



US 20200264698A1

(19) **United States**

(12) **Patent Application Publication**
Veeramani et al.

(10) **Pub. No.: US 2020/0264698 A1**

(43) **Pub. Date: Aug. 20, 2020**

(54) **KEYBOARD FOR VIRTUAL REALITY**

Publication Classification

(71) Applicant: **Intel Corporation**, Santa Clara, CA (US)

(51) **Int. Cl.**
G06F 3/01 (2006.01)
G06F 3/03 (2006.01)
G06F 1/16 (2006.01)
G06F 3/042 (2006.01)
G02B 27/01 (2006.01)

(72) Inventors: **Karthik Veeramani**, Hillsboro, OR (US); **Jianfang Zhu**, Hillsboro, OR (US); **Sayan Lahiri**, Hillsboro, OR (US); **Bo Qiu**, Hillsboro, OR (US); **Bradley A. Jackson**, Hillsboro, OR (US); **Paul S. Diefenbaugh**, Portland, OR (US); **Kim Pallister**, Portland, OR (US)

(52) **U.S. Cl.**
CPC **G06F 3/011** (2013.01); **G06F 3/0304** (2013.01); **G06F 1/1686** (2013.01); **G02B 27/017** (2013.01); **G06F 3/0425** (2013.01); **G06F 3/0426** (2013.01); **G06F 3/017** (2013.01); **G06F 1/163** (2013.01)

(73) Assignee: **Intel Corporation**, Santa Clara, CA (US)

(57) **ABSTRACT**


(21) Appl. No.: **16/793,539**

An embodiment of a graphics apparatus may include an image generator, and a gesture tracker communicatively coupled to the image generator. The image generator may be configured to generate an image of a virtual input device, the gesture tracker may be configured to determine a position of a user's finger relative to the virtual input device, and the image generator may be further configured to generate an image of a virtual finger based on the determined position of the user's finger relative to the virtual input device. Other embodiments are disclosed and claimed.

(22) Filed: **Feb. 18, 2020**

Related U.S. Application Data

(63) Continuation of application No. 15/476,988, filed on Apr. 1, 2017, now Pat. No. 10,606,343.



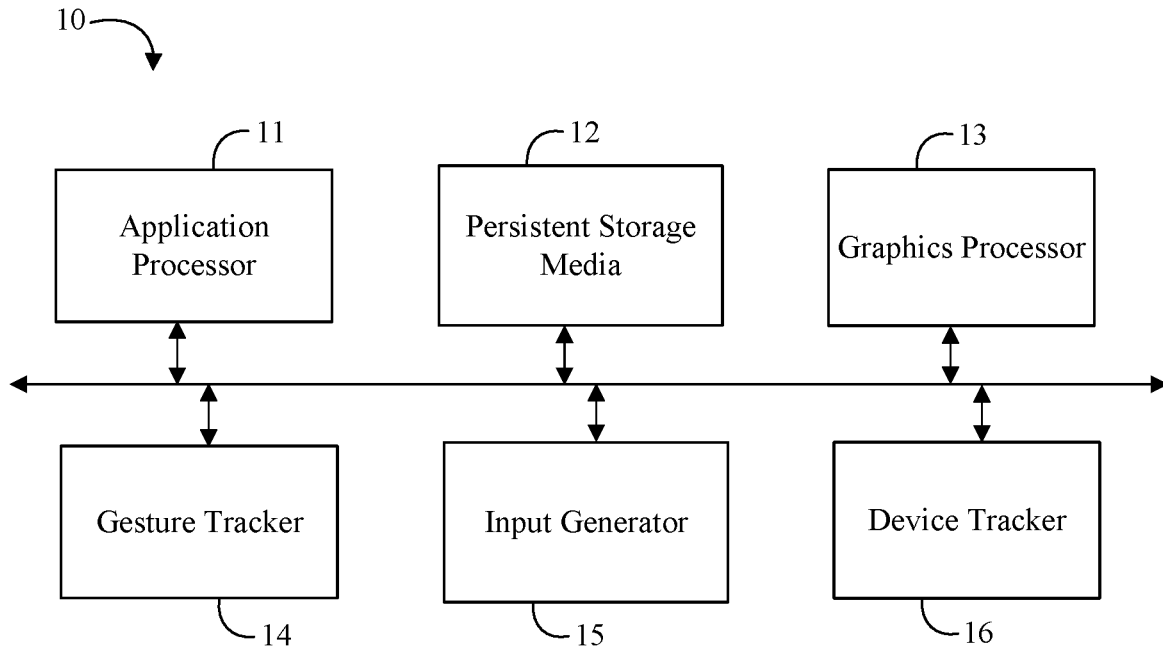


FIG. 1

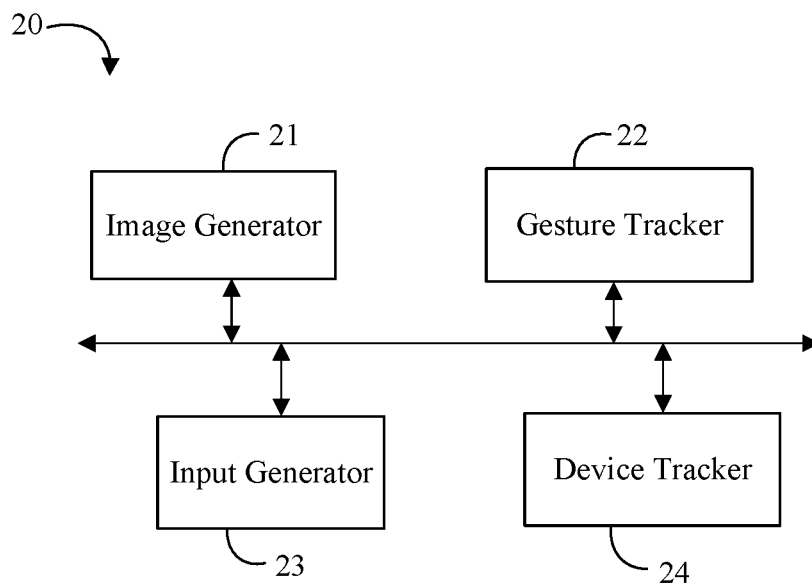


FIG. 2

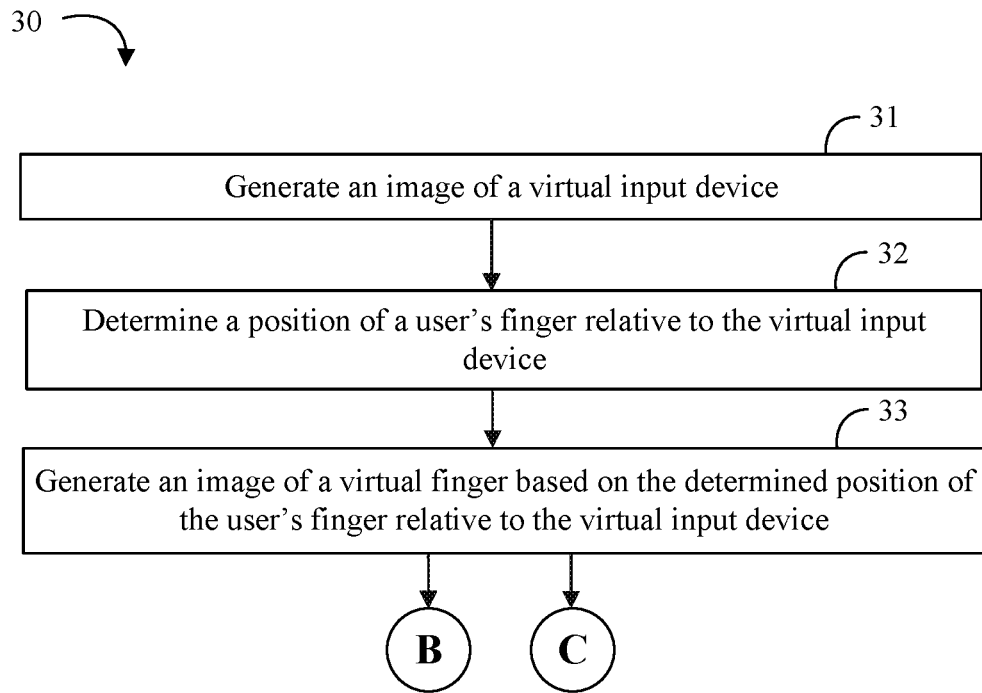


FIG. 3

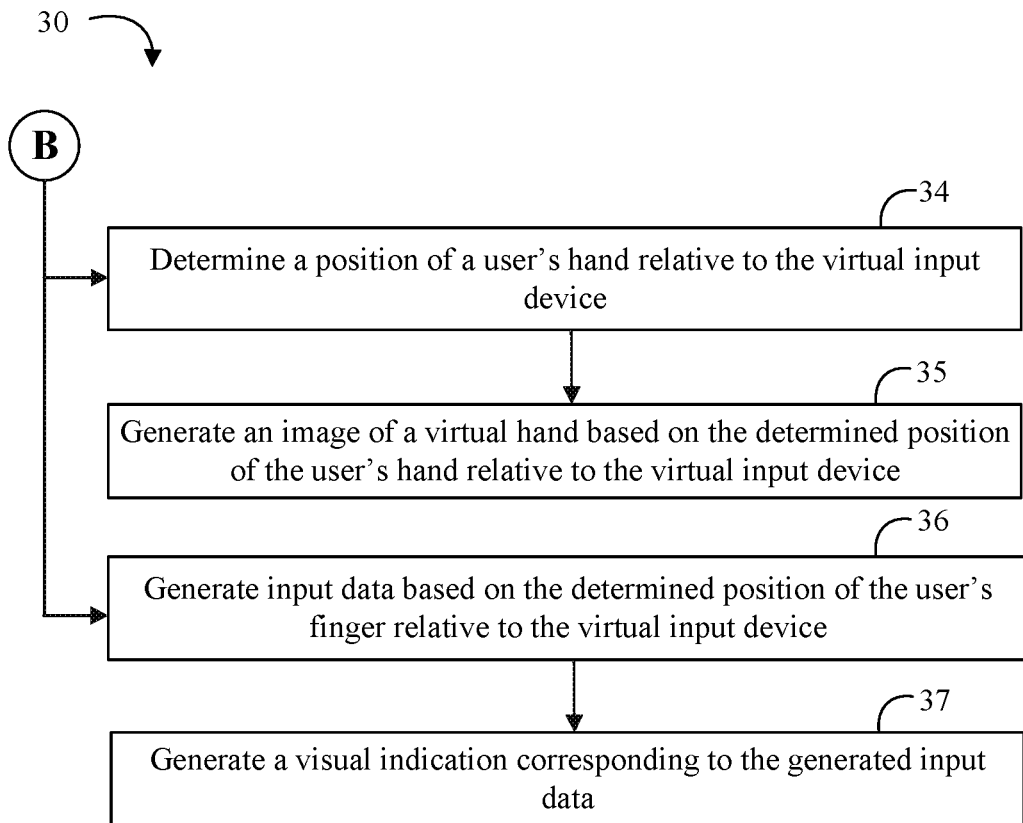


FIG. 4

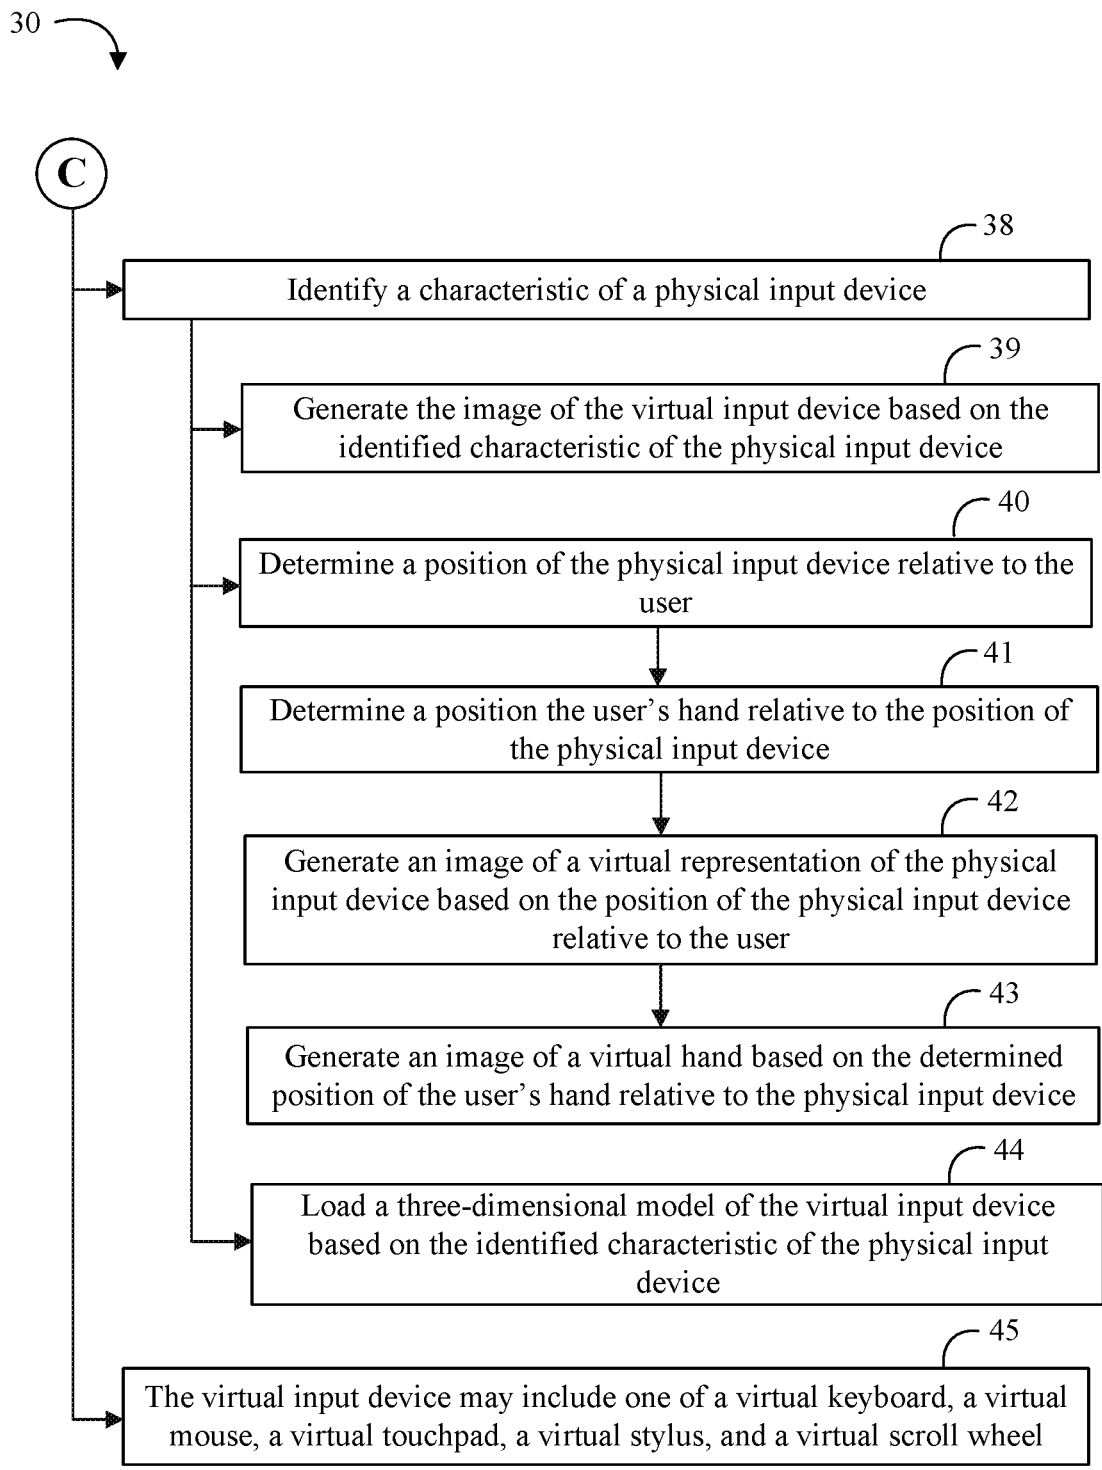


FIG. 5




FIG. 6A




FIG. 6B




FIG. 7

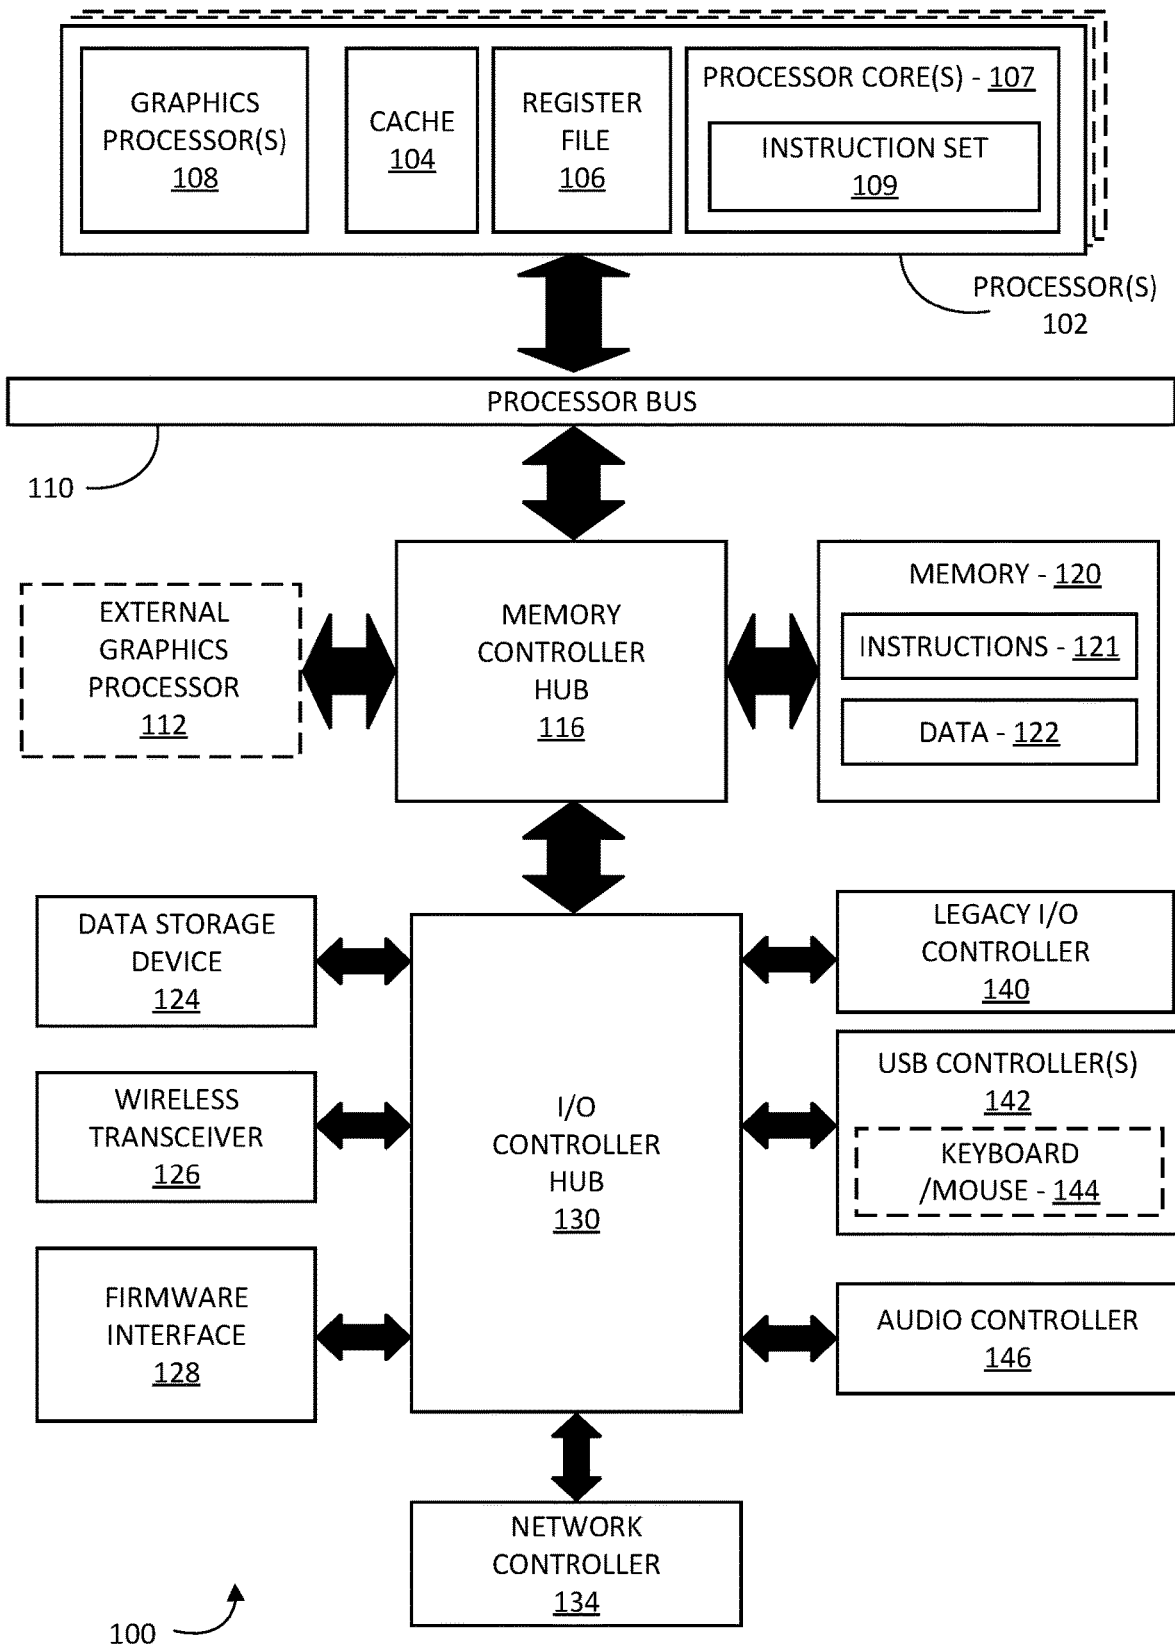


FIG. 8




FIG. 9

GRAPHICS PROCESSOR
300




FIG. 10




FIG. 11




FIG. 12

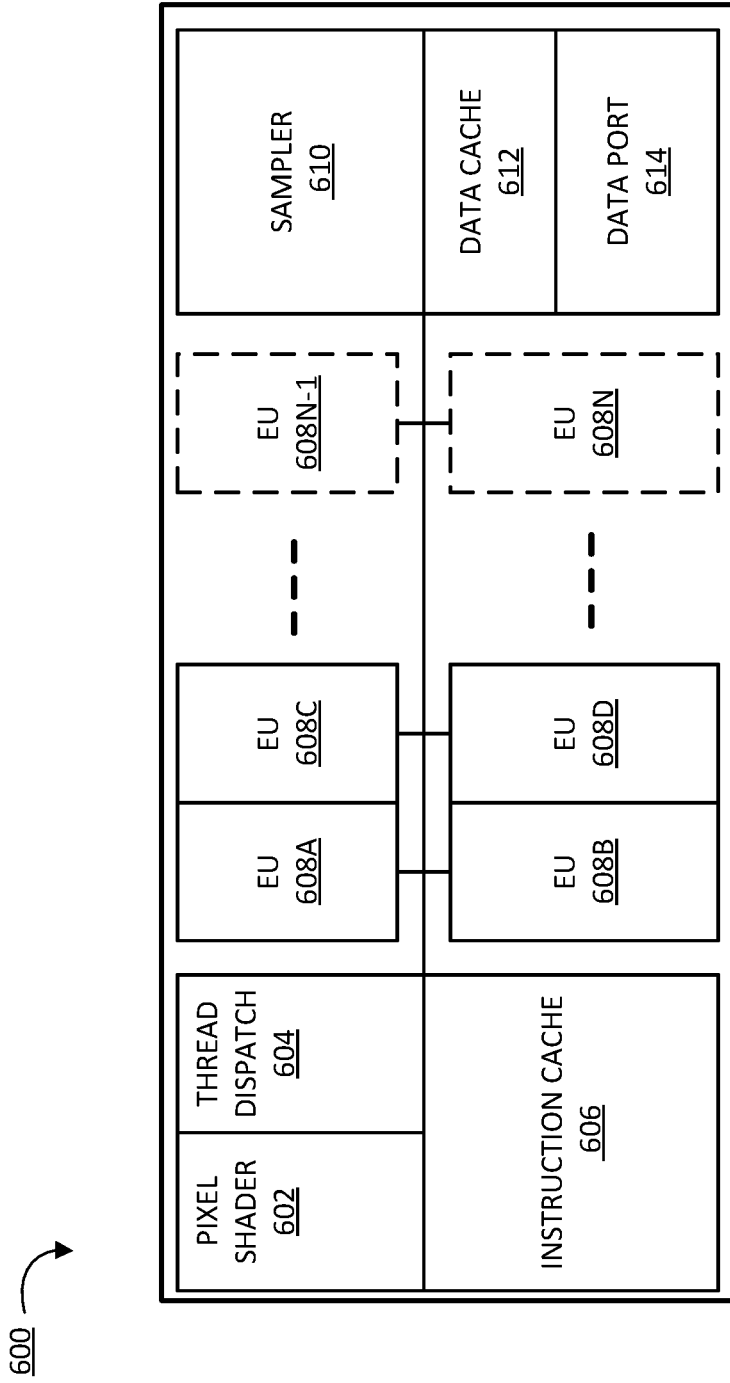


FIG. 13

GRAPHICS CORE INSTRUCTION FORMATS
700






FIG. 14




FIG. 15

Graphics Processor Command Format
900




FIG. 16A

Graphics Processor Command Sequence
910

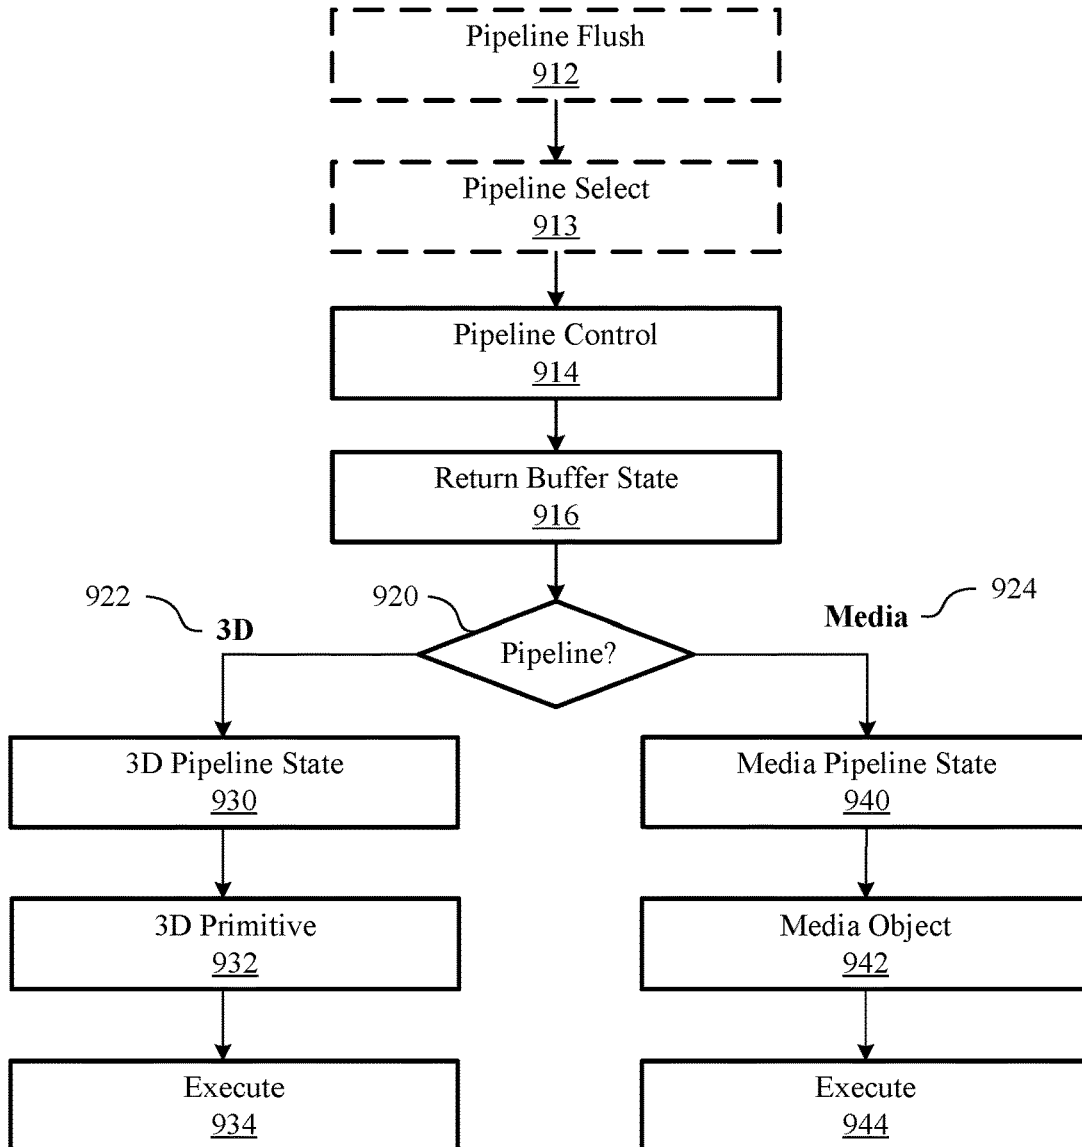


FIG. 16B

Data Processing System 1000

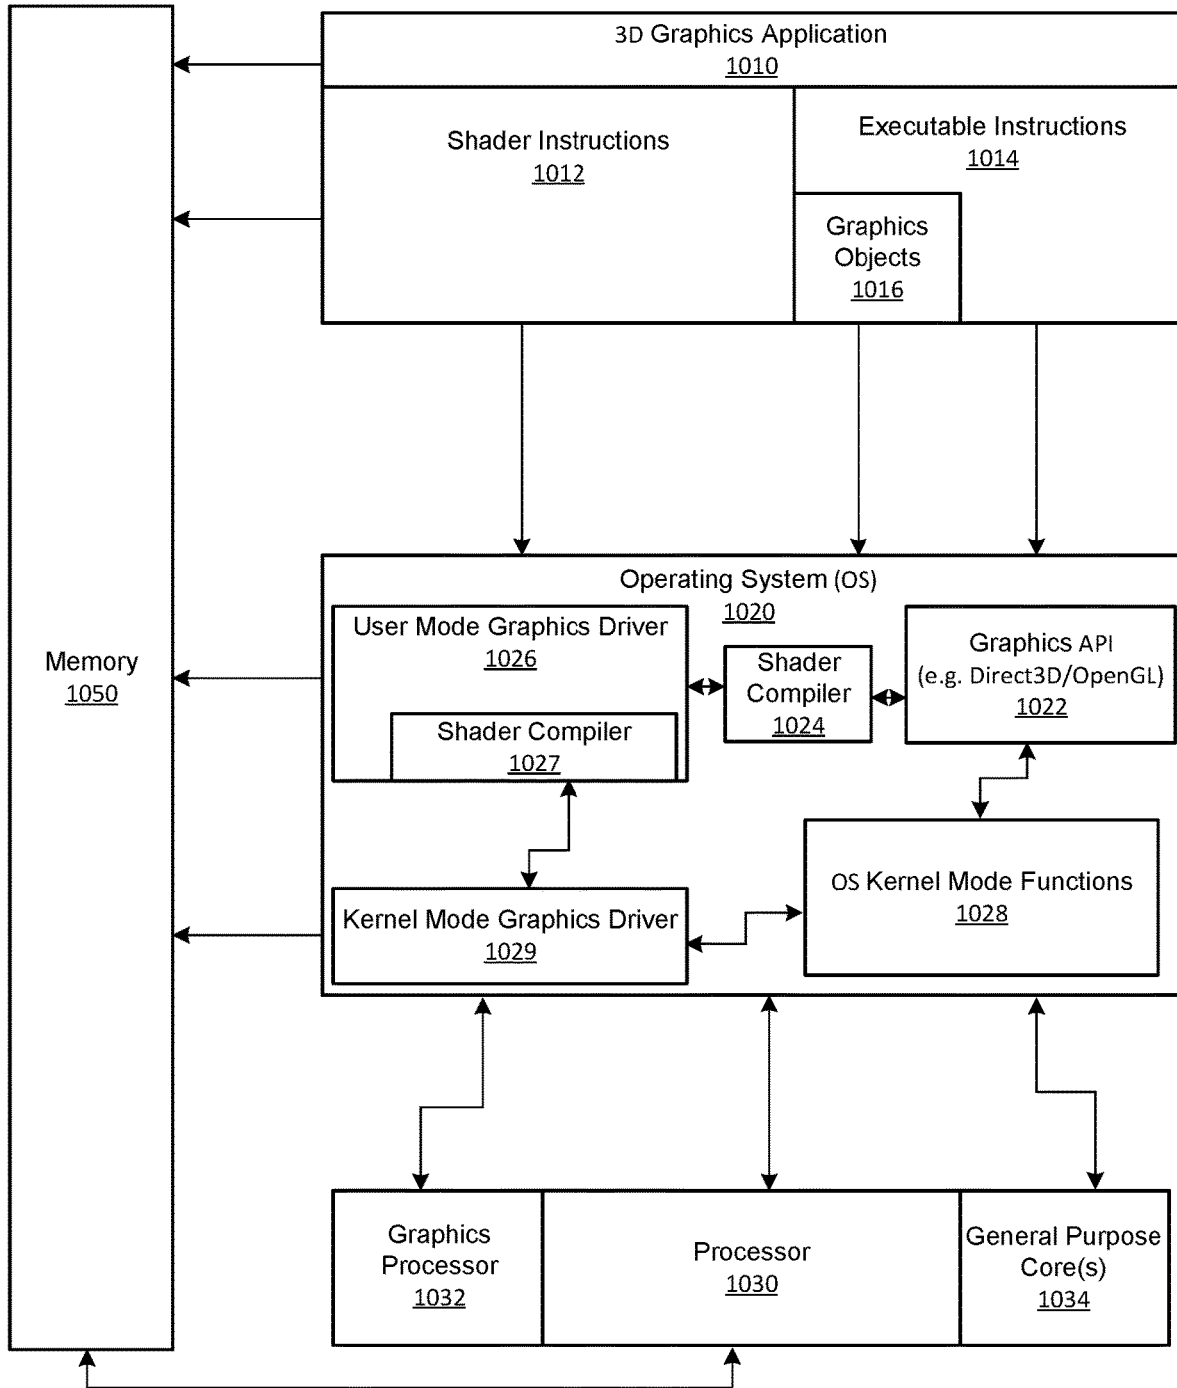


FIG. 17

IP CORE DEVELOPMENT - 1100




FIG. 18




FIG. 19

KEYBOARD FOR VIRTUAL REALITY**CROSS-REFERENCE TO RELATED APPLICATIONS**

[0001] The present application claims the benefit of priority to U.S. Non-Provisional patent application Ser. No. 15/476,988 filed on Apr. 1, 2017.

TECHNICAL FIELD

[0002] Embodiments generally relate to virtual reality (VR). More particularly, embodiments relate to a keyboard for virtual reality.

BACKGROUND

[0003] A VR environment may include a host system running a VR application connected to a head-mounted display (HMD) to immerse a user in a virtual environment. As the user moves or turns their head, the display may be updated to track the motion so the user may perceive what is displayed as appearing to be a realistic environment. In the VR environment, user input may typically take the form of sensed head movement, eye tracking or gesture recognition.

BRIEF DESCRIPTION OF THE DRAWINGS

[0004] The various advantages of the embodiments will become apparent to one skilled in the art by reading the following specification and appended claims, and by referencing the following drawings, in which:

[0005] FIG. 1 is a block diagram of an example of an electronic processing system according to an embodiment;

[0006] FIG. 2 is a block diagram of an example of a graphics apparatus according to an embodiment;

[0007] FIGS. 3 to 5 are flowcharts of an example of a method of displaying virtual objects according to an embodiment;

[0008] FIGS. 6A to 6B are illustrative diagrams of an example of a user with a VR system according to an embodiment;

[0009] FIG. 7 is an illustrative diagram of an example of displayed virtual objects according to an embodiment;

[0010] FIGS. 8-10 are block diagrams of an example of an overview of a data processing system according to an embodiment;

[0011] FIG. 11 is a block diagram of an example of a graphics processing engine according to an embodiment;

[0012] FIGS. 12-14 are block diagrams of examples of execution units according to an embodiment;

[0013] FIG. 15 is a block diagram of an example of a graphics pipeline according to an embodiment;

[0014] FIGS. 16A-16B are block diagrams of examples of graphics pipeline programming according to an embodiment;

[0015] FIG. 17 is a block diagram of an example of a graphics software architecture according to an embodiment;

[0016] FIG. 18 is a block diagram of an example of an intellectual property (IP) core development system according to an embodiment; and

[0017] FIG. 19 is a block diagram of an example of a system on a chip integrated circuit according to an embodiment.

DESCRIPTION OF EMBODIMENTS

[0018] Turning now to FIG. 1, an embodiment of an electronic processing system 10 may include an application processor 11, persistent storage media 12 communicatively coupled to the application processor 11, a graphics processor 13 communicatively coupled to the application processor 11, and a gesture tracker 14 communicatively coupled to the application processor 11. In some embodiments, the graphics processor 13 may be configured to generate an image of a virtual input device, the gesture tracker 14 may be configured to determine a position of a user's finger relative to the virtual input device, and the graphics processor 13 may be further configured to generate an image of a virtual finger based on the determined position of the user's finger relative to the virtual input device. Some embodiments of the system 10 may further include an input generator 15 communicatively coupled to the gesture tracker 14 to generate input data based on the determined position of the user's finger relative to the virtual input device. Some embodiments of the system may additionally, or alternatively, include a device tracker 16 communicatively coupled to the graphics processor 13 to identify a characteristic of a physical input device. For example, the graphics processor 13 may be configured to generate the image of the virtual input device based on the identified characteristic of the physical input device.

[0019] Embodiments of each of the above application processor 11, persistent storage media 12, graphics processor 13, gesture tracker 14, input generator 15, device tracker 16, and other system components may be implemented in hardware, software, or any suitable combination thereof. For example, hardware implementations may include configurable logic such as, for example, programmable logic arrays (PLAs), field-programmable gate arrays (FPGAs), complex programmable logic devices (CPLDs), or in fixed-functionality logic hardware using circuit technology such as, for example, ASIC, complementary metal oxide semiconductor (CMOS) or transistor-transistor logic (TTL) technology, or any combination thereof. Alternatively, or additionally, these components may be implemented in one or more modules as a set of logic instructions stored in a machine- or computer-readable storage medium such as random access memory (RAM), read only memory (ROM), programmable ROM (PROM), firmware, flash memory, etc., to be executed by a processor or computing device. For example, computer program code to carry out the operations of the components may be written in any combination of one or more operating system applicable/appropriate programming languages, including an object-oriented programming language such as PYTHON, PERL, JAVA, SMALLTALK, C++, C# or the like and conventional procedural programming languages, such as the "C" programming language or similar programming languages.

[0020] Sensor Examples

[0021] For example, a VR system may include a variety of sensors such as two dimensional (2D) cameras, three dimensional (3D) cameras, depth cameras, gyroscopes, accelerometers, inertial measurement units (IMUs), location services, microphones, proximity sensors, thermometers, biometric sensors, etc. The sensors may be distributed across multiple devices. Some embodiments may include an INTEL REALSENSE camera. The information from the sensors may include or be combined with input data from the user's other devices (e.g. smartphones, wearable devices, human interface devices (HIDs), etc.). For example, the

user's device(s) may include one or more 2D, 3D, and/or depth cameras. The user's device(s) may also include gyroscopes, accelerometers, IMUs, location services, thermometers, biometric sensors, etc. For example, the user may carry a smartphone (e.g. in the user's pocket) and/or may wear a wearable device (e.g. such as a smart watch, an activity monitor, and/or a fitness tracker). The user's device(s) may also include a microphone which may be utilized to detect if the user is speaking, on the phone, speaking to another nearby person, etc.

[0022] The sensors available to the VR system may also include some or all of the user's various other devices which are capable of capturing information related to the user's actions or activity (e.g. including an input/output (I/O) interface of the user devices which can capture keyboard/mouse/touch activity). The VR system may get information directly from the user's other devices (e.g. wired or wirelessly) or the VR system may be able to integrate information from the devices from a server or a service (e.g. information may be uploaded from a fitness tracker to a cloud service, which the VR system may download).

[0023] Gesture Tracker Examples

[0024] The VR system may be configured to perform facial recognition, gaze tracking, facial expression recognition, and/or gesture recognition including body-level gestures, arm/leg-level gestures, hand-level gestures, and/or finger-level gestures. For example, a machine vision system may be configured to classify an action of the user. In some embodiments, a suitably configured machine vision system may be able to determine if the user is present at a computer, typing at a keyboard, using a mouse, using a trackpad, using a touchscreen, using a HMD, using a VR system, sitting, standing, and/or otherwise taking some other action or activity. For example, devices such as the MICROSOFT KINECT and the MICROSOFT HOLOLENSA may provide various gesture tracking capabilities.

[0025] Device Tracker Examples

[0026] Some embodiments of a VR system, for example, may analyze and/or perform feature/object recognition on images captured by a camera. For example, machine vision and/or image processing may identify and/or recognize objects in a scene (e.g. a desk, a table, a keyboard, mouse, other HIDs, etc.). For example, the device tracker may get camera data related to a real object in a scene and may use that information to identify position and orientation of the real object. In some embodiments, the device tracker may be configured to construct a 3D model of an object in the captured scene. The device tracker may also identify when the real object in the scene moves from one position/orientation to another.

[0027] Focus/Gaze Tracker Examples

[0028] The VR system may get information from the sensors to determine focus information. At a high level, focus information may be based on 1) where the user is assumed to be looking, 2) where the user is determined to be looking, 3) where an application wants the user to look, and/or 4) where the user is predicted to be looking in the future. Some focus cues may be stronger in the focal region of where the user is looking. If the user is looking straight ahead they may see things in sharp focus. With scenes or objects towards the periphery, the user may notice motion but not details in sharp focus.

[0029] Some embodiments, for example, may include an eye tracker or get eye information from an eye tracker to

track the user's eyes. The user may go through a calibration process which may help the eye tracker provide more accurate information. When a user is wearing a VR headset, for example, a camera may capture an image of a pupil and the system may determine where the user is looking (e.g. a focus area, depth of field, and/or direction). The camera may capture pupil information and the system may infer where the user's focus area is based on that information.

[0030] Motion Tracker Examples

[0031] The VR system may also get information from the sensors, and/or the focus/gaze tracker, and other sources to provide motion information. At a high level, motion information may be based on 1) the user moving their head, 2) the user moving their eyes, 3) the user moving their body, 4) where an application wants the user to turn their head, eyes, and/or body, and/or 4) where the user is predicted to turn their head, eyes, and/or body in the future. Some motion information may be determined readily from the sensed information. For example, head position, velocity, acceleration, motion direction, etc. may be determined from an accelerometer. Eye motion information may be determined by tracking eye position information over time (e.g. if the eye tracker provides only eye position information).

[0032] Some motion information may be content-based. In a game or on-the-fly 3D content, for example, the application may know how quickly and where the objects are moving. The application may provide the information to the motion tracker (e.g. through an application programming interface (API) call). Future content-based object motion information for a next frame/scene may also be fed into the motion tracker for decision making. Some content-based motion information may be determined by image processing or machine vision processing the content.

[0033] Tracker Overlap Examples

[0034] Those skilled in the art will appreciate that aspects of various trackers described herein may overlap with other trackers and that portions of each tracker may be implemented or distributed throughout various portions of an electronic processing system. For example, the focus tracker may use motion information to provide a predicted future focus area and the motion tracker may use focus information to predict a future motion. Eye motion information may come directly from an eye tracker, may be determined/predicted by the focus tracker, and/or may be determined/predicted by the motion tracker. The examples herein should be considered as illustrative and not limiting in terms of specific implementations.

[0035] Turning now to FIG. 2, an embodiment of a graphics apparatus 20 may include an image generator 21, and a gesture tracker 22 communicatively coupled to the image generator 21. In some embodiments, the image generator 21 may be configured to generate an image of a virtual input device, the gesture tracker 22 may be configured to determine a position of a user's finger relative to the virtual input device, and the image generator 21 may also be configured to generate an image of a virtual finger based on the determined position of the user's finger relative to the virtual input device. The gesture tracker 22 may be further configured to determine a position of a user's hand relative to the virtual input device, and the image generator may be further configured to generate an image of a virtual hand based on the determined position of the user's hand relative to the virtual input device.

[0036] In some embodiments, the apparatus **20** may further include an input generator **23** communicatively coupled to the gesture tracker **22** to generate input data based on the determined position of the user's finger relative to the virtual input device. For example, the image generator **21** may generate a visual indication corresponding to the generated input data. For example, if the user presses a key in the real or virtual environments, the key may appear to move or may appear immersed as if it was pressed in the virtual environment. Additionally, or alternatively, the key may be highlighted with a different color, the key may be outlined, or some other visual indication may be generated. In addition, or alternatively, audio, haptic, or other feedback may be generated corresponding to the generated input data (e.g. a key click sound).

[0037] Some embodiments may additionally, or alternatively include a device tracker **24** communicatively coupled to the image generator **21** to identify a characteristic of a physical input device, where the image generator **21** may be further configured to generate the image of the virtual input device based on the identified characteristic of the physical input device. For example, the device tracker **24** may be configured to determine a position of the physical input device relative to the user, and the gesture tracker **22** may be configured to determine a position the user's hand relative to the position of the physical input device. The image generator may be configured to generate an image of a virtual representation of the physical input device based on the position of the physical input device relative to the user, and to generate an image of a virtual hand based on the determined position of the user's hand relative to the physical input device. In some embodiments, the image generator **21** may be configured to load a three-dimensional model of the virtual input device based on the identified characteristic of the physical input device. In any of the embodiments herein, non-limiting examples of the virtual input device may include any of a virtual keyboard, a virtual mouse, a virtual touchpad, a virtual stylus, and a virtual scroll wheel.

[0038] For example, embodiments or portions of the image generator **21** may be implemented in applications or driver software (e.g. through an API). Other embodiments or portions of the image generator **21** may be implemented in specialized code (e.g. shaders) to be executed on a graphics processor unit (GPU). Other embodiments or portions of the image generator **21** may be implemented in fixed function logic or specialized hardware (e.g. in the GPU). In some embodiments, the image generator **21** may be part of a VR application with access to a database/library of 3D models of input devices and a database/library of 3D models of fingers and/or hands. For example, the VR application may run on a general purpose processor which calls appropriate driver routines and/or graphics routines to generate images of the virtual input device(s), virtual finger(s), and/or virtual hands on the display(s). In some embodiments, the image generator **21** may be part of a VR runtime system loaded as driver routines. For example, VR applications may simply request input through an API call and the VR runtime routines may execute appropriate code and call appropriate graphics routines to generate the images of the virtual input device(s), virtual finger(s), and/or virtual hands on the display(s).

[0039] For example, embodiments or portions of the input generator **23** may be implemented in applications or driver software (e.g. through an API). Other embodiments or portions of the input generator **23** may be implemented in

specialized code (e.g. shaders) to be executed on a GPU. Other embodiments or portions of the input generator **23** may be implemented in fixed function logic or specialized hardware (e.g. in the GPU). For example, the input generator **23** may get information from a gesture tracker to indicate that the user has performed a key press action with a particular finger. The input generator **23** may determine the position of the corresponding virtual finger relative to the virtual keyboard to identify which key the user pressed (or virtually pressed). The input generator **23** may then send a signal through the IO interface corresponding to the identified key. The input generator **23** may also be configured to generate signals through the IO interface that correspond to gestures such as touch (e.g. for a touchpad or a touch screen), scrolling (e.g. for a scroll wheel on a mouse), button clicks (e.g. for mouse buttons). The input generator **23** may also be configured to generate signals through the IO interface based on information from the gesture tracker **22** and/or the device tracker **24** that correspond to movement of an input device (e.g. such as moving a physical mouse or virtually moving a virtual mouse).

[0040] Embodiments of each of the above image generator **21**, gesture tracker **22**, input generator **23**, device tracker **24**, and other components of the apparatus **20** may be implemented in hardware, software, or any combination thereof. For example, hardware implementations may include configurable logic such as, for example, PLAs, FPGAs, CPLDs, or in fixed-functionality logic hardware using circuit technology such as, for example, ASIC, CMOS, or TTL technology, or any combination thereof. Alternatively, or additionally, these components may be implemented in one or more modules as a set of logic instructions stored in a machine- or computer-readable storage medium such as RAM, ROM, PROM, firmware, flash memory, etc., to be executed by a processor or computing device. For example, computer program code to carry out the operations of the components may be written in any combination of one or more operating system applicable/appropriate programming languages, including an object-oriented programming language such as PYTHON, PERL, JAVA, SMALLTALK, C++, C# or the like and conventional procedural programming languages, such as the "C" programming language or similar programming languages.

[0041] Turning now to FIGS. **3** to **5**, an embodiment of a method **30** of displaying virtual objects may include generating an image of a virtual input device at block **31**, determining a position of a user's finger relative to the virtual input device at block **32**, and generating an image of a virtual finger based on the determined position of the user's finger relative to the virtual input device at block **33**. The method **30** may also include determining a position of a user's hand relative to the virtual input device at block **34**, and generating an image of a virtual hand based on the determined position of the user's hand relative to the virtual input device at block **35**. For example, some embodiments may include generating input data based on the determined position of the user's finger relative to the virtual input device at block **36**, and generating a visual indication corresponding to the generated input data at block **37**.

[0042] In some embodiments, the method **30** may further include identifying a characteristic of a physical input device at block **38**, and generating the image of the virtual input device based on the identified characteristic of the physical input device at block **39**. For example, the method **30** may

include determining a position of the physical input device relative to the user at block 40, determining a position the user's hand relative to the position of the physical input device at block 41, generating an image of a virtual representation of the physical input device based on the position of the physical input device relative to the user at block 42, and generating an image of a virtual hand based on the determined position of the user's hand relative to the physical input device at block 43. Some embodiments may further include loading a three-dimensional model of the virtual input device based on the identified characteristic of the physical input device at block 44. For example, the virtual input device may include one of a virtual keyboard, a virtual mouse, a virtual touchpad, a virtual stylus, and a virtual scroll wheel at block 45, among other HIDs.

[0043] Embodiments of the method 30 may be implemented in a system, apparatus, application processor, graphics processor unit (GPU), parallel processor unit (PPU), or a graphics processor pipeline apparatus such as, for example, those described herein. More particularly, hardware implementations of the method 30 may include configurable logic such as, for example, PLAs, FPGAs, CPLDs, or in fixed-functionality logic hardware using circuit technology such as, for example, ASIC, CMOS, or TTL technology, or any combination thereof. Alternatively, or additionally, the method 30 may be implemented in one or more modules as a set of logic instructions stored in a machine- or computer-readable storage medium such as RAM, ROM, PROM, firmware, flash memory, etc., to be executed by a processor or computing device. For example, computer program code to carry out the operations of the components may be written in any combination of one or more operating system applicable/appropriate programming languages, including an object-oriented programming language such as PYTHON, PERL, JAVA, SMALLTALK, C++, C# or the like and conventional procedural programming languages, such as the "C" programming language or similar programming languages. For example, the method 30 may be implemented on a computer readable medium as described in connection with Examples 19 to 25 below.

[0044] For example, embodiments or portions of the method 30 may be implemented in applications or driver software (e.g. through an API). Other embodiments or portions of the method 30 may be implemented in specialized code (e.g. shaders) to be executed on a GPU. Other embodiments or portions of the method 30 may be implemented in fixed function logic or specialized hardware (e.g. in the GPU).

[0045] While VR applications may be well suited for entertainment, conventional VR systems may not be as well suited for productivity usages such as web browsing, checking e-mail, word processing, spreadsheets, etc. Some embodiments may increase the potential for VR systems to support productivity usages. For example, an embodiment of VR system may replace a multi-monitor desktop setup in a work cubicle with a headset and some wireless input devices, thereby reducing wiring and desktop clutter. Some embodiments may provide a way to use common HIDs, such as keyboard and mouse, which may be important for user productivity. For productivity, a user may want to use their most comfortable HID (such as an ergonomic keyboard/mouse) for long time periods, the same way they might without a HMD. Also, a user who prefers to look at a keyboard while typing may have difficulty typing while

wearing the HMD (e.g. without being able to see the keyboard). Some embodiments may advantageously improve usability of a keyboard and/or mouse in a VR environment for productivity usages. For example, some embodiments may expand the use of input devices for VR applications.

[0046] For entertainment usages with limited input requirements, some applications or a VR runtime routine may draw a software keyboard, which a user may control through gaze together with a touchpad or controller to enter data. For cursor input, gaze may control a cursor movement and the touchpad/controller may act as a selection action. Gaze control for input may be a limited option for productivity usage. Advantageously, some embodiments may provide better input speed for more intensive data input applications. Some embodiments may also provide more ergonomic comfort for entering data over a longer time period.

[0047] Turning now to FIG. 6A, FIG. 6B, and FIG. 7, a VR system 60 may include a HMD 61 to be worn by a user 62. A look-down camera 63 may be positioned on the HMD 61 to capture image and/or depth image information (e.g. see FIG. 6A), which may be used for gesture tracking of the fingers and/or hands 64 of the user 62. If the user 62 addresses a keyboard 65 in the real environment by placing their hands 64 near the keyboard 65 (e.g. see FIG. 6B), a gesture tracker may recognize that gesture and the VR system 60 may generate an image of a virtual keyboard 65_v along with virtual hands 64_v in the virtual environment shown on a display 66 of the HMD 61 in response to the recognized gesture. The VR system may further recognize a physical mouse 67 in the real environment and generate an image of a virtual mouse 67_v in the virtual environment on the display 66. In some embodiments, the VR system 60 may overlay one or more productivity applications 68 (e.g. a file browser, a web browser, and email application, etc.) in front of the VR scene so the user can productively work on another application without removing the HMD 61. As illustrated in FIG. 7, a key press in the real environment may generate a visual indication of a key press 69 in the virtual environment (e.g. a movement of the pressed key in the 3D model, and/or a change in color or highlight of the pressed key).

[0048] Some embodiments may provide a VR runtime routine to render a 3D representation of the actual keyboard/mouse model that a user has setup (e.g. or which has been auto-detected). The virtual keyboard may be drawn in the virtual space close to where the user's finger would be if they were inside the virtual space. Virtual hands/fingers may also be rendered graphically, positioned appropriately on the keys/mouse the user wants to interact with. Some embodiments may texture the animated hands/fingers to match the user's skin tone/color based on corresponding information sensed by the HMD's various trackers.

[0049] In some embodiments, the VR system may identify a physical input device. For example, the VR runtime routine may present a menu to the user on a one-time basis (which can be overridden or changed), and the user may select the appropriate device (e.g. identified by model numbers of the devices). The user may also attach their physical devices to the VR system and the model or serial number may be auto-detected. For example, if the device is paired over Bluetooth to the VR runtime routine, the model of the device may be identified using the idVendor and idProduct

fields of the USB HID Device Descriptor. The VR system may load a 3D model of the identified product either from the cloud or from local storage based on the model or serial number. In some embodiments, the VR system may be able to determine a physical layout of the identified product and generate the 3D model based on the physical layout.

[0050] Some embodiments may provide the user two or more options to interact with a physical input device in the virtual environment. For example, the user may connect their physical input device to the VR system and use the device as they normally would for input. Additionally, or alternatively, the user may use a dummy physical device (e.g. not connected to the VR system but still providing haptic feedback through the feel of the device). Additionally, or alternatively, the user may just make the motions of using an input device (e.g. type on the desk or in the air as though a keyboard existed at their hand position). In any of these examples, the user would see a virtual input device in the virtual environment, and a gesture tracker (e.g. a wide angled, world-facing camera on the HMD) may constantly track the user's hand and finger movements. The VR system may update the rendered image to correspond to the user using the virtual input device. For example, if the user's finger moves from key "A" to key "L" in the real environment, animated fingers in the virtual space may be updated accordingly. If the user physically repositions the input device (e.g. to get more ergonomic comfort), the location and orientation of the 3D model of that device may be adjusted in the virtual world.

[0051] Some embodiments may also be useful for AR/MR applications. For example, if the user does not have physical access to their input device(s), they may make a gesture that the AR/MR system may recognize as addressing a keyboard/mouse/stylus/etc. (e.g. holding both hands out in front of their body with their fingers curved). The AR/MR system may then generate an image of a virtual input device proximate to the user's hands/fingers for the user to interact with in the AR/MR environment. Advantageously, the virtual input device may correspond to the user's usual physical input device and the user's muscle memory for that device may provide more productive use of the virtual input device. In some embodiments, the VR/AR/MR system may scale a size of the virtual keyboard to a detected and/or measured hand size based on the user's resting and/or natural hand position (e.g. when no physical device is present).

[0052] In some embodiments, the user may advantageously have the flexibility to switch seamlessly between their real devices and a pretend device if they don't have access to their real ones. In some embodiments, the devices may be drawn in the virtual space to closely match a user's physical setup (e.g. matching a model of a real device, positioned close to the user's view, texturing the graphical hands with the user's skin color etc.). The user may advantageously find the setup in the virtual environment very close to the productivity setup they are used to in their real space.

[0053] In some embodiments, the user may call up a virtual productivity environment on the HMD display. When the productivity environment is setup, models for virtual input devices may be identified. For example, the user may connect or pair a physical input device (e.g. keyboard/mouse/stylus/touchpad/etc.) to the VR system over a medium such as USB or Bluetooth. The connection may cause exchange of HID descriptors, which may allow the

VR system to identify the make and model of the device. The user may additionally, or alternatively, select a device from an appropriate interface. For example, the VR system presents an interface that allows the choice of make and model of their preferred input device. In both (a) and (b), the VR system remembers the choice until overridden with a different preference. In addition, or alternatively, the VR system may read a bar code, a quick-response (QR) code, radio-frequency identification (RFID) tag, or other symbol or tag on the physical input device to recognize the device. In addition, or alternatively, the VR system may use machine vision to recognize the device. In some embodiments, the VR system may be further configured to create a 3D model of the physical input device (e.g. using machine vision).

[0054] The VR system may then load a 3D graphical model of the make and model identified (or the created 3D model), either from cloud or local storage (e.g. from a database of 3D models). A world-facing, wide-angled camera on the periphery of the HMD may start tracking the user's hand, fingers, and any real input device that the user may be interacting with. The VR system may take inputs from the tracker and overlay a graphical version of the user's hands and input devices, in the orientation and position close to the real scenario (e.g. as detected by the tracking camera). The tracker input may be constantly monitored to update the graphical model with user's current hand/finger positions. Advantageously, the user may be able to type more comfortably without looking outside the HMD because the user may be able to see their virtual fingers as they type.

[0055] In some embodiments, the VR system may be part of a virtual operating system (OS) specifically designed to support VR applications together with productivity applications. When the user brings up the virtual productivity environment, they may pause a VR application to attend to other applications (e.g. e-mail). The paused VR scene may become a background image for the other application windows. For example, the user making a keyboard gesture may automatically pause the current VR application. The user may also allow the VR application to keep running and the virtual OS and/or other application windows may appear to float in front of the VR graphics.

[0056] System Overview

[0057] FIG. 8 is a block diagram of a processing system 100, according to an embodiment. In various embodiments the system 100 includes one or more processors 102 and one or more graphics processors 108, and may be a single processor desktop system, a multiprocessor workstation system, or a server system having a large number of processors 102 or processor cores 107. In an embodiment, the system 100 is a processing platform incorporated within a system-on-a-chip (SoC) integrated circuit for use in mobile, handheld, or embedded devices.

[0058] An embodiment of system 100 can include, or be incorporated within a server-based gaming platform, a game console, including a game and media console, a mobile gaming console, a handheld game console, or an online game console. In some embodiments system 100 is a mobile phone, smart phone, tablet computing device or mobile Internet device. Data processing system 100 can also include, couple with, or be integrated within a wearable device, such as a smart watch wearable device, smart eyewear device, augmented reality device, or virtual reality device. In some embodiments, data processing system 100 is a television or set top box device having one or more

processors **102** and a graphical interface generated by one or more graphics processors **108**.

[0059] In some embodiments, the one or more processors **102** each include one or more processor cores **107** to process instructions which, when executed, perform operations for system and user software. In some embodiments, each of the one or more processor cores **107** is configured to process a specific instruction set **109**. In some embodiments, instruction set **109** may facilitate Complex Instruction Set Computing (CISC), Reduced Instruction Set Computing (RISC), or computing via a Very Long Instruction Word (VLIW). Multiple processor cores **107** may each process a different instruction set **109**, which may include instructions to facilitate the emulation of other instruction sets. Processor core **107** may also include other processing devices, such as a Digital Signal Processor (DSP).

[0060] In some embodiments, the processor **102** includes cache memory **104**. Depending on the architecture, the processor **102** can have a single internal cache or multiple levels of internal cache. In some embodiments, the cache memory is shared among various components of the processor **102**. In some embodiments, the processor **102** also uses an external cache (e.g., a Level-3 (L3) cache or LLC) (not shown), which may be shared among processor cores **107** using known cache coherency techniques. A register file **106** is additionally included in processor **102** which may include different types of registers for storing different types of data (e.g., integer registers, floating point registers, status registers, and an instruction pointer register). Some registers may be general-purpose registers, while other registers may be specific to the design of the processor **102**.

[0061] In some embodiments, processor **102** is coupled to a processor bus **110** to transmit communication signals such as address, data, or control signals between processor **102** and other components in system **100**. In one embodiment the system **100** uses an exemplary 'hub' system architecture, including a memory controller hub **116** and an Input Output (I/O) controller hub **130**. A memory controller hub **116** facilitates communication between a memory device and other components of system **100**, while an I/O Controller Hub (ICH) **130** provides connections to I/O devices via a local I/O bus. In one embodiment, the logic of the memory controller hub **116** is integrated within the processor.

[0062] Memory device **120** can be a dynamic random access memory (DRAM) device, a static random access memory (SRAM) device, flash memory device, phase-change memory device, or some other memory device having suitable performance to serve as process memory. In one embodiment the memory device **120** can operate as system memory for the system **100**, to store data **122** and instructions **121** for use when the one or more processors **102** executes an application or process. Memory controller hub **116** also couples with an optional external graphics processor **112**, which may communicate with the one or more graphics processors **108** in processors **102** to perform graphics and media operations.

[0063] In some embodiments, ICH **130** enables peripherals to connect to memory device **120** and processor **102** via a high-speed I/O bus. The I/O peripherals include, but are not limited to, an audio controller **146**, a firmware interface **128**, a wireless transceiver **126** (e.g., Wi-Fi, Bluetooth), a data storage device **124** (e.g., hard disk drive, flash memory, etc.), and a legacy I/O controller **140** for coupling legacy (e.g., Personal System 2 (PS/2)) devices to the system. One

or more Universal Serial Bus (USB) controllers **142** connect input devices, such as keyboard and mouse **144** combinations. A network controller **134** may also couple to ICH **130**. In some embodiments, a high-performance network controller (not shown) couples to processor bus **110**. It will be appreciated that the system **100** shown is exemplary and not limiting, as other types of data processing systems that are differently configured may also be used. For example, the I/O controller hub **130** may be integrated within the one or more processor **102**, or the memory controller hub **116** and I/O controller hub **130** may be integrated into a discreet external graphics processor, such as the external graphics processor **112**.

[0064] FIG. **9** is a block diagram of an embodiment of a processor **200** having one or more processor cores **202A-202N**, an integrated memory controller **214**, and an integrated graphics processor **208**. Those elements of FIG. **9** having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such. Processor **200** can include additional cores up to and including additional core **202N** represented by the dashed lined boxes. Each of processor cores **202A-202N** includes one or more internal cache units **204A-204N**. In some embodiments each processor core also has access to one or more shared cache units **206**.

[0065] The internal cache units **204A-204N** and shared cache units **206** represent a cache memory hierarchy within the processor **200**. The cache memory hierarchy may include at least one level of instruction and data cache within each processor core and one or more levels of shared mid-level cache, such as a Level 2 (L2), Level 3 (L3), Level 4 (L4), or other levels of cache, where the highest level of cache before external memory is classified as the LLC. In some embodiments, cache coherency logic maintains coherency between the various cache units **206** and **204A-204N**.

[0066] In some embodiments, processor **200** may also include a set of one or more bus controller units **216** and a system agent core **210**. The one or more bus controller units **216** manage a set of peripheral buses, such as one or more Peripheral Component Interconnect buses (e.g., PCI, PCI Express). System agent core **210** provides management functionality for the various processor components. In some embodiments, system agent core **210** includes one or more integrated memory controllers **214** to manage access to various external memory devices (not shown).

[0067] In some embodiments, one or more of the processor cores **202A-202N** include support for simultaneous multi-threading. In such embodiment, the system agent core **210** includes components for coordinating and operating cores **202A-202N** during multi-threaded processing. System agent core **210** may additionally include a power control unit (PCU), which includes logic and components to regulate the power state of processor cores **202A-202N** and graphics processor **208**.

[0068] In some embodiments, processor **200** additionally includes graphics processor **208** to execute graphics processing operations. In some embodiments, the graphics processor **208** couples with the set of shared cache units **206**, and the system agent core **210**, including the one or more integrated memory controllers **214**. In some embodiments, a display controller **211** is coupled with the graphics processor **208** to drive graphics processor output to one or more coupled displays. In some embodiments, display controller

211 may be a separate module coupled with the graphics processor via at least one interconnect, or may be integrated within the graphics processor **208** or system agent core **210**.

[0069] In some embodiments, a ring based interconnect unit **212** is used to couple the internal components of the processor **200**. However, an alternative interconnect unit may be used, such as a point-to-point interconnect, a switched interconnect, or other techniques, including techniques well known in the art. In some embodiments, graphics processor **208** couples with the ring interconnect **212** via an I/O link **213**.

[0070] The exemplary I/O link **213** represents at least one of multiple varieties of I/O interconnects, including an on package I/O interconnect which facilitates communication between various processor components and a high-performance embedded memory module **218**, such as an eDRAM module. In some embodiments, each of the processor cores **202-202N** and graphics processor **208** use embedded memory modules **218** as a shared Last Level Cache.

[0071] In some embodiments, processor cores **202A-202N** are homogenous cores executing the same instruction set architecture. In another embodiment, processor cores **202A-202N** are heterogeneous in terms of instruction set architecture (ISA), where one or more of processor cores **202A-N** execute a first instruction set, while at least one of the other cores executes a subset of the first instruction set or a different instruction set. In one embodiment processor cores **202A-202N** are heterogeneous in terms of microarchitecture, where one or more cores having a relatively higher power consumption couple with one or more power cores having a lower power consumption. Additionally, processor **200** can be implemented on one or more chips or as an SoC integrated circuit having the illustrated components, in addition to other components.

[0072] FIG. **10** is a block diagram of a graphics processor **300**, which may be a discrete graphics processing unit, or may be a graphics processor integrated with a plurality of processing cores. In some embodiments, the graphics processor communicates via a memory mapped I/O interface to registers on the graphics processor and with commands placed into the processor memory. In some embodiments, graphics processor **300** includes a memory interface **314** to access memory. Memory interface **314** can be an interface to local memory, one or more internal caches, one or more shared external caches, and/or to system memory.

[0073] In some embodiments, graphics processor **300** also includes a display controller **302** to drive display output data to a display device **320**. Display controller **302** includes hardware for one or more overlay planes for the display and composition of multiple layers of video or user interface elements. In some embodiments, graphics processor **300** includes a video codec engine **306** to encode, decode, or transcode media to, from, or between one or more media encoding formats, including, but not limited to Moving Picture Experts Group (MPEG) formats such as MPEG-2, Advanced Video Coding (AVC) formats such as H.264/MPEG-4 AVC, as well as the Society of Motion Picture & Television Engineers (SMPTE) 421M/VC-1, and Joint Photographic Experts Group (JPEG) formats such as JPEG, and Motion JPEG (MJPEG) formats.

[0074] In some embodiments, graphics processor **300** includes a block image transfer (BLIT) engine **304** to perform two-dimensional (2D) rasterizer operations including, for example, bit-boundary block transfers. However, in

one embodiment, 2D graphics operations are performed using one or more components of graphics—processing engine (GPE) **310**. In some embodiments, graphics-processing engine **310** is a compute engine for performing graphics operations, including three-dimensional (3D) graphics operations and media operations.

[0075] In some embodiments, GPE **310** includes a 3D pipeline **312** for performing 3D operations, such as rendering three-dimensional images and scenes using processing functions that act upon 3D primitive shapes (e.g., rectangle, triangle, etc.). The 3D pipeline **312** includes programmable and fixed function elements that perform various tasks within the element and/or spawn execution threads to a 3D/Media sub-system **315**. While 3D pipeline **312** can be used to perform media operations, an embodiment of GPE **310** also includes a media pipeline **316** that is specifically used to perform media operations, such as video post-processing and image enhancement.

[0076] In some embodiments, media pipeline **316** includes fixed function or programmable logic units to perform one or more specialized media operations, such as video decode acceleration, video de-interlacing, and video encode acceleration in place of, or on behalf of video codec engine **306**. In some embodiments, media pipeline **316** additionally includes a thread spawning unit to spawn threads for execution on 3D/Media sub-system **315**. The spawned threads perform computations for the media operations on one or more graphics execution units included in 3D/Media sub-system **315**.

[0077] In some embodiments, 3D/Media subsystem **315** includes logic for executing threads spawned by 3D pipeline **312** and media pipeline **316**. In one embodiment, the pipelines send thread execution requests to 3D/Media subsystem **315**, which includes thread dispatch logic for arbitrating and dispatching the various requests to available thread execution resources. The execution resources include an array of graphics execution units to process the 3D and media threads. In some embodiments, 3D/Media subsystem **315** includes one or more internal caches for thread instructions and data. In some embodiments, the subsystem also includes shared memory, including registers and addressable memory, to share data between threads and to store output data.

[0078] 3D/Media Processing

[0079] FIG. **11** is a block diagram of a graphics processing engine **410** of a graphics processor in accordance with some embodiments. In one embodiment, the GPE **410** is a version of the GPE **310** shown in FIG. **10**. Elements of FIG. **11** having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such.

[0080] In some embodiments, GPE **410** couples with a command streamer **403**, which provides a command stream to the GPE 3D and media pipelines **412**, **416**. In some embodiments, command streamer **403** is coupled to memory, which can be system memory, or one or more of internal cache memory and shared cache memory. In some embodiments, command streamer **403** receives commands from the memory and sends the commands to 3D pipeline **412** and/or media pipeline **416**. The commands are directives fetched from a ring buffer, which stores commands for the 3D and media pipelines **412**, **416**. In one embodiment, the ring buffer can additionally include batch command

buffers storing batches of multiple commands. The 3D and media pipelines **412**, **416** process the commands by performing operations via logic within the respective pipelines or by dispatching one or more execution threads to an execution unit array **414**. In some embodiments, execution unit array **414** is scalable, such that the array includes a variable number of execution units based on the target power and performance level of GPE **410**.

[0081] In some embodiments, a sampling engine **430** couples with memory (e.g., cache memory or system memory) and execution unit array **414**. In some embodiments, sampling engine **430** provides a memory access mechanism for execution unit array **414** that allows execution array **414** to read graphics and media data from memory. In some embodiments, sampling engine **430** includes logic to perform specialized image sampling operations for media.

[0082] In some embodiments, the specialized media sampling logic in sampling engine **430** includes a de-noise/de-interlace module **432**, a motion estimation module **434**, and an image scaling and filtering module **436**. In some embodiments, de-noise/de-interlace module **432** includes logic to perform one or more of a de-noise or a de-interlace algorithm on decoded video data. The de-interlace logic combines alternating fields of interlaced video content into a single frame of video. The de-noise logic reduces or removes data noise from video and image data. In some embodiments, the de-noise logic and de-interlace logic are motion adaptive and use spatial or temporal filtering based on the amount of motion detected in the video data. In some embodiments, the de-noise/de-interlace module **432** includes dedicated motion detection logic (e.g., within the motion estimation engine **434**).

[0083] In some embodiments, motion estimation engine **434** provides hardware acceleration for video operations by performing video acceleration functions such as motion vector estimation and prediction on video data. The motion estimation engine determines motion vectors that describe the transformation of image data between successive video frames. In some embodiments, a graphics processor media codec uses video motion estimation engine **434** to perform operations on video at the macro-block level that may otherwise be too computationally intensive to perform with a general-purpose processor. In some embodiments, motion estimation engine **434** is generally available to graphics processor components to assist with video decode and processing functions that are sensitive or adaptive to the direction or magnitude of the motion within video data.

[0084] In some embodiments, image scaling and filtering module **436** performs image-processing operations to enhance the visual quality of generated images and video. In some embodiments, scaling and filtering module **436** processes image and video data during the sampling operation before providing the data to execution unit array **414**.

[0085] In some embodiments, the GPE **410** includes a data port **444**, which provides an additional mechanism for graphics subsystems to access memory. In some embodiments, data port **444** facilitates memory access for operations including render target writes, constant buffer reads, scratch memory space reads/writes, and media surface accesses. In some embodiments, data port **444** includes cache memory space to cache accesses to memory. The cache memory can be a single data cache or separated into multiple caches for the multiple subsystems that access

memory via the data port (e.g., a render buffer cache, a constant buffer cache, etc.). In some embodiments, threads executing on an execution unit in execution unit array **414** communicate with the data port by exchanging messages via a data distribution interconnect that couples each of the sub-systems of GPE **410**.

[0086] Execution Units

[0087] FIG. **12** is a block diagram of another embodiment of a graphics processor **500**. Elements of FIG. **12** having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such.

[0088] In some embodiments, graphics processor **500** includes a ring interconnect **502**, a pipeline front-end **504**, a media engine **537**, and graphics cores **580A-580N**. In some embodiments, ring interconnect **502** couples the graphics processor to other processing units, including other graphics processors or one or more general-purpose processor cores. In some embodiments, the graphics processor is one of many processors integrated within a multi-core processing system.

[0089] In some embodiments, graphics processor **500** receives batches of commands via ring interconnect **502**. The incoming commands are interpreted by a command streamer **503** in the pipeline front-end **504**. In some embodiments, graphics processor **500** includes scalable execution logic to perform 3D geometry processing and media processing via the graphics core(s) **580A-580N**. For 3D geometry processing commands, command streamer **503** supplies commands to geometry pipeline **536**. For at least some media processing commands, command streamer **503** supplies the commands to a video front end **534**, which couples with a media engine **537**. In some embodiments, media engine **537** includes a Video Quality Engine (VQE) **530** for video and image post-processing and a multi-format encode/decode (MFX) **533** engine to provide hardware-accelerated media data encode and decode. In some embodiments, geometry pipeline **536** and media engine **537** each generate execution threads for the thread execution resources provided by at least one graphics core **580A**.

[0090] In some embodiments, graphics processor **500** includes scalable thread execution resources featuring modular cores **580A-580N** (sometimes referred to as core slices), each having multiple sub-cores **550A-550N**, **560A-560N** (sometimes referred to as core sub-slices). In some embodiments, graphics processor **500** can have any number of graphics cores **580A** through **580N**. In some embodiments, graphics processor **500** includes a graphics core **580A** having at least a first sub-core **550A** and a second core sub-core **560A**. In other embodiments, the graphics processor is a low power processor with a single sub-core (e.g., **550A**). In some embodiments, graphics processor **500** includes multiple graphics cores **580A-580N**, each including a set of first sub-cores **550A-550N** and a set of second sub-cores **560A-560N**. Each sub-core in the set of first sub-cores **550A-550N** includes at least a first set of execution units **552A-552N** and media/texture samplers **554A-554N**. Each sub-core in the set of second sub-cores **560A-560N** includes at least a second set of execution units **562A-562N** and samplers **564A-564N**. In some embodiments, each sub-core **550A-550N**, **560A-560N** shares a set of shared resources **570A-570N**. In some embodiments, the shared resources include shared cache memory and pixel

operation logic. Other shared resources may also be included in the various embodiments of the graphics processor.

[0091] FIG. 13 illustrates thread execution logic 600 including an array of processing elements employed in some embodiments of a GPE. Elements of FIG. 13 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.

[0092] In some embodiments, thread execution logic 600 includes a pixel shader 602, a thread dispatcher 604, instruction cache 606, a scalable execution unit array including a plurality of execution units 608A-608N, a sampler 610, a data cache 612, and a data port 614. In one embodiment the included components are interconnected via an interconnect fabric that links to each of the components. In some embodiments, thread execution logic 600 includes one or more connections to memory, such as system memory or cache memory, through one or more of instruction cache 606, data port 614, sampler 610, and execution unit array 608A-608N. In some embodiments, each execution unit (e.g. 608A) is an individual vector processor capable of executing multiple simultaneous threads and processing multiple data elements in parallel for each thread. In some embodiments, execution unit array 608A-608N includes any number individual execution units.

[0093] In some embodiments, execution unit array 608A-608N is primarily used to execute “shader” programs. In some embodiments, the execution units in array 608A-608N execute an instruction set that includes native support for many standard 3D graphics shader instructions, such that shader programs from graphics libraries (e.g., Direct 3D and OpenGL) are executed with a minimal translation. The execution units support vertex and geometry processing (e.g., vertex programs, geometry programs, vertex shaders), pixel processing (e.g., pixel shaders, fragment shaders) and general-purpose processing (e.g., compute and media shaders).

[0094] Each execution unit in execution unit array 608A-608N operates on arrays of data elements. The number of data elements is the “execution size,” or the number of channels for the instruction. An execution channel is a logical unit of execution for data element access, masking, and flow control within instructions. The number of channels may be independent of the number of physical Arithmetic Logic Units (ALUs) or Floating Point Units (FPUs) for a particular graphics processor. In some embodiments, execution units 608A-608N support integer and floating-point data types.

[0095] The execution unit instruction set includes single instruction multiple data (SIMD) instructions. The various data elements can be stored as a packed data type in a register and the execution unit will process the various elements based on the data size of the elements. For example, when operating on a 256-bit wide vector, the 256 bits of the vector are stored in a register and the execution unit operates on the vector as four separate 64-bit packed data elements (Quad-Word (QW) size data elements), eight separate 32-bit packed data elements (Double Word (DW) size data elements), sixteen separate 16-bit packed data elements (Word (W) size data elements), or thirty-two

separate 8-bit data elements (byte (B) size data elements). However, different vector widths and register sizes are possible.

[0096] One or more internal instruction caches (e.g., 606) are included in the thread execution logic 600 to cache thread instructions for the execution units. In some embodiments, one or more data caches (e.g., 612) are included to cache thread data during thread execution. In some embodiments, sampler 610 is included to provide texture sampling for 3D operations and media sampling for media operations. In some embodiments, sampler 610 includes specialized texture or media sampling functionality to process texture or media data during the sampling process before providing the sampled data to an execution unit.

[0097] During execution, the graphics and media pipelines send thread initiation requests to thread execution logic 600 via thread spawning and dispatch logic. In some embodiments, thread execution logic 600 includes a local thread dispatcher 604 that arbitrates thread initiation requests from the graphics and media pipelines and instantiates the requested threads on one or more execution units 608A-608N. For example, the geometry pipeline (e.g., 536 of FIG. 12) dispatches vertex processing, tessellation, or geometry processing threads to thread execution logic 600 (FIG. 13). In some embodiments, thread dispatcher 604 can also process runtime thread spawning requests from the executing shader programs.

[0098] Once a group of geometric objects has been processed and rasterized into pixel data, pixel shader 602 is invoked to further compute output information and cause results to be written to output surfaces (e.g., color buffers, depth buffers, stencil buffers, etc.). In some embodiments, pixel shader 602 calculates the values of the various vertex attributes that are to be interpolated across the rasterized object. In some embodiments, pixel shader 602 then executes an application programming interface (API)-supplied pixel shader program. To execute the pixel shader program, pixel shader 602 dispatches threads to an execution unit (e.g., 608A) via thread dispatcher 604. In some embodiments, pixel shader 602 uses texture sampling logic in sampler 610 to access texture data in texture maps stored in memory. Arithmetic operations on the texture data and the input geometry data compute pixel color data for each geometric fragment, or discards one or more pixels from further processing.

[0099] In some embodiments, the data port 614 provides a memory access mechanism for the thread execution logic 600 output processed data to memory for processing on a graphics processor output pipeline. In some embodiments, the data port 614 includes or couples to one or more cache memories (e.g., data cache 612) to cache data for memory access via the data port.

[0100] FIG. 14 is a block diagram illustrating a graphics processor instruction formats 700 according to some embodiments. In one or more embodiment, the graphics processor execution units support an instruction set having instructions in multiple formats. The solid lined boxes illustrate the components that are generally included in an execution unit instruction, while the dashed lines include components that are optional or that are only included in a sub-set of the instructions. In some embodiments, instruction format 700 described and illustrated are macro-instructions, in that they are instructions supplied to the execution

unit, as opposed to micro-operations resulting from instruction decode once the instruction is processed.

[0101] In some embodiments, the graphics processor execution units natively support instructions in a 128-bit format **710**. A 64-bit compacted instruction format **730** is available for some instructions based on the selected instruction, instruction options, and number of operands. The native 128-bit format **710** provides access to all instruction options, while some options and operations are restricted in the 64-bit format **730**. The native instructions available in the 64-bit format **730** vary by embodiment. In some embodiments, the instruction is compacted in part using a set of index values in an index field **713**. The execution unit hardware references a set of compaction tables based on the index values and uses the compaction table outputs to reconstruct a native instruction in the 128-bit format **710**.

[0102] For each format, instruction opcode **712** defines the operation that the execution unit is to perform. The execution units execute each instruction in parallel across the multiple data elements of each operand. For example, in response to an add instruction the execution unit performs a simultaneous add operation across each color channel representing a texture element or picture element. By default, the execution unit performs each instruction across all data channels of the operands. In some embodiments, instruction control field **714** enables control over certain execution options, such as channels selection (e.g., predication) and data channel order (e.g., swizzle). For 128-bit instructions **710** an exec-size field **716** limits the number of data channels that will be executed in parallel. In some embodiments, exec-size field **716** is not available for use in the 64-bit compact instruction format **730**.

[0103] Some execution unit instructions have up to three operands including two source operands, src0 **722**, src1 **722**, and one destination **718**. In some embodiments, the execution units support dual destination instructions, where one of the destinations is implied. Data manipulation instructions can have a third source operand (e.g., SRC2 **724**), where the instruction opcode **712** determines the number of source operands. An instruction's last source operand can be an immediate (e.g., hard-coded) value passed with the instruction.

[0104] In some embodiments, the 128-bit instruction format **710** includes an access/address mode information **726** specifying, for example, whether direct register addressing mode or indirect register addressing mode is used. When direct register addressing mode is used, the register address of one or more operands is directly provided by bits in the instruction **710**.

[0105] In some embodiments, the 128-bit instruction format **710** includes an access/address mode field **726**, which specifies an address mode and/or an access mode for the instruction. In one embodiment the access mode to define a data access alignment for the instruction. Some embodiments support access modes including a 16-byte aligned access mode and a 1-byte aligned access mode, where the byte alignment of the access mode determines the access alignment of the instruction operands. For example, when in a first mode, the instruction **710** may use byte-aligned addressing for source and destination operands and when in a second mode, the instruction **710** may use 16-byte-aligned addressing for all source and destination operands.

[0106] In one embodiment, the address mode portion of the access/address mode field **726** determines whether the

instruction is to use direct or indirect addressing. When direct register addressing mode is used bits in the instruction **710** directly provide the register address of one or more operands. When indirect register addressing mode is used, the register address of one or more operands may be computed based on an address register value and an address immediate field in the instruction.

[0107] In some embodiments instructions are grouped based on opcode **712** bit-fields to simplify Opcode decode **740**. For an 8-bit opcode, bits **4**, **5**, and **6** allow the execution unit to determine the type of opcode. The precise opcode grouping shown is merely an example. In some embodiments, a move and logic opcode group **742** includes data movement and logic instructions (e.g., move (mov), compare (cmp)). In some embodiments, move and logic group **742** shares the five most significant bits (MSB), where move (mov) instructions are in the form of 0000xxxxb and logic instructions are in the form of 0001xxxxb. A flow control instruction group **744** (e.g., call, jump (jmp)) includes instructions in the form of 0010xxxxb (e.g., 0x20). A miscellaneous instruction group **746** includes a mix of instructions, including synchronization instructions (e.g., wait, send) in the form of 0011xxxxb (e.g., 0x30). A parallel math instruction group **748** includes component-wise arithmetic instructions (e.g., add, multiply (mul)) in the form of 0100xxxxb (e.g., 0x40). The parallel math group **748** performs the arithmetic operations in parallel across data channels. The vector math group **750** includes arithmetic instructions (e.g., dp4) in the form of 0101xxxxb (e.g., 0x50). The vector math group performs arithmetic such as dot product calculations on vector operands.

[0108] Graphics Pipeline

[0109] FIG. **15** is a block diagram of another embodiment of a graphics processor **800**. Elements of FIG. **15** having the same reference numbers (or names) as the elements of any other figure herein can operate or function in any manner similar to that described elsewhere herein, but are not limited to such.

[0110] In some embodiments, graphics processor **800** includes a graphics pipeline **820**, a media pipeline **830**, a display engine **840**, thread execution logic **850**, and a render output pipeline **870**. In some embodiments, graphics processor **800** is a graphics processor within a multi-core processing system that includes one or more general purpose processing cores. The graphics processor is controlled by register writes to one or more control registers (not shown) or via commands issued to graphics processor **800** via a ring interconnect **802**. In some embodiments, ring interconnect **802** couples graphics processor **800** to other processing components, such as other graphics processors or general-purpose processors. Commands from ring interconnect **802** are interpreted by a command streamer **803**, which supplies instructions to individual components of graphics pipeline **820** or media pipeline **830**.

[0111] In some embodiments, command streamer **803** directs the operation of a vertex fetcher **805** that reads vertex data from memory and executes vertex-processing commands provided by command streamer **803**. In some embodiments, vertex fetcher **805** provides vertex data to a vertex shader **807**, which performs coordinate space transformation and lighting operations to each vertex. In some embodiments, vertex fetcher **805** and vertex shader **807**

execute vertex-processing instructions by dispatching execution threads to execution units **852A**, **852B** via a thread dispatcher **831**.

[0112] In some embodiments, execution units **852A**, **852B** are an array of vector processors having an instruction set for performing graphics and media operations. In some embodiments, execution units **852A**, **852B** have an attached L1 cache **851** that is specific for each array or shared between the arrays. The cache can be configured as a data cache, an instruction cache, or a single cache that is partitioned to contain data and instructions in different partitions.

[0113] In some embodiments, graphics pipeline **820** includes tessellation components to perform hardware-accelerated tessellation of 3D objects. In some embodiments, a programmable hull shader **811** configures the tessellation operations. A programmable domain shader **817** provides back-end evaluation of tessellation output. A tessellator **813** operates at the direction of hull shader **811** and contains special purpose logic to generate a set of detailed geometric objects based on a coarse geometric model that is provided as input to graphics pipeline **820**. In some embodiments, if tessellation is not used, tessellation components **811**, **813**, **817** can be bypassed.

[0114] In some embodiments, complete geometric objects can be processed by a geometry shader **819** via one or more threads dispatched to execution units **852A**, **852B**, or can proceed directly to the clipper **829**. In some embodiments, the geometry shader operates on entire geometric objects, rather than vertices or patches of vertices as in previous stages of the graphics pipeline. If the tessellation is disabled the geometry shader **819** receives input from the vertex shader **807**. In some embodiments, geometry shader **819** is programmable by a geometry shader program to perform geometry tessellation if the tessellation units are disabled.

[0115] Before rasterization, a clipper **829** processes vertex data. The clipper **829** may be a fixed function clipper or a programmable clipper having clipping and geometry shader functions. In some embodiments, a rasterizer/depth **873** in the render output pipeline **870** dispatches pixel shaders to convert the geometric objects into their per pixel representations. In some embodiments, pixel shader logic is included in thread execution logic **850**. In some embodiments, an application can bypass the rasterizer **873** and access un-rasterized vertex data via a stream out unit **823**.

[0116] The graphics processor **800** has an interconnect bus, interconnect fabric, or some other interconnect mechanism that allows data and message passing amongst the major components of the processor. In some embodiments, execution units **852A**, **852B** and associated cache(s) **851**, texture and media sampler **854**, and texture/sampler cache **858** interconnect via a data port **856** to perform memory access and communicate with render output pipeline components of the processor. In some embodiments, sampler **854**, caches **851**, **858** and execution units **852A**, **852B** each have separate memory access paths.

[0117] In some embodiments, render output pipeline **870** contains a rasterizer and depth test component **873** that converts vertex-based objects into an associated pixel-based representation. In some embodiments, the rasterizer logic includes a windower/masker unit to perform fixed function triangle and line rasterization. An associated render cache **878** and depth cache **879** are also available in some embodiments. A pixel operations component **877** performs pixel-based operations on the data, though in some instances, pixel

operations associated with 2D operations (e.g. bit block image transfers with blending) are performed by the 2D engine **841**, or substituted at display time by the display controller **843** using overlay display planes. In some embodiments, a shared L3 cache **875** is available to all graphics components, allowing the sharing of data without the use of main system memory.

[0118] In some embodiments, graphics processor media pipeline **830** includes a media engine **837** and a video front end **834**. In some embodiments, video front end **834** receives pipeline commands from the command streamer **803**. In some embodiments, media pipeline **830** includes a separate command streamer. In some embodiments, video front-end **834** processes media commands before sending the command to the media engine **837**. In some embodiments, media engine **337** includes thread spawning functionality to spawn threads for dispatch to thread execution logic **850** via thread dispatcher **831**.

[0119] In some embodiments, graphics processor **800** includes a display engine **840**. In some embodiments, display engine **840** is external to processor **800** and couples with the graphics processor via the ring interconnect **802**, or some other interconnect bus or fabric. In some embodiments, display engine **840** includes a 2D engine **841** and a display controller **843**. In some embodiments, display engine **840** contains special purpose logic capable of operating independently of the 3D pipeline. In some embodiments, display controller **843** couples with a display device (not shown), which may be a system integrated display device, as in a laptop computer, or an external display device attached via a display device connector.

[0120] In some embodiments, graphics pipeline **820** and media pipeline **830** are configurable to perform operations based on multiple graphics and media programming interfaces and are not specific to any one application programming interface (API). In some embodiments, driver software for the graphics processor translates API calls that are specific to a particular graphics or media library into commands that can be processed by the graphics processor. In some embodiments, support is provided for the Open Graphics Library (OpenGL) and Open Computing Language (OpenCL) from the Khronos Group, the Direct3D library from the Microsoft Corporation, or support may be provided to both OpenGL and D3D. Support may also be provided for the Open Source Computer Vision Library (OpenCV). A future API with a compatible 3D pipeline would also be supported if a mapping can be made from the pipeline of the future API to the pipeline of the graphics processor.

[0121] Graphics Pipeline Programming

[0122] FIG. 16A is a block diagram illustrating a graphics processor command format **900** according to some embodiments. FIG. 16B is a block diagram illustrating a graphics processor command sequence **910** according to an embodiment. The solid lined boxes in FIG. 16A illustrate the components that are generally included in a graphics command while the dashed lines include components that are optional or that are only included in a sub-set of the graphics commands. The exemplary graphics processor command format **900** of FIG. 16A includes data fields to identify a target client **902** of the command, a command operation code (opcode) **904**, and the relevant data **906** for the command. A sub-opcode **905** and a command size **908** are also included in some commands.

[0123] In some embodiments, client 902 specifies the client unit of the graphics device that processes the command data. In some embodiments, a graphics processor command parser examines the client field of each command to condition the further processing of the command and route the command data to the appropriate client unit. In some embodiments, the graphics processor client units include a memory interface unit, a render unit, a 2D unit, a 3D unit, and a media unit. Each client unit has a corresponding processing pipeline that processes the commands. Once the command is received by the client unit, the client unit reads the opcode 904 and, if present, sub-opcode 905 to determine the operation to perform. The client unit performs the command using information in data field 906. For some commands an explicit command size 908 is expected to specify the size of the command. In some embodiments, the command parser automatically determines the size of at least some of the commands based on the command opcode. In some embodiments commands are aligned via multiples of a double word.

[0124] The flow diagram in FIG. 16B shows an exemplary graphics processor command sequence 910. In some embodiments, software or firmware of a data processing system that features an embodiment of a graphics processor uses a version of the command sequence shown to set up, execute, and terminate a set of graphics operations. A sample command sequence is shown and described for purposes of example only as embodiments are not limited to these specific commands or to this command sequence. Moreover, the commands may be issued as batch of commands in a command sequence, such that the graphics processor will process the sequence of commands in at least partially concurrence.

[0125] In some embodiments, the graphics processor command sequence 910 may begin with a pipeline flush command 912 to cause any active graphics pipeline to complete the currently pending commands for the pipeline. In some embodiments, the 3D pipeline 922 and the media pipeline 924 do not operate concurrently. The pipeline flush is performed to cause the active graphics pipeline to complete any pending commands. In response to a pipeline flush, the command parser for the graphics processor will pause command processing until the active drawing engines complete pending operations and the relevant read caches are invalidated. Optionally, any data in the render cache that is marked 'dirty' can be flushed to memory. In some embodiments, pipeline flush command 912 can be used for pipeline synchronization or before placing the graphics processor into a low power state.

[0126] In some embodiments, a pipeline select command 913 is used when a command sequence requires the graphics processor to explicitly switch between pipelines. In some embodiments, a pipeline select command 913 is required only once within an execution context before issuing pipeline commands unless the context is to issue commands for both pipelines. In some embodiments, a pipeline flush command is 912 is required immediately before a pipeline switch via the pipeline select command 913.

[0127] In some embodiments, a pipeline control command 914 configures a graphics pipeline for operation and is used to program the 3D pipeline 922 and the media pipeline 924. In some embodiments, pipeline control command 914 configures the pipeline state for the active pipeline. In one embodiment, the pipeline control command 914 is used for

pipeline synchronization and to clear data from one or more cache memories within the active pipeline before processing a batch of commands.

[0128] In some embodiments, return buffer state commands 916 are used to configure a set of return buffers for the respective pipelines to write data. Some pipeline operations require the allocation, selection, or configuration of one or more return buffers into which the operations write intermediate data during processing. In some embodiments, the graphics processor also uses one or more return buffers to store output data and to perform cross thread communication. In some embodiments, the return buffer state 916 includes selecting the size and number of return buffers to use for a set of pipeline operations.

[0129] The remaining commands in the command sequence differ based on the active pipeline for operations. Based on a pipeline determination 920, the command sequence is tailored to the 3D pipeline 922 beginning with the 3D pipeline state 930, or the media pipeline 924 beginning at the media pipeline state 940.

[0130] The commands for the 3D pipeline state 930 include 3D state setting commands for vertex buffer state, vertex element state, constant color state, depth buffer state, and other state variables that are to be configured before 3D primitive commands are processed. The values of these commands are determined at least in part based the particular 3D API in use. In some embodiments, 3D pipeline state 930 commands are also able to selectively disable or bypass certain pipeline elements if those elements will not be used.

[0131] In some embodiments, 3D primitive 932 command is used to submit 3D primitives to be processed by the 3D pipeline. Commands and associated parameters that are passed to the graphics processor via the 3D primitive 932 command are forwarded to the vertex fetch function in the graphics pipeline. The vertex fetch function uses the 3D primitive 932 command data to generate vertex data structures. The vertex data structures are stored in one or more return buffers. In some embodiments, 3D primitive 932 command is used to perform vertex operations on 3D primitives via vertex shaders. To process vertex shaders, 3D pipeline 922 dispatches shader execution threads to graphics processor execution units.

[0132] In some embodiments, 3D pipeline 922 is triggered via an execute 934 command or event. In some embodiments, a register write triggers command execution. In some embodiments execution is triggered via a 'go' or 'kick' command in the command sequence. In one embodiment command execution is triggered using a pipeline synchronization command to flush the command sequence through the graphics pipeline. The 3D pipeline will perform geometry processing for the 3D primitives. Once operations are complete, the resulting geometric objects are rasterized and the pixel engine colors the resulting pixels. Additional commands to control pixel shading and pixel back end operations may also be included for those operations.

[0133] In some embodiments, the graphics processor command sequence 910 follows the media pipeline 924 path when performing media operations. In general, the specific use and manner of programming for the media pipeline 924 depends on the media or compute operations to be performed. Specific media decode operations may be offloaded to the media pipeline during media decode. In some embodiments, the media pipeline can also be bypassed and media decode can be performed in whole or in part using resources

provided by one or more general purpose processing cores. In one embodiment, the media pipeline also includes elements for general-purpose graphics processor unit (GPGPU) operations, where the graphics processor is used to perform SIMD vector operations using computational shader programs that are not explicitly related to the rendering of graphics primitives.

[0134] In some embodiments, media pipeline 924 is configured in a similar manner as the 3D pipeline 922. A set of media pipeline state commands 940 are dispatched or placed into in a command queue before the media object commands 942. In some embodiments, media pipeline state commands 940 include data to configure the media pipeline elements that will be used to process the media objects. This includes data to configure the video decode and video encode logic within the media pipeline, such as encode or decode format. In some embodiments, media pipeline state commands 940 also support the use one or more pointers to “indirect” state elements that contain a batch of state settings.

[0135] In some embodiments, media object commands 942 supply pointers to media objects for processing by the media pipeline. The media objects include memory buffers containing video data to be processed. In some embodiments, all media pipeline states must be valid before issuing a media object command 942. Once the pipeline state is configured and media object commands 942 are queued, the media pipeline 924 is triggered via an execute command 944 or an equivalent execute event (e.g., register write). Output from media pipeline 924 may then be post processed by operations provided by the 3D pipeline 922 or the media pipeline 924. In some embodiments, GPGPU operations are configured and executed in a similar manner as media operations.

[0136] Graphics Software Architecture

[0137] FIG. 17 illustrates exemplary graphics software architecture for a data processing system 1000 according to some embodiments. In some embodiments, software architecture includes a 3D graphics application 1010, an operating system 1020, and at least one processor 1030. In some embodiments, processor 1030 includes a graphics processor 1032 and one or more general-purpose processor core(s) 1034. The graphics application 1010 and operating system 1020 each execute in the system memory 1050 of the data processing system.

[0138] In some embodiments, 3D graphics application 1010 contains one or more shader programs including shader instructions 1012. The shader language instructions may be in a high-level shader language, such as the High Level Shader Language (HLSL) or the OpenGL Shader Language (GLSL). The application also includes executable instructions 1014 in a machine language suitable for execution by the general-purpose processor core 1034. The application also includes graphics objects 1016 defined by vertex data.

[0139] In some embodiments, operating system 1020 is a Microsoft® Windows® operating system from the Microsoft Corporation, a proprietary UNIX-like operating system, or an open source UNIX-like operating system using a variant of the Linux kernel. When the Direct3D API is in use, the operating system 1020 uses a front-end shader compiler 1024 to compile any shader instructions 1012 in HLSL into a lower-level shader language. The compilation may be a just-in-time (JIT) compilation or the application can perform shader pre-compilation. In some embodiments,

high-level shaders are compiled into low-level shaders during the compilation of the 3D graphics application 1010.

[0140] In some embodiments, user mode graphics driver 1026 contains a back-end shader compiler 1027 to convert the shader instructions 1012 into a hardware specific representation. When the OpenGL API is in use, shader instructions 1012 in the GLSL high-level language are passed to a user mode graphics driver 1026 for compilation. In some embodiments, user mode graphics driver 1026 uses operating system kernel mode functions 1028 to communicate with a kernel mode graphics driver 1029. In some embodiments, kernel mode graphics driver 1029 communicates with graphics processor 1032 to dispatch commands and instructions.

[0141] IP Core Implementations

[0142] One or more aspects of at least one embodiment may be implemented by representative code stored on a machine-readable medium which represents and/or defines logic within an integrated circuit such as a processor. For example, the machine-readable medium may include instructions which represent various logic within the processor. When read by a machine, the instructions may cause the machine to fabricate the logic to perform the techniques described herein. Such representations, known as “IP cores,” are reusable units of logic for an integrated circuit that may be stored on a tangible, machine-readable medium as a hardware model that describes the structure of the integrated circuit. The hardware model may be supplied to various customers or manufacturing facilities, which load the hardware model on fabrication machines that manufacture the integrated circuit. The integrated circuit may be fabricated such that the circuit performs operations described in association with any of the embodiments described herein.

[0143] FIG. 18 is a block diagram illustrating an IP core development system 1100 that may be used to manufacture an integrated circuit to perform operations according to an embodiment. The IP core development system 1100 may be used to generate modular, re-usable designs that can be incorporated into a larger design or used to construct an entire integrated circuit (e.g., an SOC integrated circuit). A design facility 1130 can generate a software simulation 1110 of an IP core design in a high level programming language (e.g., C/C++). The software simulation 1110 can be used to design, test, and verify the behavior of the IP core. A register transfer level (RTL) design can then be created or synthesized from the simulation model 1100. The RTL design 1115 is an abstraction of the behavior of the integrated circuit that models the flow of digital signals between hardware registers, including the associated logic performed using the modeled digital signals. In addition to an RTL design 1115, lower-level designs at the logic level or transistor level may also be created, designed, or synthesized. Thus, the particular details of the initial design and simulation may vary.

[0144] The RTL design 1115 or equivalent may be further synthesized by the design facility into a hardware model 1120, which may be in a hardware description language (HDL), or some other representation of physical design data. The HDL may be further simulated or tested to verify the IP core design. The IP core design can be stored for delivery to a 3rd party fabrication facility 1165 using non-volatile memory 1140 (e.g., hard disk, flash memory, or any non-volatile storage medium). Alternatively, the IP core design may be transmitted (e.g., via the Internet) over a wired connection 1150 or wireless connection 1160. The fabrica-

tion facility 1165 may then fabricate an integrated circuit that is based at least in part on the IP core design. The fabricated integrated circuit can be configured to perform operations in accordance with at least one embodiment described herein.

[0145] FIG. 19 is a block diagram illustrating an exemplary system on a chip integrated circuit 1200 that may be fabricated using one or more IP cores, according to an embodiment. The exemplary integrated circuit includes one or more application processors 1205 (e.g., CPUs), at least one graphics processor 1210, and may additionally include an image processor 1215 and/or a video processor 1220, any of which may be a modular IP core from the same or multiple different design facilities. The integrated circuit includes peripheral or bus logic including a USB controller 1225, universal asynchronous receiver/transmitter (UART) controller 1230, a serial peripheral interface (SPI)/secure digital input output (SDIO) controller 1235, and an integrated interchip sound (I2S)/inter-integrated circuit (I2C) controller 1240. Additionally, the integrated circuit can include a display device 1245 coupled to one or more of a high-definition multimedia interface (HDMI) controller 1250 and a mobile industry processor interface (MIPI) display interface 1255. Storage may be provided by a flash memory subsystem 1260 including flash memory and a flash memory controller. Memory interface may be provided via a memory controller 1265 for access to SDRAM or SRAM memory devices. Some integrated circuits additionally include an embedded security engine 1270.

[0146] Additionally, other logic and circuits may be included in the processor of integrated circuit 1200, including additional graphics processors/cores, peripheral interface controllers, or general purpose processor cores.

[0147] Advantageously, in some embodiments any of the system 100, the processor 200, graphics processor 300, graphics processing engine 410, graphics processor 500, thread execution logic 600, graphics processor 800, data processing system 1000, and/or the integrated circuit 1200 may be integrated or configured with any of the various embodiments described herein (e.g. or portions thereof), including, for example, those described in the following Additional Notes and Examples.

[0148] Additional Notes and Examples:

[0149] Example 1 may include an electronic processing system, comprising an application processor, persistent storage media communicatively coupled to the application processor, a graphics processor communicatively coupled to the application processor, and a gesture tracker communicatively coupled to the application processor, wherein the graphics processor is to generate an image of a virtual input device, the gesture tracker is to determine a position of a user's finger relative to the virtual input device, and wherein the graphics processor is further to generate an image of a virtual finger based on the determined position of the user's finger relative to the virtual input device.

[0150] Example 2 may include the system of Example 1, further comprising an input generator communicatively coupled to the gesture tracker to generate input data based on the determined position of the user's finger relative to the virtual input device.

[0151] Example 3 may include the system of any of Examples 1 to 2, further comprising a device tracker communicatively coupled to the graphics processor to identify a characteristic of a physical input device, wherein the graph-

ics processor is further to generate the image of the virtual input device based on the identified characteristic of the physical input device.

[0152] Example 4 may include a graphics apparatus, comprising an image generator, and a gesture tracker communicatively coupled to the image generator, wherein the image generator is to generate an image of a virtual input device, the gesture tracker is to determine a position of a user's finger relative to the virtual input device, and wherein the image generator is further to generate an image of a virtual finger based on the determined position of the user's finger relative to the virtual input device.

[0153] Example 5 may include the apparatus of Example 4, wherein the gesture tracker is further to determine a position of a user's hand relative to the virtual input device, and wherein the image generator is further to generate an image of a virtual hand based on the determined position of the user's hand relative to the virtual input device.

[0154] Example 6 may include the apparatus of Example 4, further comprising an input generator communicatively coupled to the gesture tracker to generate input data based on the determined position of the user's finger relative to the virtual input device.

[0155] Example 7 may include the apparatus of Example 4, further comprising a device tracker communicatively coupled to the image generator to identify a characteristic of a physical input device, wherein the image generator is further to generate the image of the virtual input device based on the identified characteristic of the physical input device.

[0156] Example 8 may include the apparatus of Example 7, wherein the device tracker is further to determine a position of the physical input device relative to the user, the gesture tracker is further to determine a position of the user's hand relative to the position of the physical input device, the image generator is further to generate an image of a virtual representation of the physical input device based on the position of the physical input device relative to the user and wherein the image generator is further to generate an image of a virtual hand based on the determined position of the user's hand relative to the physical input device.

[0157] Example 9 may include the apparatus of Example 7, wherein the image generator is further to load a three-dimensional model of the virtual input device based on the identified characteristic of the physical input device.

[0158] Example 10 may include the apparatus of any of Examples 4 to 9, wherein the virtual input device comprises one of a virtual keyboard, a virtual mouse, a virtual touchpad, a virtual stylus, and a virtual scroll wheel.

[0159] Example 11 may include a method of displaying virtual objects, comprising generating an image of a virtual input device, determining a position of a user's finger relative to the virtual input device, and generating an image of a virtual finger based on the determined position of the user's finger relative to the virtual input device.

[0160] Example 12 may include the method of Example 11, further comprising determining a position of a user's hand relative to the virtual input device, and generating an image of a virtual hand based on the determined position of the user's hand relative to the virtual input device.

[0161] Example 13 may include the method of Example 11, further comprising generating input data based on the determined position of the user's finger relative to the virtual input device.

[0162] Example 14 may include the method of Example 13, further comprising generating a visual indication corresponding to the generated input data.

[0163] Example 15 may include the method of Example 11, further comprising identifying a characteristic of a physical input device, and generating the image of the virtual input device based on the identified characteristic of the physical input device.

[0164] Example 16 may include the method of Example 15, further comprising determining a position of the physical input device relative to the user, determining a position the user's hand relative to the position of the physical input device, generating an image of a virtual representation of the physical input device based on the position of the physical input device relative to the user, and generating an image of a virtual hand based on the determined position of the user's hand relative to the physical input device.

[0165] Example 17 may include the method of Example 15, further comprising loading a three-dimensional model of the virtual input device based on the identified characteristic of the physical input device.

[0166] Example 18 may include the method of any of Examples 11 to 17, wherein the virtual input device comprises one of a virtual keyboard, a virtual mouse, a virtual touchpad, a virtual stylus, and a virtual scroll wheel.

[0167] Example 19 may include at least one computer readable medium, comprising a set of instructions, which when executed by a computing device, cause the computing device to generate an image of a virtual input device, determine a position of a user's finger relative to the virtual input device, and generate an image of a virtual finger based on the determined position of the user's finger relative to the virtual input device.

[0168] Example 20 may include the at least one computer readable medium of Example 19, comprising a further set of instructions, which when executed by the computing device, cause the computing device to determine a position of a user's hand relative to the virtual input device, and generate an image of a virtual hand based on the determined position of the user's hand relative to the virtual input device.

[0169] Example 21 may include the at least one computer readable medium of Example 19, comprising a further set of instructions, which when executed by the computing device, cause the computing device to generate input data based on the determined position of the user's finger relative to the virtual input device.

[0170] Example 22 may include the at least one computer readable medium of Example 19, comprising a further set of instructions, which when executed by the computing device, cause the computing device to identify a characteristic of a physical input device, and generate the image of the virtual input device based on the identified characteristic of the physical input device.

[0171] Example 23 may include the at least one computer readable medium of Example 22, comprising a further set of instructions, which when executed by the computing device, cause the computing device to determine a position of the physical input device relative to the user, determine a position the user's hand relative to the position of the physical input device, generate an image of a virtual representation of the physical input device based on the position of the physical input device relative to the user, and generate an image of a virtual hand based on the determined position of the user's hand relative to the physical input device.

[0172] Example 24 may include the at least one computer readable medium of Example 22, comprising a further set of instructions, which when executed by the computing device, cause the computing device to load a three-dimensional model of the virtual input device based on the identified characteristic of the physical input device.

[0173] Example 25 may include the at least one computer readable medium of any of Examples 19 to 24, wherein the virtual input device comprises one of a virtual keyboard, a virtual mouse, a virtual touchpad, a virtual stylus, and a virtual scroll wheel.

[0174] Example 26 may include the at least one computer readable medium of Example 21, comprising a further set of instructions, which when executed by the computing device, cause the computing device to generate a visual indication corresponding to the generated input data.

[0175] Example 27 may include a graphics apparatus, comprising means for generating an image of a virtual input device, means for determining a position of a user's finger relative to the virtual input device, and means for generating an image of a virtual finger based on the determined position of the user's finger relative to the virtual input device.

[0176] Example 28 may include the apparatus of Example 27, further comprising means for determining a position of a user's hand relative to the virtual input device, and means for generating an image of a virtual hand based on the determined position of the user's hand relative to the virtual input device.

[0177] Example 29 may include the apparatus of Example 27, further comprising means for generating input data based on the determined position of the user's finger relative to the virtual input device.

[0178] Example 30 may include the apparatus of Example 27, further comprising means for identifying a characteristic of a physical input device, and means for generating the image of the virtual input device based on the identified characteristic of the physical input device.

[0179] Example 31 may include the apparatus of Example 30, further comprising means for determining a position of the physical input device relative to the user, means for determining a position the user's hand relative to the position of the physical input device, means for generating an image of a virtual representation of the physical input device based on the position of the physical input device relative to the user, and means for generating an image of a virtual hand based on the determined position of the user's hand relative to the physical input device.

[0180] Example 32 may include the apparatus of Example 30, further comprising means for loading a three-dimensional model of the virtual input device based on the identified characteristic of the physical input device.

[0181] Example 33 may include the apparatus of any of Examples 27 to 32, wherein the virtual input device comprises one of a virtual keyboard, a virtual mouse, a virtual touchpad, a virtual stylus, and a virtual scroll wheel.

[0182] Example 34 may include the apparatus of Example 29, further comprising means generating a visual indication corresponding to the generated input data.

[0183] Example 35 may include a method of displaying a virtual object, comprising identifying a characteristic of a physical input device, loading a three-dimensional (3D) model of a virtual input device based on the identified

characteristic of the physical input device, and generating an image of the virtual input device based on the loaded 3D model.

[0184] Example 36 may include the method of Example 35, further comprising identifying a serial number of the physical input device, and downloading the 3D model from a cloud-based service based on the serial number of the physical device.

[0185] Example 37 may include the method of Example 35, further comprising storing one or more 3D models in a database, and loading the 3D model from the one or more stored 3D models.

[0186] Example 38 may include the method of Example 35, further comprising identifying a physical layout of the physical input device, and generating the 3D model on the fly based on the identified physical layout of the physical input device.

[0187] Example 39 may include the method of Example 35, further comprising visually mapping the physical input device, and generating the 3D model on the fly based on the visual map of the physical input device.

[0188] Embodiments are applicable for use with all types of semiconductor integrated circuit (“IC”) chips. Examples of these IC chips include but are not limited to processors, controllers, chipset components, programmable logic arrays (PLAs), memory chips, network chips, systems on chip (SoCs), SSD/NAND controller ASICs, and the like. In addition, in some of the drawings, signal conductor lines are represented with lines. Some may be different, to indicate more constituent signal paths, have a number label, to indicate a number of constituent signal paths, and/or have arrows at one or more ends, to indicate primary information flow direction. This, however, should not be construed in a limiting manner. Rather, such added detail may be used in connection with one or more exemplary embodiments to facilitate easier understanding of a circuit. Any represented signal lines, whether or not having additional information, may actually comprise one or more signals that may travel in multiple directions and may be implemented with any suitable type of signal scheme, e.g., digital or analog lines implemented with differential pairs, optical fiber lines, and/or single-ended lines.

[0189] Example sizes/models/values/ranges may have been given, although embodiments are not limited to the same. As manufacturing techniques (e.g., photolithography) mature over time, it is expected that devices of smaller size could be manufactured. In addition, well known power/ground connections to IC chips and other components may or may not be shown within the figures, for simplicity of illustration and discussion, and so as not to obscure certain aspects of the embodiments. Further, arrangements may be shown in block diagram form in order to avoid obscuring

embodiments, and also in view of the fact that specifics with respect to implementation of such block diagram arrangements are highly dependent upon the platform within which the embodiment is to be implemented, i.e., such specifics should be well within purview of one skilled in the art. Where specific details (e.g., circuits) are set forth in order to describe example embodiments, it should be apparent to one skilled in the art that embodiments can be practiced without, or with variation of, these specific details. The description is thus to be regarded as illustrative instead of limiting.

[0190] The term “coupled” may be used herein to refer to any type of relationship, direct or indirect, between the components in question, and may apply to electrical, mechanical, fluid, optical, electromagnetic, electromechanical or other connections. In addition, the terms “first”, “second”, etc. may be used herein only to facilitate discussion, and carry no particular temporal or chronological significance unless otherwise indicated.

[0191] As used in this application and in the claims, a list of items joined by the term “one or more of” may mean any combination of the listed terms. For example, the phrases “one or more of A, B or C” may mean A; B; C; A and B; A and C; B and C; or A, B and C.

[0192] Those skilled in the art will appreciate from the foregoing description that the broad techniques of the embodiments can be implemented in a variety of forms. Therefore, while the embodiments have been described in connection with particular examples thereof, the true scope of the embodiments should not be so limited since other modifications will become apparent to the skilled practitioner upon a study of the drawings, specification, and following claims.

We claim:

1. An electronic processing system, comprising:
 - an application processor;
 - persistent storage media communicatively coupled to the application processor;
 - a graphics processor communicatively coupled to the application processor; and
 - a gesture tracker communicatively coupled to the application processor, wherein:
 - the graphics processor is to generate an image of a virtual input device,
 - the gesture tracker is to determine a position of a user’s finger relative to the virtual input device, and wherein
 - the graphics processor is further to generate an image of a virtual finger based on the determined position of the user’s finger relative to the virtual input device.

* * * * *