context manager 128 returns the context associated with the queried hash identifier (e.g., hash identifier ABC as provided by conditional execute packet 302 of FIG. 3) to the command processor 122 based on the query hit from context manager 128 determining that hash identifier ABC is present in the hash table 130. At block 416, the command processor 122 drains the encapsulated state 312 (e.g., skipping and discarding state update packets 304-308) and assigns the context matching the query hit for execution of the draw call packet 310 using an existing context hash identifier. Subsequently, at block 418, the command processor 122 executes the command packet following the bypassed encapsulated state 312. By determining that a desired hash identifier is available for use, switching to the context for that hash identifier, and discarding any command packets that would have redefined that context, the GPU 104 improves operating efficiency by realizing that the requested hash identifier is currently available at GPU hardware and switching to it rather than forcing the processor to stall from lack of available contexts.

In various embodiments, a context done event is sent down the graphics pipeline when a context is rolled. The context manager increments a counter indicating that there is a context done event outstanding. Further, in various embodiments, a context is only allowed to recycle when there are no pending context done events in the graphics pipeline.

For illustrative purposes, and not limitation, a graphics context bouncing system in accordance with various embodiments is described herein in terms of a GPU that executes a stream of draw commands based on context states of a graphics pipeline of the GPU. It is to be appreciated, however, that such a graphics context bouncing system is applicable to other types of ASICs or systems that execute a stream of commands based on the state of the respective systems. Based on the description contained herein, a person skilled in the relevant art(s) will understand how to implement embodiments of the present invention in such other types of systems. Additionally, those skilled in the art will recognize that numerous details are set forth herein (e.g., specific numbers and arrangements of processor configurations,

microarchitectural details, logic partitioning/integration details, sequences of operations, types, interrelationships of system components, numbers of unique context state sets stored at the GPU, and the like). However, it is understood that the graphics context bouncing described herein may be performed for any number of individual registers, register banks, and/or unique sets of context state without departing from the scope of this disclosure.

.As disclosed herein, in some embodiments a method includes: receiving, at a command processor of a graphics processing unit (GPU), a conditional execute packet providing an identifier corresponding to an encapsulated state, wherein the encapsulated state includes one or more context state packets following the conditional execute packet; and executing, based at least in part on determining whether the identifier of the encapsulated state matches one of a plurality of identifiers of active context states currently stored at the GPU, a command packet following the encapsulated state. In one aspect, determining whether the identifier of the encapsulated state matches one of a plurality of identifiers includes querying a context manager to search an identifier table storing the plurality of identifiers of active context states currently stored at the GPU. In another aspect, the method includes: draining, in response to determining the identifier matching one of the plurality of identifiers of active context states, the encapsulated state; and assigning the matching one of the plurality of identifiers of active context states to the command packet following the encapsulated state.

In one aspect, the method includes: allocating, in response to determining the identifier does not match any of the plurality of identifiers of active context states, a new context state set; and executing the one or more context state packets of the encapsulated state prior to executing the command packet. In another aspect, the method includes: retiring, from the identifier table, one the plurality of identifiers of active context states currently stored at the GPU; and storing, at the identifier table, the identifier corresponding to the encapsulated state for the new context state set. In still another aspect, the command packet includes a draw call. In yet another aspect, the command packet designates a size of the encapsulated state for a number of state update packets.

In some embodiments, a processor includes: a plurality of state registers; and a command processor, wherein the command processor is configured to: receive a conditional execute packet providing an identifier corresponding to an encapsulated state, wherein the encapsulated state includes one or more context state packets
5 following the conditional execute packet; and execute, based at least in part on determining whether the identifier of the encapsulated state matches one of a plurality of identifiers of active context states currently stored at the processor, a command packet following the encapsulated state. In one aspect, the processor includes: a context manager including an identifier table storing the plurality of
10 identifiers of active context states currently stored at the processor. In another aspect, the processor is configured to query the context manager to search the identifier table for the provided identifier.

In one aspect, the processor is further configured to: drain, in response to determining the identifier matching one of the plurality of identifiers of active context
15 states, the encapsulated state; and assign the matching one of the plurality of identifiers of active context states to the command packet following the encapsulated state. In another aspect, the processor is further configured to: allocate, in response to determining the identifier does not match any of the plurality of identifiers of active context states, a new context state set; and execute the one or more context state
20 packets of the encapsulated state prior to executing the command packet. In yet another aspect, the processor is further configured to: retire, from the identifier table, one the plurality of identifiers of active context states currently stored at the processor; and store, at the identifier table, the identifier corresponding to the encapsulated state for the new context state set. In still another aspect, the
25 command packet designates a size of the encapsulated state for a number of state update packets.

In some embodiments, a non-transitory computer readable medium embodies a set of executable instructions, the set of executable instructions to manipulate at least one processor to: receive, at a command processor of the processor, a conditional
30 execute packet providing an identifier corresponding to an encapsulated state, wherein the encapsulated state includes one or more context state packets following the conditional execute packet; and execute, based at least in part on determining

whether the identifier of the encapsulated state matches one of a plurality of identifiers of active context states currently stored at the processor, a command packet following the encapsulated state. In one aspect, the set of executable instructions is to further manipulate at least one processor to query a context manager to search an identifier table storing the plurality of identifiers of active context states currently stored at the processor. In another aspect, the set of executable instructions is to further manipulate at least one processor to: drain, in response to determining the identifier matching one of the plurality of identifiers of active context states, the encapsulated state; and assign the matching one of the plurality of identifiers of active context states to the command packet following the encapsulated state.

In one aspect, the set of executable instructions is to further manipulate at least one processor to: allocate, in response to determining the identifier does not match any of the plurality of identifiers of active context states, a new context state set; and execute the one or more context state packets of the encapsulated state prior to executing the command packet. In another aspect, the set of executable instructions is to further manipulate at least one processor to: retire, from the identifier table, one the plurality of identifiers of active context states currently stored at the processor; and store, at the identifier table, the identifier corresponding to the encapsulated state for the new context state set. In still another aspect, the command packet designates a size of the encapsulated state for a number of state update packets.

A computer readable storage medium may include any non-transitory storage medium, or combination of non-transitory storage media, accessible by a computer system during use to provide instructions and/or data to the computer system. Such storage media can include, but is not limited to, optical media (e.g., compact disc (CD), digital versatile disc (DVD), Blu-Ray disc), magnetic media (e.g., floppy disc, magnetic tape, or magnetic hard drive), volatile memory (e.g., random access memory (RAM) or cache), non-volatile memory (e.g., read-only memory (ROM) or Flash memory), or microelectromechanical systems (MEMS)-based storage media. The computer readable storage medium may be embedded in the computing system (e.g., system RAM or ROM), fixedly attached to the computing system (e.g., a magnetic hard drive), removably attached to the computing system (e.g., an optical

disc or Universal Serial Bus (USB)-based Flash memory), or coupled to the computer system via a wired or wireless network (e.g., network accessible storage (NAS)).

In some embodiments, certain aspects of the techniques described above may implemented by one or more processors of a processing system executing
5 software. The software includes one or more sets of executable instructions stored or otherwise tangibly embodied on a non-transitory computer readable storage medium. The software can include the instructions and certain data that, when executed by the one or more processors, manipulate the one or more processors to perform one or more aspects of the techniques described above. The non-transitory
10 computer readable storage medium can include, for example, a magnetic or optical disk storage device, solid state storage devices such as Flash memory, a cache, random access memory (RAM) or other non-volatile memory device or devices, and the like. The executable instructions stored on the non-transitory computer readable storage medium may be in source code, assembly language code, object code, or
15 other instruction format that is interpreted or otherwise executable by one or more processors.

Note that not all of the activities or elements described above in the general description are required, that a portion of a specific activity or device may not be required, and that one or more further activities may be performed, or elements
20 included, in addition to those described. Still further, the order in which activities are listed are not necessarily the order in which they are performed. Also, the concepts have been described with reference to specific embodiments. However, one of ordinary skill in the art appreciates that various modifications and changes can be made without departing from the scope of the present disclosure as set forth in the
25 claims below. Accordingly, the specification and figures are to be regarded in an illustrative rather than a restrictive sense, and all such modifications are intended to be included within the scope of the present disclosure.

Benefits, other advantages, and solutions to problems have been described above with regard to specific embodiments. However, the benefits, advantages,
30 solutions to problems, and any feature(s) that may cause any benefit, advantage, or solution to occur or become more pronounced are not to be construed as a critical, required, or essential feature of any or all the claims. Moreover, the particular

embodiments disclosed above are illustrative only, as the disclosed subject matter may be modified and practiced in different but equivalent manners apparent to those skilled in the art having the benefit of the teachings herein. No limitations are intended to the details of construction or design herein shown, other than as

5 described in the claims below. It is therefore evident that the particular embodiments disclosed above may be altered or modified and all such variations are considered within the scope of the disclosed subject matter. Accordingly, the protection sought herein is as set forth in the claims below.

WHAT IS CLAIMED IS:

1. A method, comprising:

receiving, at a command processor [122] of a graphics processing unit (GPU) [104], a conditional execute packet [202] providing an identifier corresponding to an encapsulated state, wherein the encapsulated state includes one or more context state packets [204, 206] following the conditional execute packet; and

executing, based at least in part on determining whether the identifier of the encapsulated state matches one of a plurality of identifiers of active context states currently stored at the GPU, a command packet [210] following the encapsulated state.

2. The method of claim 1, wherein determining whether the identifier of the encapsulated state matches one of a plurality of identifiers comprises:

querying a context manager [128] to search an identifier table [130] storing the plurality of identifiers of active context states currently stored at the GPU.

3. The method of claim 2, further comprising:

draining, in response to determining the identifier matching one of the plurality of identifiers of active context states, the encapsulated state; and assigning the matching one of the plurality of identifiers of active context states to the command packet following the encapsulated state.

4. The method of claim 2, further comprising:

allocating, in response to determining the identifier does not match any of the plurality of identifiers of active context states, a new context state set; and

executing the one or more context state packets of the encapsulated state prior to executing the command packet.

5. The method of claim 4, further comprising:
retiring, from the identifier table, one the plurality of identifiers of active context
states currently stored at the GPU; and
storing, at the identifier table, the identifier corresponding to the encapsulated
5 state for the new context state set.
6. The method of claim 1, wherein the command packet includes a draw call.
7. The method of claim 1, wherein the command packet designates a size of the
encapsulated state for a number of state update packets.
8. A processor [100], comprising:
10 a plurality of state registers [216]; and
a command processor [122], wherein the command processor is configured to:
receive a conditional execute packet [202] providing an identifier
corresponding to an encapsulated state, wherein the
encapsulated state includes one or more context state packets
15 [204, 206] following the conditional execute packet; and
execute, based at least in part on determining whether the identifier of
the encapsulated state matches one of a plurality of identifiers of
active context states currently stored at the processor, a
command packet [210] following the encapsulated state.
- 20 9. The processor of claim 8, further comprising:
a context manager [128] including an identifier table [130] storing the plurality
of identifiers of active context states currently stored at the processor.
10. The processor of claim 9, wherein the processor is configured to query the
context manager to search the identifier table for the provided identifier.
- 25 11. The processor of claim 9, wherein the processor is further configured to:
drain, in response to determining the identifier matching one of the plurality of
identifiers of active context states, the encapsulated state; and

assign the matching one of the plurality of identifiers of active context states to the command packet following the encapsulated state.

12. The processor of claim 9, wherein the processor is further configured to:

allocate, in response to determining the identifier does not match any of the

5 plurality of identifiers of active context states, a new context state set;

and

execute the one or more context state packets of the encapsulated state prior

to executing the command packet.

13. The processor of claim 12, wherein the processor is further configured to:

10 retire, from the identifier table, one the plurality of identifiers of active context states currently stored at the processor; and

store, at the identifier table, the identifier corresponding to the encapsulated

state for the new context state set.

14. The processor of claim 8, wherein the command packet designates a size of the

15 encapsulated state for a number of state update packets.

15. A non-transitory computer readable medium embodying a set of executable instructions, the set of executable instructions to manipulate at least one processor [100] to:

receive, at a command processor [122] of the processor, a conditional execute

20 packet [202] providing an identifier corresponding to an encapsulated

state, wherein the encapsulated state includes one or more context

state packets [204, 206] following the conditional execute packet; and

execute, based at least in part on determining whether the identifier of the

encapsulated state matches one of a plurality of identifiers of active

25 context states currently stored at the processor, a command packet

[210] following the encapsulated state.

16. The non-transitory computer readable medium of claim 15, the set of executable instructions to further manipulate at least one processor to:

query a context manager [128] to search an identifier table [130] storing the plurality of identifiers of active context states currently stored at the processor.

17. The non-transitory computer readable medium of claim 16, the set of executable
5 instructions to further manipulate at least one processor to:

drain, in response to determining the identifier matching one of the plurality of identifiers of active context states, the encapsulated state; and assign the matching one of the plurality of identifiers of active context states to the command packet following the encapsulated state.

10 18. The non-transitory computer readable medium of claim 16, the set of executable instructions to further manipulate at least one processor to:

allocate, in response to determining the identifier does not match any of the plurality of identifiers of active context states, a new context state set; and

15 execute the one or more context state packets of the encapsulated state prior to executing the command packet.

19. The non-transitory computer readable medium of claim 18, the set of executable instructions to further manipulate at least one processor to:

20 retire, from the identifier table, one the plurality of identifiers of active context states currently stored at the processor; and

store, at the identifier table, the identifier corresponding to the encapsulated state for the new context state set.

20. The non-transitory computer readable medium of claim 15, wherein the
25 command packet designates a size of the encapsulated state for a number of state update packets.

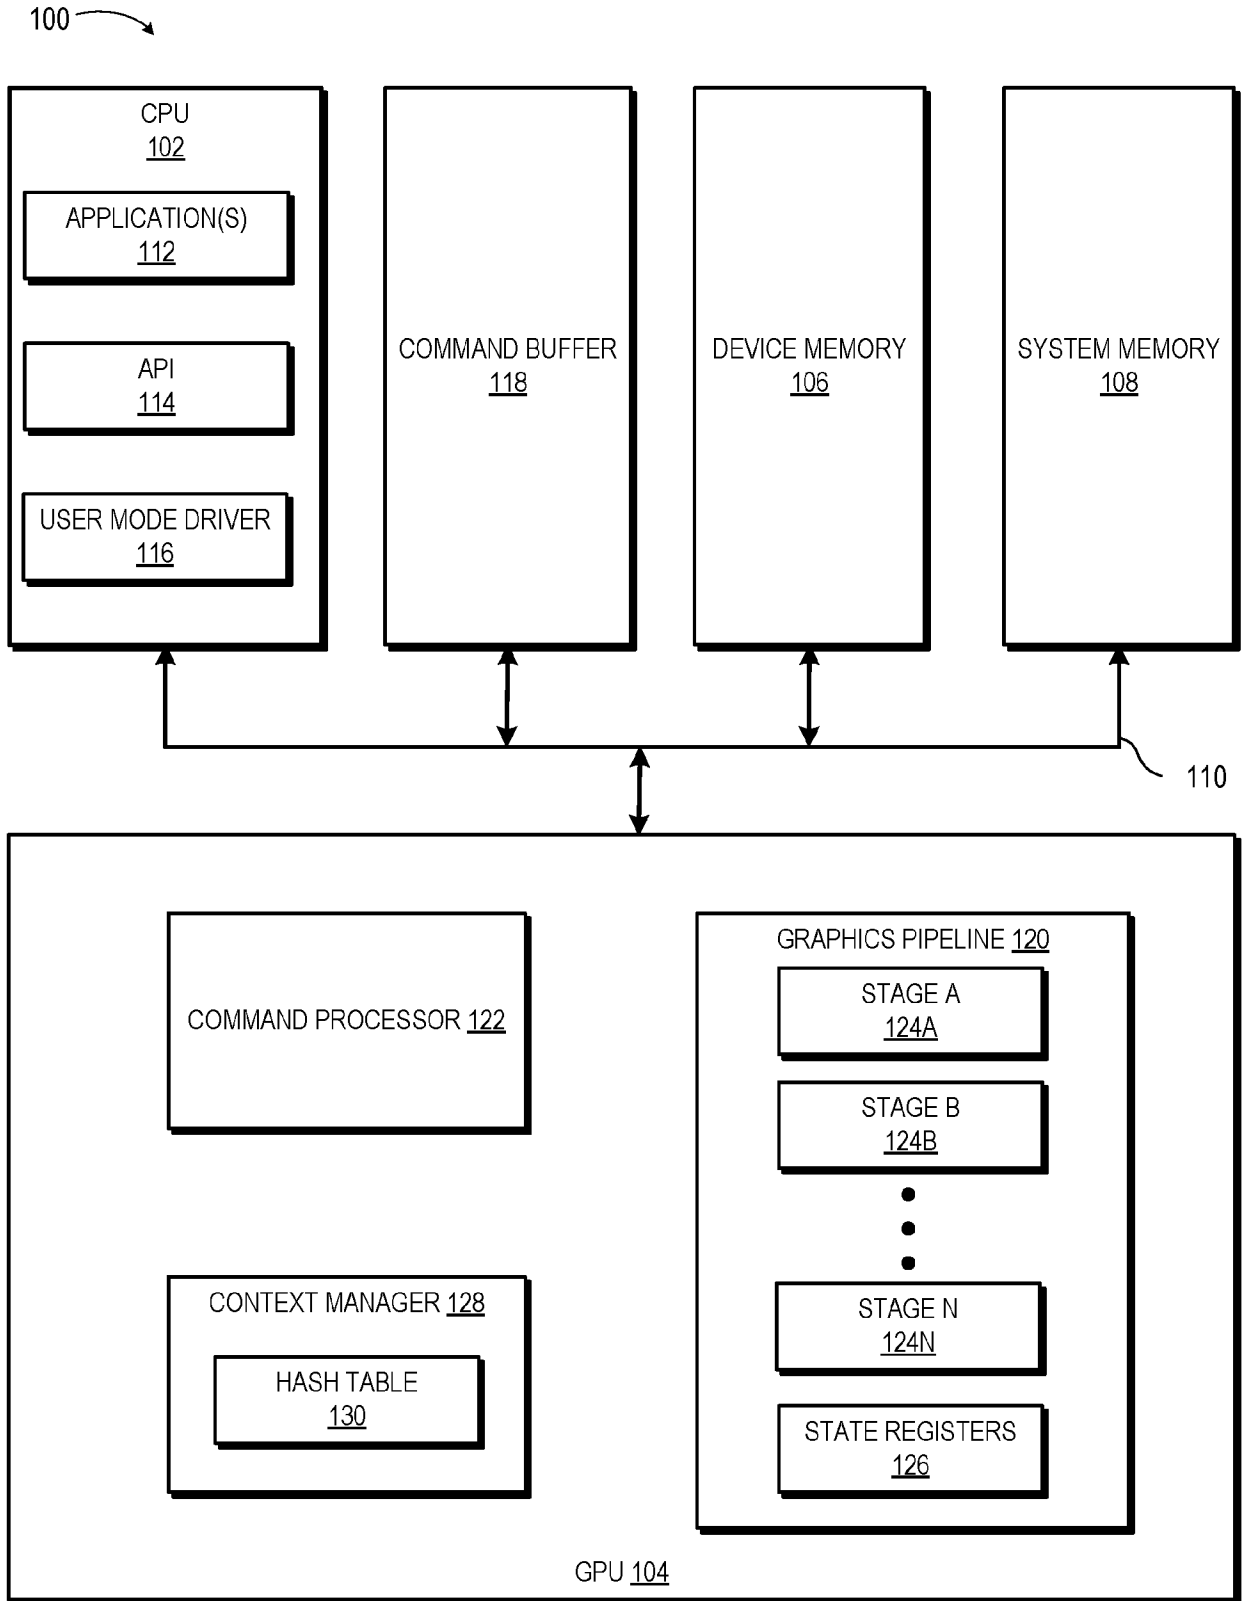


FIG. 1

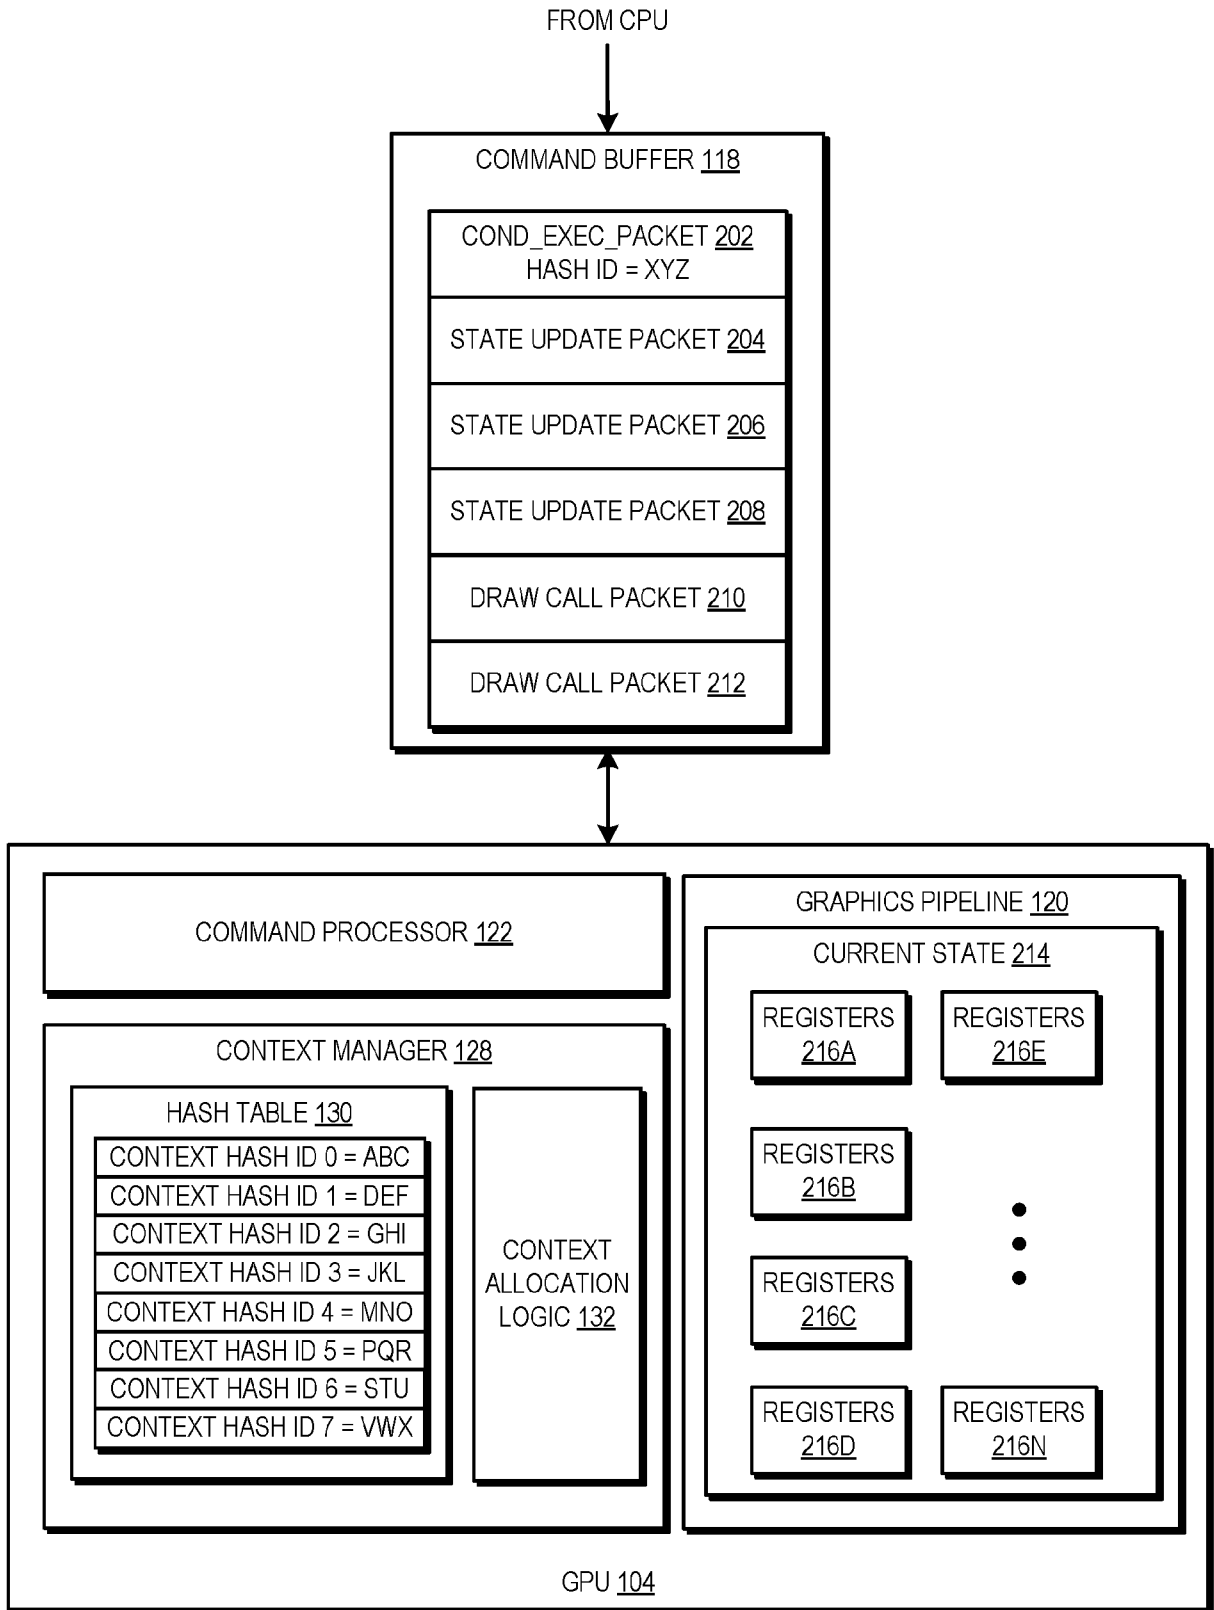


FIG. 2

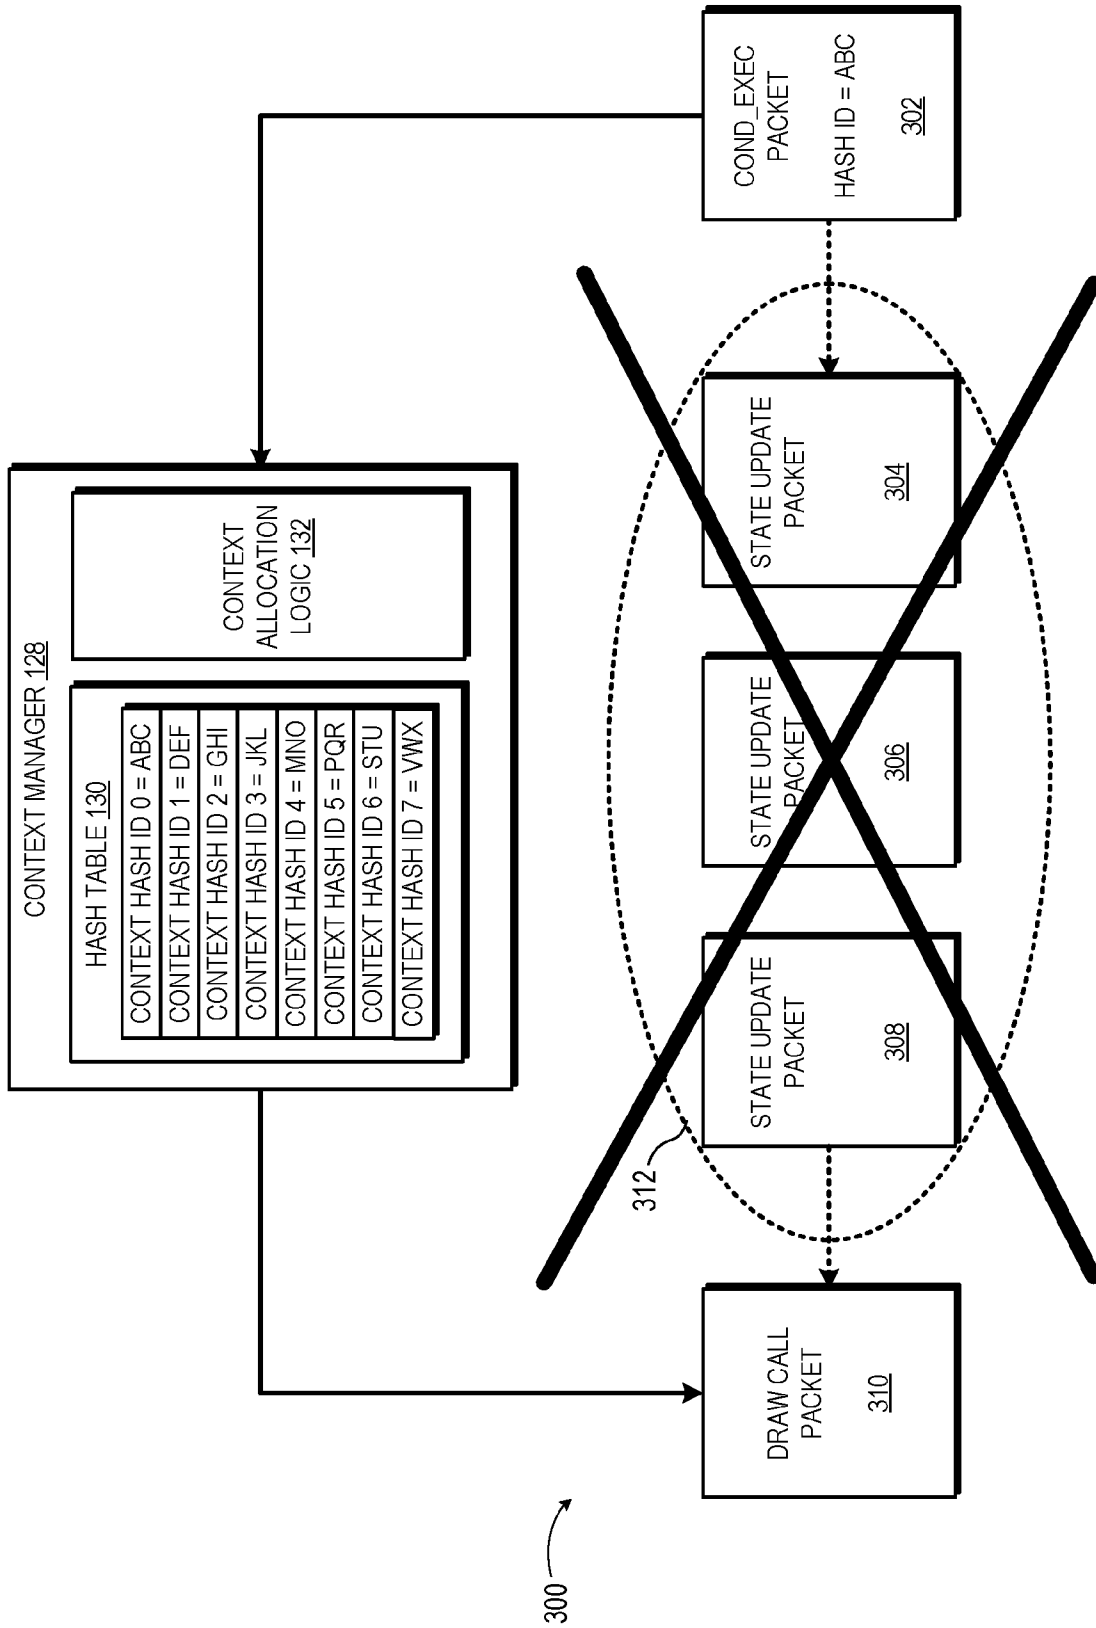


FIG. 3

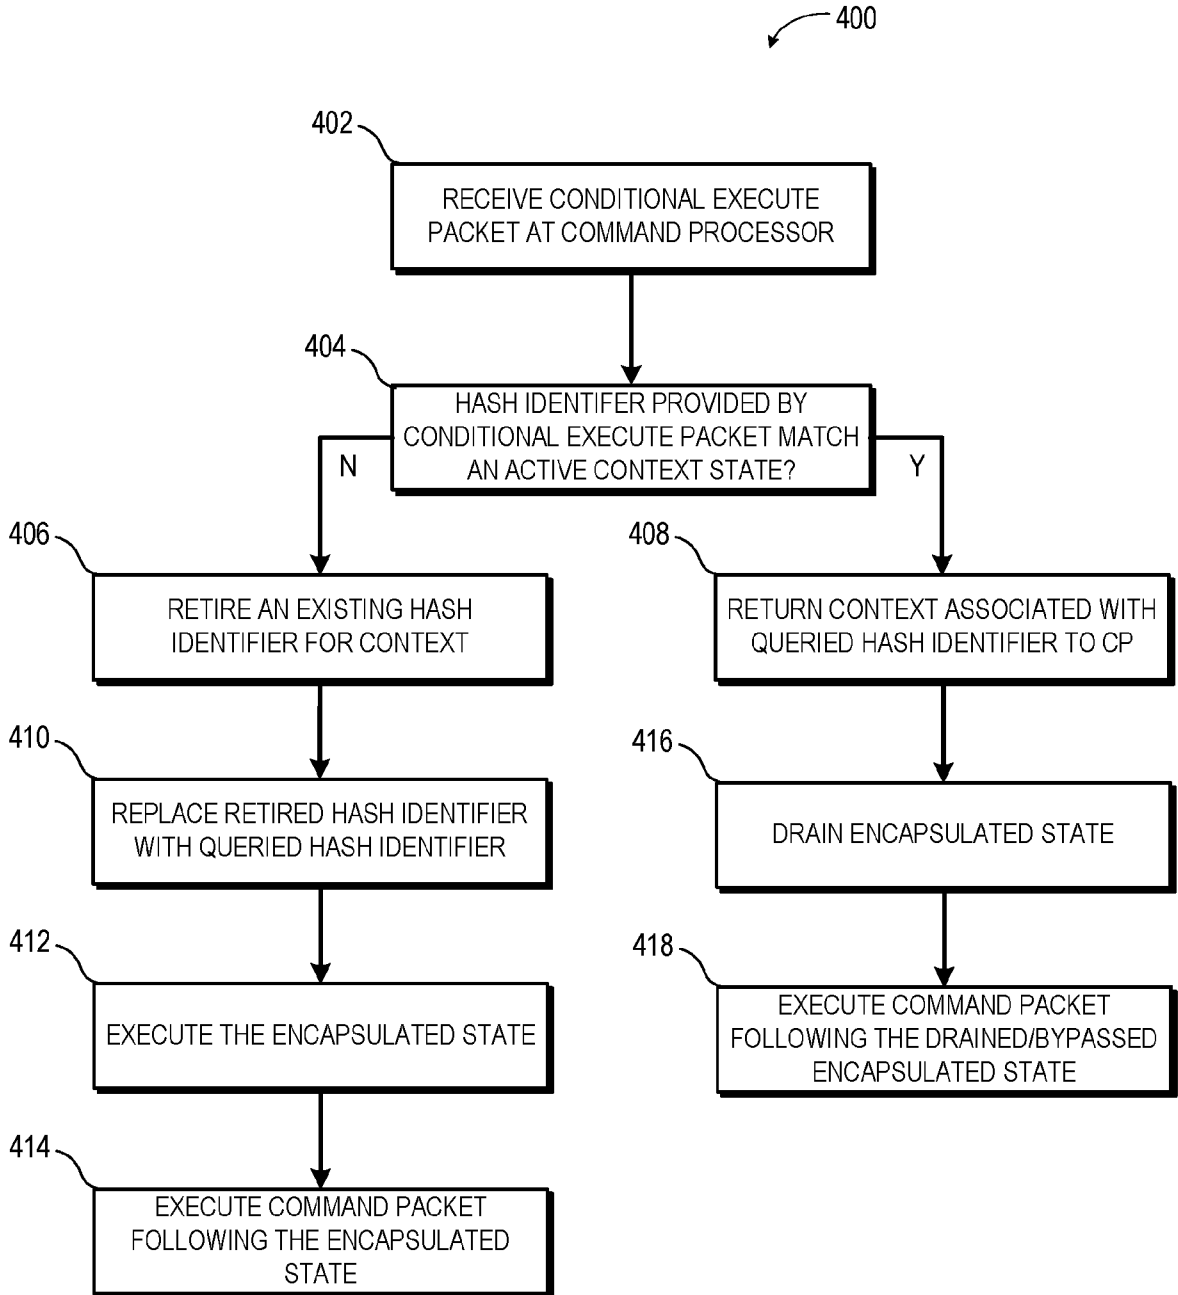


FIG. 4

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2020/035203**A. CLASSIFICATION OF SUBJECT MATTER****G06F 9/38(2006.01)i, G06F 9/48(2006.01)i, G06T 1/20(2006.01)i, G06F 12/1018(2016.01)i**

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 9/38; G06F 15/16; G06F 9/30; G06F 9/40; G06F 9/46; G06T 1/20; G06T 1/60; G06T 15/00; G06F 9/48; G06F 12/1018

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

Korean utility models and applications for utility models
Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: graphics processing unit (GPU), command, packet, identifier, context, match, encapsulate

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2016-0055609 A1 (ADVANCED MICRO DEVICES, INC.) 25 February 2016 paragraphs [0003], [0023]-[0025], [0031], [0033]; claims 5, 7; and figures 2-3	1-20
Y	US 2010-0131549 A1 (KINDLA H. KRAMER et al.) 27 May 2010 paragraphs [0407], [0425], [0573]	1-20
A	US 2018-0285109 A1 (INTEL CORPORATION) 04 October 2018 claims 1-4; and figures 1-8C	1-20
A	US 2015-0123980 A1 (HEMA CHAND NALLURI et al.) 07 May 2015 paragraphs [0026]-[0071]; and figures 1A-3C	1-20
A	US 7979683 B1 (JOHN M. DANSKIN et al.) 12 July 2011 column 13, line 53 - column 16, line 39; and figures 6-7	1-20

 Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:

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

"D" document cited by the applicant in the international application

"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

02 September 2020 (02.09.2020)

Date of mailing of the international search report

03 September 2020 (03.09.2020)

Name and mailing address of the ISA/KR

International Application Division
Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

YANG JEONG ROK

Telephone No. +82-42-481-5709



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2020/035203

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2016-0055609 A1	25/02/2016	CN 105469354 A	06/04/2016
US 2010-0131549 A1	27/05/2010	CN 101689152 B	05/10/2016
		CN 105122790 A	02/12/2015
		EP 1851750 A2	07/11/2007
		EP 2150893 A2	10/02/2010
		JP 2008-530661 A	07/08/2008
		JP 2010-525494 A	22/07/2010
		JP 2011-523112 A	04/08/2011
		JP 5905662 B2	20/04/2016
		JP 6116064 B2	19/04/2017
		KR 10-0960577 B1	03/06/2010
		KR 10-2015-0105308 A	16/09/2015
		NO 20074200 A	07/11/2007
		US 2015-0309581 A1	29/10/2015
		US 2017-0235617 A1	17/08/2017
		US 9842670 B2	12/12/2017
		WO 2010-129599 A1	11/11/2010
		WO 2010-138743 A2	02/12/2010
US 2018-0285109 A1	04/10/2018	US 10572258 B2	25/02/2020
US 2015-0123980 A1	07/05/2015	KR 10-1693662 B1	17/01/2017
		KR 10-2015-0051868 A	13/05/2015
		US 9563466 B2	07/02/2017
US 7979683 B1	12/07/2011	US 8095782 B1	10/01/2012