



(51) International Patent Classification: **G06F 21/10** (2013.01)

(21) International Application Number: PCT/US2016/043903

(22) International Filing Date: 25 July 2016 (25.07.2016)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data: 14/821,174 7 August 2015 (07.08.2015) US

(71) Applicant: **QUALCOMM INCORPORATED** [US/US];
Attn: International IP Administration, 5775 Morehouse Drive, San Diego, California 92121-1714 (US).

(72) Inventors: **SHARP, Colin Christopher**; 5775 Morehouse Drive, San Diego, California 92121-1714 (US).
VISWANATHAN, Ramesh; 5775 Morehouse Drive, San Diego, California 92121-1714 (US).

(74) Agent: **PRIEM, David F.**; Shumaker & Sieffert, P.A., 1625 Radio Drive, Woodbury, MN 55125 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY,

BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.


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

Declarations under Rule 4.17:

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

Published:

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

(54) Title: HARDWARE ENFORCED CONTENT PROTECTION FOR GRAPHICS PROCESSING UNITS

(57) Abstract: This disclosure proposes techniques for graphics processing. In one example, a graphics processing unit (GPU) is configured to access a memory according to one of an unsecure mode and a secure mode. The GPU may include a memory access controller configured to direct memory transactions from at least one hardware unit of the GPU to an unsecure memory unit or a secure memory unit based on the unsecure mode or secure mode and a resource descriptor associated with a memory resource.

HARDWARE ENFORCED CONTENT PROTECTION FOR GRAPHICS PROCESSING UNITS

TECHNICAL FIELD

[0001] This disclosure relates to techniques for graphics processing, and more specifically to techniques for content protection.

BACKGROUND

[0002] Modern operating systems, including open platforms (e.g., Android or other open source platforms) and closed platforms (e.g., Microsoft Windows®), are not typically trusted in terms of protecting secure content which is streamed to, or processed by, such open platforms. While modern operating systems provide a level of security via the user-kernel mode separation, ultimately components of kernel mode, both in closed platforms, and particularly in open platforms do not provide a strong level of trust. Kernel mode drivers can easily be installed, and a malicious kernel mode driver naturally bypasses the security boundary. Kernel mode hardware drivers in such open platforms are used to control the operation of hardware (e.g., graphics processing units (GPUs)) that may process secure content. However, because such drivers are often open source, and/or not considered to be “secure” in relation to protected content, they are more susceptible to alteration by third parties. Such alterations may cause the protected content (e.g., digital rights managed (DRM) content) that is streamed through or processed by the hardware controlled by such drivers to be stored in unsecure memories and copied. As such, control of secure content on open platforms is often difficult.

SUMMARY

[0003] In general, this disclosure describes techniques for hardware enforced content protection for a graphics processing unit (GPU). To control secure content on a hardware platform, access to secure memory may be controlled by hardware such as a GPU.

[0004] In one example of the disclosure, an apparatus for graphics processing comprises a GPU configured to access a memory according to one of an unsecure mode and a secure mode, the GPU comprising a memory access controller configured to direct

memory transactions from at least one hardware unit of the GPU to a secure context bank in a memory controller when the GPU is operating in a secure mode, and configured to direct memory transactions from the at least one hardware unit of the GPU to an unsecure context bank in the memory controller when the GPU is operating in the unsecure mode.

[0005] In another example of the disclosure, a GPU comprises one or more hardware units configured to access a memory according to one of an unsecure mode and a secure mode of the GPU, and a memory access controller configured to direct memory transactions from at least one of the one or more hardware units of the GPU to a secure context bank in a memory controller when the GPU is operating in a secure mode, and configured to direct memory transactions from the at least one of the one or more hardware units of the GPU to an unsecure context bank in the memory controller when the GPU is operating in the unsecure mode.

[0006] In another example of the disclosure, a method for graphics processing comprises accessing an unsecure portion of a memory, with a GPU, according an unsecure mode by directing memory transactions from at least one hardware unit of the GPU to an unsecure context bank in a memory controller, and accessing a secure portion of the memory, with the GPU, according to a secure mode by directing memory transactions from the at least one hardware unit of the GPU to a secure context bank in the memory controller.

[0007] In another example of the disclosure, an apparatus for graphics processing comprises means for accessing an unsecure portion of a memory according an unsecure mode by directing memory transactions from at least one hardware unit of a GPU to an unsecure context bank in a memory controller, and means for accessing a secure portion of the memory according to a secure mode by directing memory transactions from the at least one hardware unit of the GPU to a secure context bank in the memory controller.

[0008] In another example of the disclosure, an apparatus for graphics processing comprises a GPU configured to access a first memory unit according to one of an unsecure mode and a secure mode and a respective resource descriptor associated with each of a plurality of memory resources, the GPU comprising a memory access controller configured to read the respective resource descriptor associated with each of the plurality of memory resources, the memory access controller configured to receive a request for a memory transaction to the first memory unit, the memory access controller configured to, in response to the request, direct all read and write memory transactions

relating to memory resources of the plurality of memory resources for which the respective resource descriptor is a secure resource descriptor to a secure portion of the first memory unit when the GPU is operating according to the secure mode, the memory access controller configured to, in response to the request, direct all read memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is an unsecure resource descriptor to an unsecure portion of the first memory unit when the GPU is operating according to the secure mode, and the memory access controller configured to, in response to the request, drop all write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is the unsecure resource descriptor when the GPU is operating according to the secure mode.

[0009] In another example of the disclosure, a method comprises reading a respective resource descriptor for a respective memory resource of a plurality of memory resources, receiving a request for a memory transaction to a first memory unit, directing, in response to the request, read and write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is a secure resource descriptor to a secure portion of the first memory unit when a GPU is operating according to a secure mode, directing, in response to the request, read memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is an unsecure resource descriptor to an unsecure portion of the first memory unit when the GPU is operating according to the secure mode, and dropping, in response to the request, write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is the unsecure resource descriptor when the GPU is operating according to the secure mode.

[0010] In another example of the disclosure, an apparatus for graphics processing comprises means for reading a respective resource descriptor for a respective memory resource of a plurality of memory resources, means for receiving a request for a memory transaction to a first memory unit, means for directing, in response to the request, read and write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is a secure resource descriptor to a secure portion of the first memory unit when a GPU is operating according to a secure mode, means for directing, in response to the request, read memory transactions relating to memory resources of the plurality of memory resources for which the respective

resource descriptor is an unsecure resource descriptor to an unsecure portion of the first memory unit when the GPU is operating according to the secure mode, and means for dropping, in response to the request, write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is the unsecure resource descriptor when the GPU is operating according to the secure mode.

[0011] In another example, this disclosure describes a computer-readable storage medium storing instructions that, when executed, cause one or more processors to read a respective resource descriptor for a respective memory resource of a plurality of memory resources, receive a request for a memory transaction to a first memory unit, direct, in response to the request, read and write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is a secure resource descriptor to a secure portion of the first memory unit when a GPU is operating according to a secure mode, direct, in response to the request, read memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is an unsecure resource descriptor to an unsecure portion of the first memory unit when the GPU is operating according to the secure mode, and drop, in response to the request, write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is the unsecure resource descriptor when the GPU is operating according to the secure mode.

[0012] The details of one or more examples are set forth in the accompanying drawings and the description below. Other features, objects, and advantages will be apparent from the description and drawings, and from the claims.

BRIEF DESCRIPTION OF DRAWINGS

[0013] FIG. 1 is a block diagram showing an example computing device configured to use the techniques of this disclosure.

[0014] FIG. 2 is a conceptual diagram illustrating an example physical page of a system memory of FIG. 1.

[0015] FIG. 3 is a block diagram showing example processing units configured to use the techniques of this disclosure.

[0016] FIG. 4 is a block diagram showing an example structure configured to perform the hardware enforced content protection techniques of this disclosure.

[0017] FIG. 5 is a block diagram showing another example structure configured to perform the hardware enforced content protection techniques of this disclosure.

[0018] FIG. 6 is a block diagram showing another example structure configured to perform the hardware enforced content protection techniques of this disclosure.

[0019] FIG. 7A is a block diagram showing another example structure configured to perform the hardware enforced content protection techniques of this disclosure.

[0020] FIG. 7B is a block diagram showing another example structure configured to perform the hardware enforced content protection techniques of this disclosure.

[0021] FIG. 8 is a block diagram showing cache clearing techniques according to one example of this disclosure.

[0022] FIG. 9 is a block diagram showing cache clearing techniques according to another example of this disclosure.

[0023] FIG. 10 is a flowchart illustrating a method according to one example of the disclosure.

[0024] FIG. 11 is a flowchart illustrating another example method according to another example of the disclosure.

DETAILED DESCRIPTION

[0025] This disclosure relates to techniques for graphics processing, and more specifically to techniques for hardware enforced content protection for a graphics processing unit (GPU).

[0026] Modern operating systems, including open platforms (e.g., Android or other open source platforms) and closed platforms (e.g., Microsoft Windows®), are not typically trusted in terms of protecting secure content which is streamed to, or processed by, such open platforms. While modern operating systems provide a level of security via the user-kernel mode separation, ultimately components of kernel mode, both in closed platforms, and particularly in open platforms, do not provide a strong level of trust. Kernel mode drivers can easily be installed, and a malicious kernel mode driver naturally bypasses the security boundary. Kernel mode hardware drivers in such open platforms are used to control the operation of hardware (e.g., graphics processing units (GPUs)) that may process secure content. However, because such drivers are often

open source, and/or not considered to be “secure” in relation to protected content, they are more susceptible to alteration by third parties. Such alterations may cause the protected content (e.g., digital rights managed (DRM) content) that is streamed through or processed by the hardware controlled by such drivers to be stored in unsecure memories and copied. As such, control of secure content on open platforms is often difficult. To address this problem, this disclosure proposes a method and apparatus whereby access to secure memory is controlled by the hardware itself (e.g., by a GPU). **[0027]** Rather than controlling hardware access to secure or unsecure memory directly through driver code, this disclosure proposes, in one example, using the graphics driver (e.g., an open source unsecure driver) to only place the GPU in either a secure mode or an unsecure mode. Once placed in the secure mode, the GPU components may be configured such that read and/or write access to secure and unsecure memory by the GPU may be restricted based on the mode (i.e., secure or unsecure mode) of the GPU. For example, in secure mode, certain GPU components may be configured such that they are restricted to only making writes into the secure memory region. This prevents an untrusted driver from using the GPU to copy memory content from the secure memory region to an unsecure memory region. Other techniques for restricting GPU access to secure memory in the secure mode, placing the GPU into one of an unsecure mode or secure mode, and associating certain data resources with secure memory or unsecure memory will be discussed in more detail below.

[0028] In one example of the disclosure, in this secure mode, the GPU may be configured to read both secure (e.g., copy protected (CP)) content as well as unsecure content (e.g., content stored in an unsecured memory). In the unsecure mode, the GPU may be configured such that GPU components are denied all access to secure memory. In this way, even if the unsecure driver were altered to place the GPU in an unsecure mode, the GPU itself would be prevented from reading any data from a secure memory. As such, access to secure content in the secure memory is prevented.

[0029] FIG. 1 is a block diagram illustrating an example computing device 2 that may be used to implement the techniques of this disclosure for hardware enforced content protection for a GPU. Computing device 2 may comprise, for example, a personal computer, a desktop computer, a laptop computer, a tablet computer, a computer workstation, a video game platform or console, a mobile telephone such as, e.g., a cellular or satellite telephone, a landline telephone, an Internet telephone, a so-called smartphone, a handheld device such as a portable video game device or a personal

digital assistant (PDA), a personal music player, a video player, a display device, a television, a television set-top box, a server, an intermediate network device, a mainframe computer, any mobile device, or any other type of device that processes and/or displays graphical data.

[0030] As illustrated in the example of FIG. 1, computing device 2 may include a user input interface 4, a central processing unit (CPU) 6, one or more memory controllers 8, a system memory 10, a graphics processing unit (GPU) 12, a graphics memory 14, a display interface 16, a display 18 and buses 20 and 22. Note that, in some examples, graphics memory 14 may be “on-chip” with GPU 12. In some cases, all hardware elements shown in FIG. 1 may be on-chip, for example, in a system on a chip (SoC) design. User input interface 4, CPU 6, memory controllers 8, GPU 12 and display interface 16 may communicate with each other using bus 20. Memory controllers 8 and system memory 10 may also communicate with each other using bus 22. Buses 20, 22 may be any of a variety of bus structures, such as a third generation bus (e.g., a HyperTransport bus or an InfiniBand bus), a second generation bus (e.g., an Advanced Graphics Port bus, a Peripheral Component Interconnect (PCI) Express bus, or an Advanced eXtensible Interface (AXI) bus) or another type of bus or device interconnect. It should be noted that the specific configuration of buses and communication interfaces between the different components shown in FIG. 1 is merely exemplary, and other configurations of computing devices and/or other graphics processing systems with the same or different components may be used to implement the techniques of this disclosure.

[0031] CPU 6 may comprise a general-purpose or a special-purpose processor that controls operation of computing device 2. A user may provide input to computing device 2 to cause CPU 6 to execute one or more software applications. The software applications that execute on CPU 6 may include, for example, an operating system, a word processor application, an email application, a spread sheet application, a media player application, a video game application, a graphical user interface application or another program. Additionally, CPU 6 may execute a GPU driver 7 for controlling the operation of GPU 12. The user may provide input to computing device 2 via one or more input devices (not shown) such as a keyboard, a mouse, a microphone, a touch pad, a touch screen, or another input device that is coupled to computing device 2 via user input interface 4.

[0032] The software applications that execute on CPU 6 may include one or more graphics rendering instructions that instruct CPU 6 to cause the rendering of graphics data to display 18. In some examples, the software instructions may conform to a graphics application programming interface (API), such as, e.g., an Open Graphics Library (OpenGL[®]) API, an Open Graphics Library Embedded Systems (OpenGL ES) API, an Open Computing Language (OpenCL[®]) API, a Direct3D API, an X3D API, a RenderMan API, a WebGL API or any other public or proprietary standard graphics API. In order to process the graphics rendering instructions, CPU 6 may issue one or more graphics rendering commands to GPU 12 (e.g., through GPU driver 7) to cause GPU 12 to perform some or all of the rendering of the graphics data. In some examples, the graphics data to be rendered may include a list of graphics primitives, e.g., points, lines, triangles, quadrilaterals, triangle strips, etc.

[0033] Memory controllers 8 facilitate the transfer of data going into and out of system memory 10. For example, memory controllers 8 may receive memory read and write commands, and service such commands with respect to system memory 10 in order to provide memory services for the components in computing device 2. Memory controllers 8 are communicatively coupled to system memory 10 via memory bus 22. Although memory controllers 8 are illustrated in FIG. 1 as being a processing module that is separate from both CPU 6 and system memory 10, in other examples, some or all of the functionality of memory controllers 8 may be implemented on one or any of CPU 6, GPU 12 and system memory 10. System memory 10 may comprise one or more memory units. The memory units may be divided physically (e.g., separate physical disks or solid state memory units) or may be divided by memory address range. In particular, system memory 10 may be divided into two or more memory units consisting of “secure” memory units and “unsecure” memory units. In some examples, secure memory units may utilize encryption and/or other digital rights management (DRM) techniques to prevent access, copying, or deciphering of data stored thereon.

[0034] Memory controllers 8 may also include one or more memory management units (MMUs), including an IOMMU (i.e., input/output MMU) for controlling IO device access (e.g., a GPU) to system memory 10. The memory management units may implement a virtual memory system. The virtual memory space may be divided into a plurality of virtual pages. These virtual pages may be contiguous, but the physical pages in system memory 10 to which these virtual pages correspond may not be

contiguous in system memory 10. Pages may be considered as the minimum units that an MMU may be able to manage.

[0035] Modern operating systems (OS) that run on central processing units (CPUs) typically use a virtual memory scheme for allocating memory to multiple programs operating on the CPU. Virtual memory is a memory management technique that virtualizes a computer system's physical memory (e.g., RAM, disk storage, etc.) so that an application need only refer to one set of memory (i.e., the virtual memory). Virtual memory consists of contiguous address spaces that are mapped to locations in physical memory. In this way, the fragmentation of physical memory is "hidden" from the applications, which instead may interact with contiguous blocks of virtual memory. The contiguous blocks in virtual memory are typically arranged into "pages." Each page is some fixed length of contiguous blocks of virtual memory addresses. Mapping from the virtual memory to the physical memory is often handled by a memory management unit (MMU). Virtual memory space that is currently mapped to locations in physical memory is considered to be "backed" to physical memory.

[0036] The mapping of locations in virtual memory space to physical memory is stored with a translation lookaside buffer (TLB). The TLB is used by the MMU to quickly translate virtual addresses to physical addresses. The TLB may be implemented as a content-addressable memory (CAM) that uses a virtual memory address as an input and outputs a physical memory address. The MMU may then quickly retrieve the requested data using the output physical memory address.

[0037] FIG. 2 is a conceptual diagram illustrating an example physical page of system memory 10. For example, FIG. 2 illustrates an IOMMU 40 including a virtual page 42 which includes four sections (sections 0-3). It should be understood that virtual page 42 is a virtual construct that is illustrated in FIG. 2 for ease of understanding. In FIG. 2, system memory 10 may include a physical page 44 that corresponds to virtual page 42.

[0038] Physical page 44 may be stored across multiple memory units of system memory 10. For example, physical page 44 may encompass both memory unit 11A and memory unit 11N. In one example, memory unit 11A is a "secure" memory unit and memory unit 11N is an "unsecure" memory unit. Memory unit 11A may store a portion of physical page 44, indicated as portion 44A, and memory unit 11N may store a portion of physical page 44, indicated as portion 44B. As illustrated, memory unit 11A stores section 0 and section 2 of physical page 44, and memory unit 11N stores section 1 and section 3 of physical page 44.

[0039] The example of FIG. 2 only includes two memory units for purposes of illustration, but any number of memory units may be used. For instance, referring back to FIG. 1, GPU driver 7 may transmit instructions that cause GPU 12 to store pixel values or any other computed value, and may transmit the virtual addresses for where the pixel values are to be stored. GPU 12, in turn, may request IOMMU 40 to store the pixel values in accordance with the virtual addresses. IOMMU 40, in turn, may map the virtual addresses to physical addresses and store the pixel values in pages of system memory 10 in an interleaving manner based on the physical addresses.

[0040] Returning to FIG. 1, system memory 10 may store program modules and/or instructions that are accessible for execution by CPU 6 and/or data for use by the programs executing on CPU 6. For example, system memory 10 may store a window manager application that is used by CPU 6 to present a graphical user interface (GUI) on display 18. In addition, system memory 10 may store user applications and application surface data associated with the applications. System memory 10 may additionally store information for use by and/or generated by other components of computing device 2. For example, system memory 10 may act as a device memory for GPU 12 and may store data to be operated on by GPU 12 as well as data resulting from operations performed by GPU 12. For example, system memory 10 may store DRM protected game content or decoded video produced by GPU 12. In this situation, such DRM-protected content is preferably stored in a secure memory unit of system memory 10. As other examples, system memory 10 may store other graphics data such as any combination of texture buffers, depth buffers, stencil buffers, vertex buffers, frame buffers, or the like. System memory 10 may include one or more volatile or non-volatile memories or storage devices, such as, for example, random access memory (RAM), static RAM (SRAM), dynamic RAM (DRAM), read-only memory (ROM), erasable programmable ROM (EPROM), electrically erasable programmable ROM (EEPROM), Flash memory, a magnetic data media or an optical storage media.

[0041] GPU 12 may be configured to perform graphics operations to render one or more graphics primitives to display 18. Thus, when one of the software applications executing on CPU 6 requires graphics processing, CPU 6 may provide graphics commands and graphics data to GPU 12 for rendering to display 18. The graphics data may include, e.g., drawing commands, state information, primitive information, texture information, etc. GPU 12 may, in some instances, be built with a highly-parallel structure that provides more efficient processing of complex graphic-related operations

than CPU 6. For example, GPU 12 may include a plurality of processing elements that are configured to operate on multiple vertices or pixels in a parallel manner. The highly parallel nature of GPU 12 may, in some instances, allow GPU 12 to draw graphics images (e.g., GUIs and two-dimensional (2D) and/or three-dimensional (3D) graphics scenes) onto display 18 more quickly than drawing the scenes directly to display 18 using CPU 6.

[0042] GPU 12 may, in some instances, be integrated into a motherboard of computing device 2. In other instances, GPU 12 may be present on a graphics card that is installed in a port in the motherboard of computing device 2 or may be otherwise incorporated within a peripheral device configured to interoperate with computing device 2. GPU 12 may include one or more processors, such as one or more microprocessors, application specific integrated circuits (ASICs), field programmable gate arrays (FPGAs), digital signal processors (DSPs), or other equivalent integrated or discrete logic circuitry.

[0043] GPU 12 may be directly coupled to graphics memory 14. Thus, GPU 12 may read data from and write data to graphics memory 14 without using bus 20. In other words, GPU 12 may process data locally using a local storage, instead of using other, slower system memory. This allows GPU 12 to operate in a more efficient manner by eliminating the need of GPU 12 to read and write data via system bus 20, which may experience heavy bus traffic. In some instances, however, GPU 12 may not include a separate memory, but instead utilize system memory 10 via bus 20. Graphics memory 14 may include one or more volatile or non-volatile memories or storage devices, such as, e.g., random access memory (RAM), static RAM (SRAM), dynamic RAM (DRAM), erasable programmable ROM (EPROM), electrically erasable programmable ROM (EEPROM), Flash memory, a magnetic data media or an optical storage media.

[0044] CPU 6 and/or GPU 12 may store rendered image data in a frame buffer 15. Typically, frame buffer 15 would be allocated within system memory 10, but may in some circumstances be an independent memory. Display interface 16 may retrieve the data from frame buffer 15 and configure display 18 to display the image represented by the rendered image data. In some examples, display interface 16 may include a digital-to-analog converter (DAC) that is configured to convert the digital values retrieved from the frame buffer into an analog signal consumable by display 18. In other examples, display interface 16 may pass the digital values directly to display 18 for processing. Display 18 may include a monitor, a television, a projection device, a liquid crystal display (LCD), a plasma display panel, a light emitting diode (LED) array, such as an

organic LED (OLED) display, a cathode ray tube (CRT) display, electronic paper, a surface-conduction electron-emitted display (SED), a laser television display, a nanocrystal display or another type of display unit. Display 18 may be integrated within computing device 2. For instance, display 18 may be a screen of a mobile telephone or tablet computer. Alternatively, display 18 may be a stand-alone device coupled to computing device 2 via a wired or wireless communications link. For instance, display 18 may be a computer monitor or flat panel display connected to a personal computer via a cable or wireless link.

[0045] FIG. 3 is a block diagram illustrating example implementations of CPU 6, GPU 12, and system memory 10 of FIG. 1 in further detail. CPU 6 may include at least one software application 24, a graphics API 26, and a GPU driver 7, each of which may be one or more software applications or services that execute on CPU 6. GPU 12 may include a 3D graphics processing pipeline 30 that includes a plurality of graphics processing stages that operate together to execute graphics processing commands. GPU 12 may be configured to execute graphics processing pipeline 30 in a variety of rendering modes, including a binning rendering mode (also called a tile-based or deferred rendering mode) and a direct rendering mode. GPU 12 may also be operable to execute a general purpose shader 39 for performing more general computations applicable to be executed by the highly parallel nature of GPU hardware. Such general-purpose applications may be a so-called general-purpose graphics processing unit (GPGPU) and may conform to a general-purpose API, such as OpenCL.

[0046] As shown in FIG. 3, graphics processing pipeline 30 may include a command engine 32, a geometry processing stage 34, a rasterization stage 36, and a pixel processing pipeline 38. Each of the components in graphics processing pipeline 30 may be implemented as fixed-function components, programmable components (e.g., as part of a shader program executing on a programmable shader unit), or as a combination of fixed-function and programmable components. Memory available to CPU 6 and GPU 12 may include system memory 10, that may itself include frame buffer 15. Frame buffer 15 may store rendered image data.

[0047] Software application 24 may be any application that utilizes the functionality of GPU 12. For example, software application 24 may be a GUI application, an operating system, a portable mapping application, a computer-aided design program for engineering or artistic applications, a video game application, or another type of software application that uses 2D or 3D graphics. Software application 24 may also be

an application that uses the GPU to perform more general calculations, such as in a GPGPU application.

[0048] Software application 24 may include one or more drawing instructions that instruct GPU 12 to render a graphical user interface (GUI) and/or a graphics scene. For example, the drawing instructions may include instructions that define a set of one or more graphics primitives to be rendered by GPU 12. In some examples, the drawing instructions may, collectively, define all or part of a plurality of windowing surfaces used in a GUI. In additional examples, the drawing instructions may, collectively, define all or part of a graphics scene that includes one or more graphics objects within a model space or world space defined by the application.

[0049] Software application 24 may invoke GPU driver 7, via graphics API 26, to issue one or more commands to GPU 12 for rendering one or more graphics primitives into displayable graphics images. For example, software application 24 may invoke GPU driver 7, via graphics API 26, to provide primitive definitions to GPU 12. In some instances, the primitive definitions may be provided to GPU 12 in the form of a list of drawing primitives, e.g., triangles, rectangles, triangle fans, triangle strips, etc. The primitive definitions may include vertex specifications that specify one or more vertices associated with the primitives to be rendered. The vertex specifications may include positional coordinates for each vertex and, in some instances, other attributes associated with the vertex, such as, e.g., color coordinates, normal vectors, and texture coordinates. The primitive definitions may also include primitive type information (e.g., triangle, rectangle, triangle fan, triangle strip, etc.), scaling information, rotation information, and the like. Based on the instructions issued by software application 24 to GPU driver 7, GPU driver 7 may formulate one or more commands that specify one or more operations for GPU 12 to perform in order to render the primitive. When GPU 12 receives a command from CPU 6, graphics processing pipeline 30 decodes the command and configures one or more processing elements within graphics processing pipeline 30 to perform the operation specified in the command. After performing the specified operations, graphics processing pipeline 30 outputs the rendered data to frame buffer 15 associated with a display device. Graphics processing pipeline 30 may be configured to execute in one of a plurality of different rendering modes, including a binning rendering mode and a direct rendering mode.

[0050] GPU driver 7 may be further configured to compile one or more shader programs, and to download the compiled shader programs onto one or more

programmable shader units contained within GPU 12. The shader programs may be written in a high level shading language, such as, e.g., an OpenGL Shading Language (GLSL), a High Level Shading Language (HLSL), a C for Graphics (Cg) shading language, etc. The compiled shader programs may include one or more instructions that control the operation of a programmable shader unit within GPU 12. For example, the shader programs may include vertex shader programs and/or pixel shader programs. A vertex shader program may control the execution of a programmable vertex shader unit or a unified shader unit, and include instructions that specify one or more per-vertex operations. A pixel shader program may include pixel shader programs that control the execution of a programmable pixel shader unit or a unified shader unit, and include instructions that specify one or more per-pixel operations. In accordance with some examples of this disclosure, a pixel shader program may also include instructions that selectively cause texture values to be retrieved for source pixels based on corresponding destination alpha values for the source pixels.

[0051] Graphics processing pipeline 30 may be configured to receive one or more graphics processing commands from CPU 6, via graphics driver 7, and to execute the graphics processing commands to generate displayable graphics images. As discussed above, graphics processing pipeline 30 includes a plurality of stages that operate together to execute graphics processing commands. It should be noted, however, that such stages need not necessarily be implemented in separate hardware blocks. For example, portions of geometry processing stage 34 and pixel processing pipeline 38 may be implemented as part of a unified shader unit. Again, graphics processing pipeline 30 may be configured to execute in one of a plurality of different rendering modes, including a binning rendering mode and a direct rendering mode.

[0052] Command engine 32 may receive graphics processing commands and configure the remaining processing stages within graphics processing pipeline 30 to perform various operations for carrying out the graphics processing commands. The graphics processing commands may include, for example, drawing commands and graphics state commands. The drawing commands may include vertex specification commands that specify positional coordinates for one or more vertices and, in some instances, other attribute values associated with each of the vertices, such as, e.g., color coordinates, normal vectors, texture coordinates and fog coordinates. The graphics state commands may include primitive type commands, transformation commands, lighting commands, etc. The primitive type commands may specify the type of primitive to be rendered

and/or how the vertices are combined to form a primitive. The transformation commands may specify the types of transformations to perform on the vertices. The lighting commands may specify the type, direction and/or placement of different lights within a graphics scene. Command engine 32 may cause geometry processing stage 34 to perform geometry processing with respect to vertices and/or primitives associated with one or more received commands.

[0053] Geometry processing stage 34 may perform per-vertex operations and/or primitive setup operations on one or more vertices in order to generate primitive data for rasterization stage 36. Each vertex may be associated with a set of attributes, such as, e.g., positional coordinates, color values, a normal vector, and texture coordinates. Geometry processing stage 34 modifies one or more of these attributes according to various per-vertex operations. For example, geometry processing stage 34 may perform one or more transformations on vertex positional coordinates to produce modified vertex positional coordinates. Geometry processing stage 34 may, for example, apply one or more of a modeling transformation, a viewing transformation, a projection transformation, a ModelView transformation, a ModelViewProjection transformation, a viewport transformation and a depth range scaling transformation to the vertex positional coordinates to generate the modified vertex positional coordinates. In some instances, the vertex positional coordinates may be model space coordinates, and the modified vertex positional coordinates may be screen space coordinates. The screen space coordinates may be obtained after the application of the modeling, viewing, projection and viewport transformations. In some instances, geometry processing stage 34 may also perform per-vertex lighting operations on the vertices to generate modified color coordinates for the vertices. Geometry processing stage 34 may also perform other operations including, e.g., normal transformations, normal normalization operations, view volume clipping, homogenous division and/or backface culling operations.

[0054] Geometry processing stage 34 may produce primitive data that includes a set of one or more modified vertices that define a primitive to be rasterized as well as data that specifies how the vertices combine to form a primitive. Each of the modified vertices may include, for example, modified vertex positional coordinates and processed vertex attribute values associated with the vertex. The primitive data may collectively correspond to a primitive to be rasterized by further stages of graphics processing pipeline 30. Conceptually, each vertex may correspond to a corner of a primitive where

two edges of the primitive meet. Geometry processing stage 34 may provide the primitive data to rasterization stage 36 for further processing.

[0055] In some examples, all or part of geometry processing stage 34 may be implemented by one or more shader programs executing on one or more shader units. For example, geometry processing stage 34 may be implemented, in such examples, by a vertex shader, a geometry shader or any combination thereof. In other examples, geometry processing stage 34 may be implemented as a fixed-function hardware processing pipeline or as a combination of fixed-function hardware and one or more shader programs executing on one or more shader units.

[0056] Rasterization stage 36 is configured to receive, from geometry processing stage 34, primitive data that represents a primitive to be rasterized, and to rasterize the primitive to generate a plurality of source pixels that correspond to the rasterized primitive. In some examples, rasterization stage 36 may determine which screen pixel locations are covered by the primitive to be rasterized, and generate a source pixel for each screen pixel location determined to be covered by the primitive. Rasterization stage 36 may determine which screen pixel locations are covered by a primitive by using techniques known to those of skill in the art, such as, e.g., an edge-walking technique, evaluating edge equations, etc. Rasterization stage 36 may provide the resulting source pixels to pixel processing pipeline 38 for further processing.

[0057] The source pixels generated by rasterization stage 36 may correspond to a screen pixel location, e.g., a destination pixel, and be associated with one or more color attributes. All of the source pixels generated for a specific rasterized primitive may be said to be associated with the rasterized primitive. The pixels that are determined by rasterization stage 36 to be covered by a primitive may conceptually include pixels that represent the vertices of the primitive, pixels that represent the edges of the primitive and pixels that represent the interior of the primitive.

[0058] Pixel processing pipeline 38 is configured to receive a source pixel associated with a rasterized primitive, and to perform one or more per-pixel operations on the source pixel. Per-pixel operations that may be performed by pixel processing pipeline 38 include, e.g., alpha test, texture mapping, color computation, pixel shading, per-pixel lighting, fog processing, blending, a pixel ownership test, a source alpha test, a stencil test, a depth test, a scissors test and/or stippling operations. In addition, pixel processing pipeline 38 may execute one or more pixel shader programs to perform one or more per-pixel operations. The resulting data produced by pixel processing pipeline

38 may be referred to herein as destination pixel data and stored in frame buffer 15. The destination pixel data may be associated with a destination pixel in frame buffer 15 that has the same display location as the source pixel that was processed. The destination pixel data may include data such as, e.g., color values, destination alpha values, depth values, etc.

[0059] Frame buffer 15 stores destination pixels for GPU 12. Each destination pixel may be associated with a unique screen pixel location. In some examples, frame buffer 15 may store color components and a destination alpha value for each destination pixel. For example, frame buffer 15 may store Red, Green, Blue, Alpha (RGBA) components for each pixel where the “RGB” components correspond to color values and the “A” component corresponds to a destination alpha value. Pixel values may also be represented by a luma component (Y) and one or more chroma components (e.g., U and V). Although frame buffer 15 and system memory 10 are illustrated as being separate memory units, in other examples, frame buffer 15 may be part of system memory 10.

[0060] General purpose shader 39 may be any application executable on GPU 12 to perform calculations. Typically, such calculations are of the type that takes advantage of the highly parallel structure of GPU processing cores, including arithmetic logic units (ALUs). An example general purpose shader 39 may conform to the OpenCL API. OpenCL is an API that allows an application to have access across multiple processors in a heterogeneous system (e.g., a system including a CPU, GPU, DSP, etc.). Typically, in an OpenCL conforming application, GPU 12 would be used to perform non-graphical computing. Examples of non-graphical computing applications may include physics-based simulations, fast Fourier transforms, audio signal processing, digital image processing, video processing, image post filtering, computational camera, climate research, weather forecasting, neural networks, cryptography, and massively parallel data crunching, among many others.

[0061] FIG. 4 is a block diagram showing an example device configured to implement the hardware enforced content protection techniques of this disclosure. In the example of FIG. 4, GPU 12 may be configured to operate according to a secure mode or an unsecure mode. In one example of the disclosure, in secure mode, GPU 12 is restricted from writing output data (e.g., game data, video, etc.) to unsecure memory 56. Rather, in secure mode, GPU 12 may only write output data to secure memory 57. While in secure mode, GPU 12 may read data from either secure memory 57 or unsecure memory 56. In unsecure mode, for this example, GPU 12 is restricted from reading any data

from secure memory 57. Rather, in unsecure mode, GPU 12 may only read data from unsecure memory 56. Likewise, while in unsecure mode, GPU 12 may only write data to unsecure memory 56.

[0062] Unsecure memory 56 and secure memory 57 may be any type of memory, including one or more volatile or non-volatile memories or storage devices. Example memory and storage devices include RAM, SRAM, DRAM, ROM, EPROM, EEPROM, Flash memory, magnetic data media or optical storage media. Secure memory 57 includes additional features not found in unsecure memory 56. For example, secure memory 57 may utilize encryption, authentication and/or other digital rights management techniques to prevent access to, copying of, or deciphering of data stored thereon. In general, secure memory 57 may be considered to be one portion of system memory 10 and unsecure memory 56 may be considered to be another portion of system memory 10.

[0063] In accordance with one or more examples of the disclosure described below, GPU 12 may be configured to control or otherwise affect where data is read from and written to using memory access controller 53. Memory access controller 53 is responsive to the mode GPU 12 is operating under (i.e., secure mode or unsecure mode), and makes read/write decisions based on the mode. In general, memory access controller 53 may be configured to impose a restriction on the nature of the transactions between GPU 12 and memory controller 50. Memory access controller 53 may be configured to be responsive to the mode (i.e., the secure mode or unsecure mode) in which GPU 12 is currently operating, and may impose restrictions to memory transactions in accordance with the examples of the disclosure below.

[0064] In one example of the disclosure, the GPU memory mode (e.g., secure mode or unsecure mode) is set by GPU driver 7 operating on CPU 6. GPU driver 7 may change the memory mode in GPU 12 in several different ways. In one example, GPU driver 7 may directly write a value into a register in GPU 12 that indicates to GPU 12 which memory mode to use (e.g., secure mode or unsecure mode). In another example, GPU 12 may include one or more instructions in a command stream executable by GPU 12 that instruct GPU 12 itself to write a certain value to a register that indicates which memory mode to use. In this way, GPU driver 7 may only select the memory mode that the GPU operates under, and does not make any direct instructions that specifies which data is to be written to which memory. As such, even if GPU driver 7 were altered to place GPU 12 in an unsecure mode, through the function of memory access controller

53, GPU 12 would prevent any read access from secure memory 57, as memory access controller 53 is only able to read from unsecure memory 56 in the unsecure mode.

Likewise, even if GPU driver 7 were altered to place GPU 12 in a secure mode, through the function of memory access controller 53, GPU 12 would prevent any write access to unsecure memory 56, as memory access controller 53 is only able to write to secure memory 57 in the secure mode. As such, the techniques of this disclosure may still prevent copying of data to unsecure memory 56, even if GPU driver 7 were altered to place GPU 12 in a secure mode.

[0065] In one example of the disclosure, memory access controller 53 is configured to access secure memory 57 and unsecure memory 56 via secure and unsecure memory management unit (MMU) page tables, respectively. In this example, virtual address ranges are provided to GPU 12 by GPU driver 7. The virtual address ranges include a range of virtual addresses for the secure memory and a range of virtual addresses for the unsecure memory. When placed in secure mode by GPU driver 7, GPU 12 utilizes the range of virtual addresses for the secure memory to perform reads and writes. GPU 12 would also be able to use the range of virtual addresses for the unsecure memory to perform reads in the secure mode, but not to perform writes, thereby preventing unauthorized copying of protected data from the secure memory. When placed in unsecure mode by GPU driver 7, GPU 12 would utilize the range of virtual addresses for the unsecure memory to perform reads and writes.

[0066] In one example, memory access controller 53 routes reads and writes to the appropriate memory units (e.g., secure memory 57 or unsecure memory 56) by determining if the virtual address used in the read or write request are within an unsecure range of virtual memory addresses or within a secure range of virtual addresses. Based on the range determination, memory access controller utilizes one of unsecure IOMMU 51 or secure IOMMU 52 in memory controller 50. Memory controller 50 is configured to facilitate the transfer of data going into and out of system memory 10. To effectively handle any such transaction, memory controller 50 may include one or more MMUs for controlling device access, such as GPU 12, to system memory 10. Unsecure IOMMU 51 and secure IOMMU 52 contain mappings for virtualized memory addresses providing a continuous view of pages for its client. In this example, a client may be any entity that binds or provides one or more resources to GPU 12 (e.g., an application executed by GPU 12 or an application executing on CPU 6). A resource is a container of information (e.g., a memory or buffer) for GPU 12 to

use in some manner. In some examples, a resource may have a descriptor that provides information about how the memory is to be used.

[0067] In one example of the disclosure, unsecure IOMMU 51 is an IOMMU that is configured to map virtual memory addresses to physical memory addresses in unsecure memory 56. Secure IOMMU 52 is an IOMMU that is configured to map virtual memory addresses to physical memory addresses in secure memory 57. Unsecure IOMMU 51 performs the mappings to unsecure memory 56 using an unsecure page table. The unsecure page table is a page table that maps a range of virtual memory addresses (e.g., the range provided by GPU driver 7) to locations in unsecure memory 56. Likewise, secure IOMMU 52 performs the mappings to secure memory 57 using a secure page table. The secure page table is a page table that maps a range of virtual memory addresses (e.g., the range provided by GPU driver 7) to locations in secure memory 57. As depicted in FIG. 4, unsecure IOMMU 51 and secure IOMMU 52 are part of a single memory controller 50. Memory controller 50 may be one of memory controllers 8 depicted in FIG. 1. In effect, memory controller 50 becomes a secure IOMMU when it is operating with a secure page table, and becomes an unsecure IOMMU when it is operating with an unsecure page table. In other examples, unsecure IOMMU 51 and secure IOMMU 52 may be physically separate MMUs.

[0068] In one example of the disclosure, both secure and unsecure page tables are provided to secure IOMMU 52 and unsecure IOMMU 51 by secure operating system (OS) 54 executing on CPU 6. A secure OS is an OS that operates alongside a normal “rich” OS (e.g., Apple iOS, Google Android, Microsoft Windows, etc.). The secure OS provides security applications to protect and separate a secure kernel and any secure peripherals (e.g., secure IOMMU 52) from any code running on the rich OS (e.g., GPU driver 7). An example of a secure OS is the TrustZone software made by ARM Holdings. In general, a secure OS is considered to be much less susceptible to alteration and attack than software running on a rich OS, including software such as graphics drivers. In accordance with the techniques of this disclosure, only the secure OS is allowed to update the page tables for mapping virtual memory address ranges to physical memory addresses. As such, any attempt to alter the graphics driver, including the virtual address ranges provided by the driver, will not result in secure content being stored in unsecure memory, as only the secure OS provides the ultimate mappings to secure and unsecure memory.

[0069] In the example where both the secure and unsecure page tables are available at memory controller 50 (e.g., memory controller 50 includes both unsecure IOMMU 51 and secure IOMMU 52), GPU 12 is able to read data from both unsecure memory 56 and secure memory 57 in secure mode. The other read/write restrictions still apply. That is, in secure mode, writes are only made to secure memory 57 by GPU 12, and in unsecure mode both reads and writes by GPU 12 are limited to unsecure memory 56.

[0070] In another example of the disclosure, rather than having both a secure and unsecure IOMMU available to the GPU, where data traffic is directed to either the secure or unsecure IOMMU via memory access controller 53, only one IOMMU (i.e., either unsecure IOMMU 51 or secure IOMMU 52) would be made available to GPU 12 depending on the selected memory mode. That is, if the memory mode is the unsecure mode, secure OS 54 only provides page table mappings for the unsecure IOMMU 51. In this situation, the secure IOMMU 52 would be unavailable. If the memory mode is the secure mode, secure OS 54 only provides page table mappings for the secure IOMMU 52. In this situation, the unsecure IOMMU 51 would be unavailable. This example of only having one IOMMU available per memory mode would provide a more simple implementation where both reads and writes were restricted per memory mode. That is, only reads and writes to secure memory 57 by GPU 12 would be allowed in secure mode, while only reads and writes to unsecure memory 56 by GPU 12 would be allowed in unsecure mode. This differs slightly from the approach described above where both IOMMUs may be available, in that the secure mode would no longer allow for reads for unsecure memory 56.

[0071] Even when in secure mode, there are some writes, other than the ultimate output product of GPU 12, which would be better for GPU to write to unsecure memory. These writes include the communication tokens between GPU 12 and graphics driver 7. Such data includes timestamps and other ancillary data and control data, such as counter data and query data. GPU 12 uses memory (e.g., unsecure memory 56) to communicate such timestamps and data back to the driver. Since the graphics driver 7 is untrusted, the memory involved in the communication path needs to be unsecure (e.g., unsecure memory 56). As one example, when GPU 12 reaches a certain point in processing, GPU 12 writes a timestamp/sequential marker to memory. Graphics driver 7 uses this information to determine how far the GPU has proceeded in a specific command stream. This determination, for example, allows graphics driver 7 to release memory objects that GPU 12 is operating on, once GPU 12 finishes. There are many other types of

signaling and communication paths GPU 12 may use for providing information to graphics driver 7. As another example, graphics driver 7 can request GPU 12 to report performance counters after a drawcall. GPU 12 then writes these performance counters to a memory location (e.g., in unsecure memory 56) specified by graphics driver 7.

[0072] To solve this exception to the general rule above that GPU 12 does not write to unsecure memory in secure mode, GPU 12 hardware may be modified such that certain hardware blocks are configured to have unsecure memory accesses, while also not having access to data paths and caches that connect to or contain secure content when the GPU is running in secure mode.

[0073] FIG. 5 depicts an example implementation where certain hardware blocks of GPU 12 only have direct access to unsecure memory through the memory interface block (VBIF 60) of GPU 12, and then through unsecure IOMMU 51, even when GPU 12 is in secure mode. One example of such a hardware block is a command processor (CP) 62 block at the front end of the GPU. CP 62 may execute a command engine, such as command engine 32 shown in FIG. 3. CP 62 is responsible for sending messages (via unsecure memory) back to the GPU driver 7. As shown in FIG. 5, CP 62 is configured to have only one physical path to memory (in this case, unsecure memory), through unsecure IOMMU 51. As such, regardless of whether any other hardware blocks of GPU 12 are operating on secure content, CP 62 may never gain access to such secure content. To further ensure that CP 62 has no access to secure content, CP 62 may also be physically isolated from (e.g., have no connections to) any registers that may be used to store secure content, including a debug bus. As shown in FIG. 5, CP 62 has no direct access to L2 cache 61 and graphics memory (GMEM) 70. GMEM 70 is fast memory (often SRAM) that GPU 12 uses as a render target or framebuffer when rendering content for display in some operational modes of GPU 12. L2 cache 61 is a secondary cache that is used to store recently addressed data or frequently used data so that the number of accesses to main memory (e.g., secure memory) may be reduced. L2 cache 61 may also be used to buffer program instructions. Typically, L2 cache 61 is larger than GMEM 70.

[0074] Other hardware blocks of GPU 12 may also be configured to only have access to unsecure memory. For example, a primitive control (PC) unit and a visibility stream compressor (VSC) may be configured to only have access to unsecure memory. A PC unit controls how a primitive (e.g., a triangle) progresses or “walks” through a graphics pipeline (e.g., graphics 3D processing pipeline 30 of FIG. 3). A VSC is used in a tile-

based or deferred rendering scheme to compress and manage a visibility stream. In general, it may be beneficial, in some circumstances, to avoid requiring certain hardware blocks to write to secure memory. Such circumstances include situations where hardware blocks are not writing secure content, and when hardware blocks are writing control data needed by a graphics driver.

[0075] Other hardware blocks in FIG. 5, store content to unsecure memory or secure memory based on the techniques described above. That is, in unsecure mode, data may only be read from or written to unsecure memory. No data may be read from secure memory in unsecure mode. In secure mode, data may only be written to secure memory. No data may be written to unsecure memory in secure mode. However, in secure mode in some examples, data may be read from both secure memory and unsecure memory. These additional hardware blocks of GPU 12 that may access memory according to the memory mode include vertex fetch decode (VFD) unit 65, high level sequencer (HLSQ) 66, vertex shader (VS) 67, pixel shader (PS) 68, and render backend (RB) 69. VFD 65 is responsible for fetching vertex data at the request of CP 62. HLSQ 66 controls the shader processors (i.e., the programmable processors on the GPU that execute shader code), populating the correct state for the job being executed and launching jobs into the shader processor. VS 67 is a vertex shader executing on the shader processor. For example, VS 67 may include vertex shader code that executes geometry processing stage 34 of graphics 3D processing pipeline 30 of FIG. 3. PS 68 is a pixel shader executing on the shader processor. For example, PS 68 may include pixel shader code that executes pixel processing pipeline 38 of graphics 3D processing pipeline 30 of FIG. 3. Render backend (RB) 69 is responsible for writing and reading pixels for the depth buffer and stencil buffer.

[0076] FIG. 6 is a block diagram showing another example structure configured to perform the hardware enforced content protection techniques of this disclosure. In the example of FIG. 6, GPU 12 and memory controller 50 are the same as that described above in FIG. 5, except for the operation of memory access controller 53. In addition, for simplification, various hardware units present in GPU 12 have been generally labeled as such as GPU hardware blocks 71. GPU hardware blocks 71 may include one or more of VFD unit 65, HLSQ 66, VS 67, PS 68, and RB 69.

[0077] In the example of FIG. 6, memory access controller 53 may be configured to direct data into a memory unit (e.g., unsecure memory 56 or secure memory 57) based on the memory mode (i.e., unsecure mode or secure mode) of GPU 12 and a resource

descriptor associated with a memory resource (e.g., a buffer or cache line storing data, as shown in client 73 of FIG. 6). A resource descriptor, e.g., called a “secure tag,” may be associated with each resource to indicate whether the data for the resource should be routed according to the secure mode (e.g., through secure IOMMU 52) or routed according to the unsecure mode (e.g., through unsecure IOMMU 51). As shown in FIG. 6, the resource descriptor may indicate a trusted “T” resource that would use secure IOMMU 52, and an untrusted “U” resource that would use unsecure IOMMU 51.

[0078] Using the resource descriptor and the memory mode of GPU 12, memory access controller 53 may be configured to direct memory reads and writes from GPU hardware blocks 71 through L2 cache 61 based on the resource descriptor. Each cache line in L2 cache 61 may include resource descriptor information. In one example, memory access controller 53 may be configured to examine the secure tag information present in the resource for a particular memory transaction (e.g., a read or a write) and determine which of unsecure IOMMU 51 or secure IOMMU 52 to use to route the transaction.

[0079] For example, when GPU 12 is set to operate in the secure mode, memory access controller 53 may examine secure tag information in the resource descriptor for the resources in the memory transaction. If the secure tag indicates a trusted resource “T,” memory access controller 53 will direct both reads and writes of such secure resources to secure IOMMU 52. In some examples, memory access controller will direct all read and writes of resource having the T resource descriptor to secure IOMMU 52. If the secure tag information indicates an untrusted resource “U,” memory access controller 53 will direct (e.g., some or all) reads of such unsecure resources to unsecure IOMMU 51, but will drop or not allow requests (e.g., some or all) for writes of the unsecure resource.

[0080] In accordance with the above examples, GPU 12 may be configured to access a first memory unit (e.g., system memory 10) according to one of an unsecure mode and a secure mode and a respective resource descriptor associated with each of a plurality of memory resources. Memory access controller 53 may be configured to read the resource descriptor of the plurality of memory resources and receive a request for a memory transaction to the first memory unit.

[0081] Memory access controller 53 may be further configured to, in response to the request, direct read and write memory transactions relating to memory resources of the plurality of memory resources having a secure resource descriptor to a secure portion of the first memory unit when GPU 12 is operating according to the secure mode.

Memory access controller 53 may be further configured to, in response to the request, direct read memory transactions relating to memory resources of the plurality of memory resources having an unsecure resource descriptor to an unsecure portion of the first memory unit when GPU 12 is operating according to the secure mode. Memory access controller 53 may be further configured to, in response to the request, drop write memory transactions relating to memory resources of the plurality of memory resources having the unsecure resource descriptor when the GPU is operating according to the secure mode.

[0082] In another example of the disclosure, memory access controller 53 may be further configured to, in response to the request, direct read and write memory transactions relating to memory resources of the plurality of memory resources having the unsecure resource descriptor to an unsecure portion of the first memory unit when the GPU is operating according to the unsecure mode. Memory access controller 53 may be further configured to, in response to the request, drop read and write memory transactions relating to memory resources of the plurality of memory resources having the secure resource descriptor when the GPU is operating according to the unsecure mode.

[0083] FIGS. 7A and 7B are block diagrams showing other example structures configured to perform the hardware enforced content protection techniques of this disclosure. In the examples of FIGS. 7A and 7B, memory controller 100 may include one or more MMUs. As described above, an MMU implements a virtualized memory scheme providing a continuous view of pages for its client. The virtual memory space may be divided into virtual pages. The MMUs may implement one or more context banks to maintain these virtual page tables. The context banks may include both page table (PT) entries that map virtual memory addresses to physical memory addresses, as well as rules that indicate whether reads, writes, or both reads and writes are allowed for the particular PT entries in each context bank.

[0084] In the example of FIGS. 7A and 7B, an MMU of memory controller 100 may include an unsecure context bank 102 and a secure context bank 105. Unsecure context bank 102 may include unsecure PT entries 104 that are mapped for read-only access. Unsecure PT entries 104 may include mappings from virtual memory addresses to physical memory addresses in unsecure memory 56. Since unsecure PT entries 104 are mapped for read-only access, memory controller 100, using unsecure context bank 102, is only able to read from unsecure memory 56. Secure context bank 105 may include

unsecure PT entries 106 that are mapped for read-only access, and secure PT entries 108 that are mapped for both read and write access (R/W). Unsecure PT entries 106 may include mappings from virtual memory addresses to physical memory addresses in unsecure memory 56. Secure PT entries 106 may include mappings from virtual memory addresses to physical memory addresses in secure memory 57.

[0085] When GPU 12 is placed in secure mode using one of the techniques discussed above, memory access controller 53 may be configured to direct memory transactions for GPU hardware blocks 71 to secure context bank 105 of memory controller 100. If GPU 12, or an instruction from a client using GPU 12, attempts to perform a write into an unsecure resource (e.g., unsecure memory 56), memory controller 100 is configured to issue a page fault since the PT entries in unsecure context bank 102 are mapped as read only in secure context bank 105. The page fault indicates to the client that such a memory transaction is not allowed.

[0086] In one example of the disclosure, CP 62 may be configured to always operate in unsecure mode irrespective of the memory mode of GPU 12. That is, CP 62 may be configured to always use unsecure context bank 102. FIG. 7A shows the flow of memory transactions from GPU 12 to unsecure context bank 102 and secure context bank 105 when GPU 12 is in secure mode. FIG. 7B shows the flow of memory transactions from GPU 12 to unsecure context bank 102 and secure context bank 105 when GPU 12 is in unsecure mode.

[0087] To reiterate, GPU 12 may be configured to access a memory (e.g., unsecure memory 56 or secure memory 57) according to one of an unsecure mode and a secure mode. GPU 12 may include memory access controller 53 that is configured to direct memory transactions from at least one hardware unit (e.g., one or more of GPU hardware blocks 71) of the GPU 12 to secure context bank 105 in memory controller 100 when GPU 12 is operating in a secure mode. Memory access controller 53 may also be configured to direct memory transactions from the at least one hardware unit of GPU 12 to unsecure context bank 102 in memory controller 100 when GPU 12 is operating in the unsecure mode.

[0088] As described above, secure context bank 105 may include read-only page table entries to the unsecure portion of the memory (e.g., unsecure memory 56) and read/write page table entries to the secure portion of the memory (e.g., secure memory 57). Unsecure context bank 102 may include read-only page table entries to the unsecure portion of the memory (e.g., unsecure memory 56). In one example, memory

controller 100 may be configured to issue a page fault when a request to write data into an address contained within the read-only page table entries of secure context bank 105 is received.

[0089] In any of the examples described above, when GPU 12 transitions from secure mode to unsecure mode, there may be secure content remaining within various caches, memories and registers of GPU 12. In one example of the disclosure, a mechanism is provided to clear and/or invalidate the various storage units of GPU 12 that may hold secure content before allowing an unsecure job using the unsecure memory mode to launch on GPU 12.

[0090] In this context, clearing a memory means that data stored in the memory is erased and/or allowed to be overwritten. In practice, clearing may involve de-allocating all memory addresses for the memory unit such that all data in the memory unit may be overwritten. In other examples, clearing may involve overwriting all data in the memory unit (e.g., with all 1's or all 0's), such that any previously stored data is no longer available. If a memory unit is not cleared, an unsecure job could copy the trailing remains of secure data to unsecure memory. This problem can be solved via secure software techniques, hardware techniques, or a combination of both techniques. Regardless, the clearing and transition to unsecure may be an atomic operation, since this operation is triggered by the unsecure driver. In this context, an atomic operation includes the clearing of internal GPU 12 memories together (i.e., atomically) with the transition back to unsecure mode. For example, there must be a single “command” that does both (changing modes and clearing internal memories); otherwise malicious software could just perform the transition back to an unsecure mode, and not execute the clearing operation.

[0091] In some examples, it may not be necessary to clear all storage units of GPU 12 when transitioning from secure mode to unsecure mode. Instead, only a portion of the storage units need to be cleared to effectively prevent unauthorized access to secure content. As one example, only half the content stored may be cleared. As another example, every other chunk of data (e.g., every other 32 bytes of data) may be cleared.

[0092] FIG. 8 is a block diagram showing cache clearing techniques according to one example of this disclosure. In the example of FIG. 8, a secure software solution is used to transition the GPU between secure and unsecure modes. In one example, a GPU register (e.g., clear register 74) is under the control of secure software (e.g., secure OS 54) running on the host CPU 6. If the GPU driver 7 switches the memory mode of GPU

12 from unsecure mode to secure mode, GPU driver 7 would also call secure software in secure OS 54 to clear any secure content remaining on caches, memories or registers of GPU 12, including L2 cache 61, GMEM 70, and other registers 72. At that point, the secure OS 54 could first launch a job on the GPU 12 by writing memory clear and/or invalidate instructions into clear register 74. Such instructions would result in all of the remaining secure data in GPU 12 being cleared. Such instructions may be a combination of a shader program, memory writes and/or register programmings (e.g., a GPU L2 cache invalidate).

[0093] FIG. 9 is a block diagram showing cache clearing techniques according to another example of this disclosure. In the example of FIG. 9, a hardware solution is used to transition the GPU 12 between secure and unsecure modes. In this example, an externally visible (e.g., memory-mapped input/output (MMIO)) or internal (e.g., command stream) register 76 may be configured so that it is directly written to by the graphics driver 7. Hardware of GPU 12 may be configured such that, when register 76 is written (e.g., when going from secure mode to unsecure mode), the hardware of GPU 12 completes the current secure job, drains the pipeline (i.e., removes any remaining secure content being processed), and clears and/or invalidates all registers, memories, and caches that could contain secure content, including L2 cache 61, GMEM 70, and other registers 72. This clearing process may include using hardwired or securely loaded and protected shader code resident on GPU 12.

[0094] FIG. 10 is a flowchart illustrating a method according to one example of the disclosure. GPU 12, including memory access controller 53, and memory controller 100 may be configured to perform the techniques of FIG. 10.

[0095] In one example of the disclosure, memory access controller 53 may be configured to access an unsecure portion of a memory (e.g., system memory 10) according to an unsecure mode by directing memory transactions from at least one hardware unit of GPU 12 to an unsecure context bank in memory controller 100 (202). Memory access controller 53 may be further configured to access a secure portion of the memory according to a secure mode by directing memory transactions from the at least one hardware unit of GPU 12 to a secure context bank in memory controller 100 (204). In one example of the disclosure, the secure context bank includes read-only page table entries to the unsecure portion of the memory and read/write page table entries to the secure portion of the memory, and the unsecure context bank includes read-only page table entries to the unsecure portion of the memory. In another example of the

disclosure, memory controller 100 may be configured to issue a page fault (208) when a request to write data into an address contained within the read-only page table entries of the secure context bank is received (206).

[0096] In another example of the disclosure, the at least one hardware unit of the GPU includes one or more of a vertex fetch decode unit, a high level sequencer, a vertex shader, a pixel shader, and a render backend unit.

[0097] In another example of the disclosure, GPU 12 may be configured to access the unsecure portion of the memory, with a front end command processor, through the unsecure context bank regardless of whether GPU 12 is operating in the unsecure mode or the secure mode.

[0098] In another example of the disclosure, GPU driver 7 may be configured to place GPU 12 in the secure mode or the unsecure mode. In a further example of the disclosure, GPU 12 may be configured to receive, from GPU driver 7, an instruction to a command stream register of GPU 12 to clear and invalidate one or more internal memories of GPU 12. GPU 12 may be further configured to clear and invalidate at least some content from the one or more internal memories of GPU 12 when GPU 12 is transitioned from the secure mode to the unsecure mode based on the instruction in the command stream register.

[0099] In another example of the disclosure, GPU 12 may be configured to receive, at a clear register, an indication to clear and invalidate one or more internal memories of the GPU, and clear and invalidate at least some content from the one or more internal memories of GPU 12 when GPU 12 is transitioned from the secure mode to the unsecure mode based on the indication in the clear register.

[0100] FIG. 11 is a flowchart illustrating a method according to one example of the disclosure. GPU 12, including memory access controller 53, may be configured to perform the techniques of FIG. 11.

[0101] In one example of the disclosure, GPU 12 may be configured to access a first memory unit (e.g., system memory 10) according to one of an unsecure mode and a secure mode and a respective resource descriptor associated with each of a plurality of memory resources. Memory access controller 53 may be configured to read the respective resource descriptor associated with each of the plurality of memory resources (302) and receive a request for a memory transaction to the first memory unit (304).

[0102] Memory access controller 53 may be further configured to determine whether the resource descriptor associated with the memory resource related to the request for

the memory transaction is a secure or unsecure resource descriptor (306). Memory access controller 53 may be further configured to direct, in response to the request, read and write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is a secure resource descriptor to a secure portion of the first memory unit when GPU 12 is operating according to the secure mode (312). Memory access controller 53 may be further configured to direct, in response to the request, read memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is an unsecure resource descriptor to an unsecure portion of the first memory unit when GPU 12 is operating according to the secure mode (308). Memory access controller 53 may also be configured to drop, in response to the request, write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is the unsecure resource descriptor when GPU 12 is operating according to the secure mode (310).

[0103] In another example of the disclosure, memory access controller 53 is further configured to, in response to the request, direct read and write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is the unsecure resource descriptor to an unsecure portion of the first memory unit when GPU 12 is operating according to the unsecure mode, and in response to the request, drop read and write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is the secure resource descriptor when GPU 12 is operating according to the unsecure mode.

[0104] In another example of the disclosure, memory access controller 53 is configured to write data to the secure portion of the first memory unit by utilizing a secure memory management unit, the secure memory management unit utilizing a secure page table containing address ranges for the secure portion of the first memory unit. In another example of the disclosure, memory access controller 53 is configured to read data from the unsecure portion of the first memory unit by utilizing an unsecure memory management unit, the unsecure memory management unit utilizing an unsecure page table containing address ranges for the unsecure portion of the first memory unit.

[0105] In another example of the disclosure, memory access controller 53 is configured to read and write data according to a virtual memory address from a range of virtual memory address, wherein the range of virtual memory addresses includes a first range

of virtual memory addresses relating to entries in the secure page table utilized by the secure memory management unit, and a second range of virtual memory addresses relating to entries in the unsecure page table utilized by the unsecure memory management unit.

[0106] In one or more examples, the functions described above may be implemented in hardware, software, firmware, or any combination thereof. If implemented in software, the functions may be stored as one or more instructions or code on an article of manufacture comprising a non-transitory computer-readable medium. Computer-readable media may include computer data storage media. Data storage media may be any available media that can be accessed by one or more computers or one or more processors to retrieve instructions, code and/or data structures for implementation of the techniques described in this disclosure. By way of example, and not limitation, such computer-readable media can comprise RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage, or other magnetic storage devices, flash memory, or any other medium that can be used to carry or store desired program code in the form of instructions or data structures and that can be accessed by a computer. Disk and disc, as used herein, includes compact disc (CD), laser disc, optical disc, digital versatile disc (DVD), floppy disk and Blu-ray disc where disks usually reproduce data magnetically, while discs reproduce data optically with lasers. Combinations of the above should also be included within the scope of computer-readable media.

[0107] The code may be executed by one or more processors, such as one or more DSPs, general purpose microprocessors, ASICs, FPGAs, or other equivalent integrated or discrete logic circuitry. In addition, in some aspects, the functionality described herein may be provided within dedicated hardware and/or software modules. Also, the techniques could be fully implemented in one or more circuits or logic elements.

[0108] The techniques of this disclosure may be implemented in a wide variety of devices or apparatuses, including a wireless handset, an integrated circuit (IC) or a set of ICs (e.g., a chip set). Various components, modules, or units are described in this disclosure to emphasize functional aspects of devices configured to perform the disclosed techniques, but do not necessarily require realization by different hardware units. Rather, as described above, various units may be combined in a codec hardware unit or provided by a collection of interoperative hardware units, including one or more processors as described above, in conjunction with suitable software and/or firmware.

[0109] Various examples have been described. These and other examples are within the scope of the following claims.

WHAT IS CLAIMED IS:

1. An apparatus for graphics processing comprising:
a graphics processing unit (GPU) configured to access a first memory unit according to one of an unsecure mode and a secure mode and a respective resource descriptor associated with each of a plurality of memory resources, the GPU comprising:
 - a memory access controller configured to read the respective resource descriptor associated with each of the plurality of memory resources,
 - the memory access controller configured to receive a request for a memory transaction to the first memory unit,
 - the memory access controller configured to, in response to the request, direct all read and write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is a secure resource descriptor to a secure portion of the first memory unit when the GPU is operating according to the secure mode,
 - the memory access controller configured to, in response to the request, direct all read memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is an unsecure resource descriptor to an unsecure portion of the first memory unit when the GPU is operating according to the secure mode, and
 - the memory access controller configured to, in response to the request, drop all write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is the unsecure resource descriptor when the GPU is operating according to the secure mode.

2. The apparatus of claim 1, wherein the memory access controller is further configured to, in response to the request, direct read and write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is the unsecure resource descriptor to an unsecure portion of the first memory unit when the GPU is operating according to the unsecure mode, and

the memory access controller is further configured to, in response to the request, drop read and write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is the secure resource descriptor when the GPU is operating according to the unsecure mode.

3. The apparatus of claim 1, wherein the memory access controller is configured to write data to the secure portion of the first memory unit utilizing a secure memory management unit, the secure memory management unit utilizing a secure page table containing address ranges for the secure portion of the first memory unit, and

wherein the memory access controller is configured to read data from the unsecure portion of the first memory unit utilizing an unsecure memory management unit, the unsecure memory management unit utilizing an unsecure page table containing address ranges for the unsecure portion of the first memory unit.

4. The apparatus of claim 3, wherein the memory access controller reads and writes data according to a virtual memory address from a range of virtual memory addresses, wherein the range of virtual memory addresses includes a first range of virtual memory addresses relating to entries in the secure page table utilized by the secure memory management unit, and a second range of virtual memory addresses relating to entries in the unsecure page table utilized by the unsecure memory management unit.

5. The apparatus of claim 4, further comprising:

a second memory unit storing a graphics driver, the graphics driver configured to place the GPU in a secure mode or an unsecure mode.

6. The apparatus of claim 5, further comprising:
the secure memory management unit;
the unsecure memory management unit; and
a central processing unit (CPU) executing a secure operating system and the graphics driver, the secure operating system configured to supply the secure page table to the secure memory management unit and the unsecure page table to the unsecure memory management unit.
7. The apparatus of claim 6, wherein the GPU further comprises a clear register and one or more internal memories, and wherein the secure operating system is configured to send an instruction to the clear register that causes the GPU to clear and invalidate at least some content from the one or more internal memories when the GPU is transitioned from the secure mode to the unsecure mode.
8. The apparatus of claim 6, wherein the GPU further comprises a command stream register and one or more internal memories, and wherein the graphics driver is configured to send an instruction to the command stream register that causes the GPU to clear and invalidate at least some content from the one or more internal memories when the GPU is transitioned from the secure mode to the unsecure mode.
9. The apparatus of claim 1, wherein the GPU further comprises:
one or more hardware blocks configured to write data to the unsecure portion of the first memory regardless of whether the GPU is in the unsecure mode or the secure mode, wherein the one or more hardware blocks do not have read access to the secure portion of the first memory unit.
10. The apparatus of claim 9, wherein the one or more hardware blocks includes a front end command processor.

11. A method comprising:

reading a respective resource descriptor for a respective memory resource of a plurality of memory resources;

receiving a request for a memory transaction to a first memory unit;

directing, in response to the request, read and write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is a secure resource descriptor to a secure portion of the first memory unit when a graphics processing unit (GPU) is operating according to a secure mode;

directing, in response to the request, read memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is an unsecure resource descriptor to an unsecure portion of the first memory unit when the GPU is operating according to the secure mode; and

dropping, in response to the request, write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is the unsecure resource descriptor when the GPU is operating according to the secure mode.

12. The method of claim 11, further comprising:

directing, in response to the request, read and write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is the unsecure resource descriptor to an unsecure portion of the first memory unit when the GPU is operating according to the unsecure mode; and

dropping, in response to the request, read and write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is the secure resource descriptor when the GPU is operating according to the unsecure mode.

13. The method of claim 11, further comprising:
 - writing data to the secure portion of the first memory unit utilizing a secure memory management unit, the secure memory management unit utilizing a secure page table containing address ranges for the secure portion of the first memory unit; and
 - reading data from the unsecure portion of the first memory unit utilizing an unsecure memory manage unit, the unsecure memory management unit utilizing an unsecure page table containing address ranges for the unsecure portion of the first memory unit.
14. The method of claim 13, further comprising:
 - reading and writing data according to a virtual memory address from a range of virtual memory addresses, wherein the range of virtual memory addresses includes a first range of virtual memory addresses relating to entries in the secure page table utilized by the secure memory management unit, and a second range of virtual memory addresses relating to entries in the unsecure page table utilized by the unsecure memory management unit.
15. The method of claim 14, further comprising:
 - placing the GPU in a secure mode or an unsecure mode.
16. The method of claim 15, further comprising:
 - supplying the secure page table to the secure memory management unit and the unsecure page table to the unsecure memory management unit.
17. The method of claim 16, further comprising:
 - sending an instruction to a clear register of the GPU; and
 - clearing and invalidating, in response to the instruction, at least some content from the one or more internal memories when the GPU is transitioned from the secure mode to the unsecure mode.

18. The method of claim 16, further comprising:
 - sending an instruction to a command stream register of the GPU; and
 - clearing and invalidating, in response to the instruction, at least some content from the one or more internal memories when the GPU is transitioned from the secure mode to the unsecure mode.
19. The method of claim 11, further comprising:
 - writing data, one or more hardware blocks of the GPU, to the unsecure portion of the first memory regardless of whether the GPU is in the unsecure mode or the secure mode, wherein the one or more hardware blocks do not have read access to the secure portion of the first memory unit.
20. The method of claim 19, wherein the one or more hardware blocks includes a front end command processor.
21. An apparatus for graphics processing comprising:
 - means for reading a respective resource descriptor for a respective memory resource of a plurality of memory resources;
 - means for receiving a request for a memory transaction to a first memory unit;
 - means for directing, in response to the request, read and write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is a secure resource descriptor to a secure portion of the first memory unit when a graphics processing unit (GPU) is operating according to a secure mode;
 - means for directing, in response to the request, read memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is an unsecure resource descriptor to an unsecure portion of the first memory unit when the GPU is operating according to the secure mode; and
 - means for dropping, in response to the request, write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is the unsecure resource descriptor when the GPU is operating according to the secure mode.

22. The apparatus of claim 21, further comprising:

means for directing, in response to the request, read and write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is the unsecure resource descriptor to an unsecure portion of the first memory unit when the GPU is operating according to the unsecure mode; and

means for dropping, in response to the request, read and write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is the secure resource descriptor when the GPU is operating according to the unsecure mode.

23. The apparatus of claim 21, further comprising:

means for writing data to the secure portion of the first memory unit utilizing a secure memory management unit, the secure memory management unit utilizing a secure page table containing address ranges for the secure portion of the first memory unit; and

means for reading data from the unsecure portion of the first memory unit utilizing an unsecure memory manage unit, the unsecure memory management unit utilizing an unsecure page table containing address ranges for the unsecure portion of the first memory unit.

24. The apparatus of claim 23, further comprising:

means for reading and writing data according to a virtual memory address from a range of virtual memory addresses, wherein the range of virtual memory addresses includes a first range of virtual memory addresses relating to entries in the secure page table utilized by the secure memory management unit, and a second range of virtual memory addresses relating to entries in the unsecure page table utilized by the unsecure memory management unit.

25. The apparatus of claim 24, further comprising:

means for placing the GPU in a secure mode or an unsecure mode.

26. A computer-readable storage medium storing instructions that, when executed, cause one or more processors to:

read a respective resource descriptor for a respective memory resource of a plurality of memory resources;

receive a request for a memory transaction to a first memory unit;

direct, in response to the request, read and write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is a secure resource descriptor to a secure portion of the first memory unit when a graphics processing unit (GPU) is operating according to a secure mode;

direct, in response to the request, read memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is an unsecure resource descriptor to an unsecure portion of the first memory unit when the GPU is operating according to the secure mode; and

drop, in response to the request, write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is the unsecure resource descriptor when the GPU is operating according to the secure mode.

27. The computer-readable storage medium of claim 26, wherein the instructions further cause the one or more processors to:

direct, in response to the request, read and write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is the unsecure resource descriptor to an unsecure portion of the first memory unit when the GPU is operating according to the unsecure mode; and

drop, in response to the request, read and write memory transactions relating to memory resources of the plurality of memory resources for which the respective resource descriptor is the secure resource descriptor when the GPU is operating according to the unsecure mode.

28. The computer-readable storage medium of claim 26, wherein the instructions further cause the one or more processors to:

write data to the secure portion of the first memory unit utilizing a secure memory management unit, the secure memory management unit utilizing a secure page table containing address ranges for the secure portion of the first memory unit; and


read data from the unsecure portion of the first memory unit utilizing an unsecure memory manage unit, the unsecure memory management unit utilizing an unsecure page table containing address ranges for the unsecure portion of the first memory unit.


29. The computer-readable storage medium of claim 28, wherein the instructions further cause the one or more processors to:

read and write data according to a virtual memory address from a range of virtual memory addresses, wherein the range of virtual memory addresses includes a first range of virtual memory addresses relating to entries in the secure page table utilized by the secure memory management unit, and a second range of virtual memory addresses relating to entries in the unsecure page table utilized by the unsecure memory management unit.

30. The computer-readable storage medium of claim 29, wherein the instructions further cause the one or more processors to:

place the GPU in a secure mode or an unsecure mode.

**FIG. 1**

**FIG. 2**




FIG. 3




FIG. 4




FIG.5





FIG. 6


..... ➤ UNSECURE
———— ➤ SECURE



SECURE MODE

FIG. 7A

..... ➤ UNSECURE
———— ➤ SECURE



UNSECURE MODE

FIG. 7B





FIG. 8

**FIG. 9**




FIG. 10




FIG. 11

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2016/043903

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F21/10
ADD.

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F

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

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

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
|-----------|--|-----------------------|
| X | US 2013/305388 A1 (KOTTILINGAL SUDEEP RAVI [US] ET AL) 14 November 2013 (2013-11-14) paragraphs [0034], [0055] - [0058]; figure 4 ----- US 2015/002523 A1 (ZENG THOMAS [US] ET AL) 1 January 2015 (2015-01-01) paragraphs [0018] - [0025]; figures 1,2 ----- EP 1 801 725 A2 (NVIDIA CORP [US]) 27 June 2007 (2007-06-27) abstract; figures 1,7 ----- | 1-30 1-30 1-30 |



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier application or patent but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

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

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

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

"&" document member of the same patent family

| | |
|--|--|
| Date of the actual completion of the international search | Date of mailing of the international search report |
| 13 October 2016 | 26/10/2016 |
| Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016 | Authorized officer Kerschbaumer, J |

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2016/043903

| Patent document cited in search report | Publication date | Patent family member(s) | | | Publication date |
|--|------------------|-------------------------|--|--|--|
| US 2013305388 | A1 | 14-11-2013 | US WO | 2013305388 A1 2013169434 A1 | 14-11-2013 14-11-2013 |
| US 2015002523 | A1 | 01-01-2015 | CA CN EP JP KR US WO | 2912929 A1 105393258 A 3017396 A1 2016524257 A 20160025554 A 2015002523 A1 2015002851 A1 | 08-01-2015 09-03-2016 11-05-2016 12-08-2016 08-03-2016 01-01-2015 08-01-2015 |
| EP 1801725 | A2 | 27-06-2007 | CN EP JP JP KR | 1984298 A 1801725 A2 4740830 B2 2007215159 A 20070063465 A | 20-06-2007 27-06-2007 03-08-2011 23-08-2007 19-06-2007 |