

(51) International Patent Classification:
G06F 9/48 (2006.01)(21) International Application Number:
PCT/US2011/064550(22) International Filing Date:
13 December 2011 (13.12.2011)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/422,595 13 December 2010 (13.12.2010) US
13/289,730 4 November 2011 (04.11.2011) US(71) Applicant (for all designated States except US): **ADVANCED MICRO DEVICES, INC.** [US/US]; One AMD Place, Sunnyvale, CA 94088 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **HARTOG, Scott** [US/US]; One AMD Place, Sunnyvale, CA 94088 (US). **TAYLOR, Clay** [US/US]; One AMD Place, Sunnyvale, CA 94088 (US). **MANTOR, Michael** [US/US]; 1620 Pinar Drive, Orlando, FL 32825 (US). **McGRATH, Kevin**[US/US]; One AMD Place, Sunnyvale, CA 94088 (US). **NUSSBAUM, Sebastien** [US/US]; One AMD Place, Sunnyvale, CA 94088 (US). **JAYASENA, Nuwan** [US/US]; One AMD Place, Sunnyvale, CA 94088 (US). **McCRARY, Rex** [US/US]; 758 Summer Oaks Court, Oviedo, FL 32765 (US). **LEATHER, Mark** [GB/US]; 265 Montclair Rd., Los Gatos, CA 95032 (US). **ROGERS, Philip, J.** [US/US]; 158 Heald Street, Pepperell, MA 01463 (US). **WOLLER, Tom** [US/US]; One AMD Place, Sunnyvale, CA 94088 (US).(74) Agents: **WOOD, Theodore, A.** et al.; Sterne, Kessler, Goldstein & Fox PLLC, 1100 New York Avenue, N.W., Washington, DC 20005-3934 (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, 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, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD,

[Continued on next page]

(54) Title: PROCESSING DEVICE CONTEXT SWITCHING




FIG. 2

(57) Abstract: Methods, systems, and computer readable media embodiments are disclosed for preemptive context-switching of processes running on an accelerated processing device. A method includes, responsive to an exception upon access to a memory by a process running on a accelerated processing device, whether to preempt the process based on the exception, and preempting, based upon the determining, the process from running on the accelerated processing device.



SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR,
TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM,
GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

(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, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, 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,

Published:

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

PROCESSING DEVICE CONTEXT SWITCHING

BACKGROUND

Field of the Invention

[0001] The present invention is generally directed to computer systems. More particularly, the present invention is directed to the execution of processes on an accelerated processing device within a computer system.

Background Art

[0002] The desire to use a graphics processing unit (GPU) for general computation has become much more pronounced recently due to the GPU's exemplary performance per unit power and/or cost. The computational capabilities for GPUs, generally, have grown at a rate exceeding that of the corresponding central processing unit (CPU) platforms. This growth, coupled with the explosion of the mobile computing market (e.g., notebooks, mobile smart phones, tablets, etc.) and its necessary supporting server/enterprise systems, has been used to provide a specified quality of desired user experience. Consequently, the combined use of CPUs and GPUs for executing workloads with data parallel content is becoming a volume technology.

[0003] However, GPUs have traditionally operated in a constrained programming environment, available primarily for the acceleration of graphics. These constraints arose from the fact that GPUs did not have as rich a programming ecosystem as CPUs. Their use, therefore, has been mostly limited to two dimensional (2D) and three dimensional (3D) graphics and a few leading edge multimedia applications, which are already accustomed to dealing with graphics and video application programming interfaces(APIs).

[0004] With the advent of multi-vendor supported OpenCL® and DirectCompute®, standard APIs and supporting tools, the limitations of the GPUs in traditional applications has been extended beyond traditional graphics. Although OpenCL and DirectCompute are a promising start, there are many hurdles remaining to creating an environment and ecosystem that allows the combination of a CPU and a GPU to be used as fluidly as the CPU for most programming tasks.

- [0005] Existing computing systems often include multiple processing devices. For example, some computing systems include both a CPU and a GPU on separate chips (e.g., the CPU might be located on a motherboard and the GPU might be located on a graphics card) or in a single chip package. Both of these arrangements, however, still include significant challenges associated with (i) separate memory systems, (ii) efficient scheduling, (iii) providing quality of service (QoS) guarantees between processes, (iv) programming model, and (v) compiling to multiple target instruction set architectures (ISAs) – all while minimizing power consumption.
- [0006] For example, the discrete chip arrangement forces system and software architects to utilize chip to chip interfaces for each processor to access memory. While these external interfaces (e.g., chip to chip) negatively affect memory latency and power consumption for cooperating heterogeneous processors, the separate memory systems (i.e., separate address spaces) and driver managed shared memory create overhead that becomes unacceptable for fine grain offload.
- [0007] Both the discrete and single chip arrangements can limit the types of commands that can be sent to the GPU for execution. This performance based limitation exists because the CPU may relatively quickly require the results of the operations performed by these computational commands. However, because of the high overhead of dispatching work to the GPU in current systems and the fact that these commands may have to wait in line for other previously-issued commands to be executed first, the latency incurred by sending computational commands to the GPU is often unacceptable.
- [0008] Given that a traditional GPU may not efficiently execute some computational commands, the commands must then be executed within the CPU. Having to execute the commands on the CPU increases the processing burden on the CPU and can hamper overall system performance.
- [0009] Although GPUs provide excellent opportunities for computational offloading, traditional GPUs may not be suitable for system-software-driven process management that is desired for efficient operation in some multi-processor environments. These limitations can create several problems.
- [0010] For example, since processes cannot be efficiently identified and/or preempted, a rogue process can occupy the GPU for arbitrary amounts of time. The occupying of the GPU by rogue processes for arbitrary amounts of time can prevent the effective utilization of the available system capacity, and can prevent or significantly reduce the

processing progress of the system. In other cases, the ability to context switch off the hardware is severely constrained - occurring at very coarse granularity and only at a very limited set of points in a program's execution.

SUMMARY OF EMBODIMENTS

[0011] Therefore, what is needed is a method and system for efficiently preempting processes from a graphics processing device, such as a GPU, and context switching another process onto the graphics processing device.

[0012] Although GPUs, accelerated processing units (APUs), and general purpose use of the graphics processing unit (GPGPU) are commonly used terms in this field, the expression "accelerated processing device (APD)" is considered to be a broader expression. For example, APD refers to any cooperating collection of hardware and/or software that performs those functions and computations associated with accelerating graphics processing tasks, data parallel tasks, or nested data parallel tasks in an accelerated manner with respect to resources such as conventional CPUs, conventional GPUs, and/or combinations thereof.

[0013] An embodiment of the present invention includes, responsive to an exception upon access to a memory by a process running on an APD, determining whether to preempt the process based on the exception, and preempting the process from running on the accelerated processing device.

[0014] Further features and advantages of the invention, as well as the structure and operation of various embodiments of the invention, are described in detail below with reference to the accompanying drawings. It is noted that the invention is not limited to the specific embodiments described herein. Such embodiments are presented herein for illustrative purposes only. Additional embodiments will be apparent to persons skilled in the relevant art(s) based on the teachings contained herein.

BRIEF DESCRIPTION OF THE DRAWINGS/FIGURES

[0015] The accompanying drawings, which are incorporated herein and form part of the specification, illustrate the present invention and, together with the description, further serve to explain the principles of the invention and to enable a person skilled in the pertinent art to make and use the invention. Various embodiments of the present

invention are described below with reference to the drawings, wherein like reference numerals are used to refer to like elements throughout.

[0016] FIG. 1A is an illustrative block diagram of a processing system, in accordance with embodiments of the present invention.

[0017] FIG. 1B is an illustrative block diagram illustration of the accelerated processing device (APD) illustrated in FIG. 1A.

[0018] FIG. 2 is a flowchart illustrating a method for APD context switching, according to an embodiment of the present invention.

[0019] FIG. 3 is a flowchart illustrating a method for detecting a page fault/memory exception, according to an embodiment of the present invention.

[0020] FIG. 4 is a flowchart illustrating a method for the APD to notify the OS of a page fault, according to an embodiment of the present invention.

[0021] FIG. 5 is a flowchart illustrating a method for determining if a APD should be context switched, according to an embodiment of the present invention.

DETAILED DESCRIPTION OF EMBODIMENTS OF THE INVENTION

[0022] In the detailed description that follows, references to "one embodiment," "an embodiment," "an example embodiment," etc., indicate that the embodiment described may include a particular feature, structure, or characteristic, but every embodiment may not necessarily include the particular feature, structure, or characteristic. Moreover, such phrases are not necessarily referring to the same embodiment. Further, when a particular feature, structure, or characteristic is described in connection with an embodiment, it is submitted that it is within the knowledge of one skilled in the art to affect such feature, structure, or characteristic in connection with other embodiments whether or not explicitly described.

[0023] The term "embodiments of the invention" does not require that all embodiments of the invention include the discussed feature, advantage or mode of operation. Alternate embodiments may be devised without departing from the scope of the invention, and well-known elements of the invention may not be described in detail or may be omitted so as not to obscure the relevant details of the invention. In addition, the terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting of the invention. For example, as used herein, the singular forms "a", "an" and "the" are intended to include the plural forms as well, unless the context

clearly indicates otherwise. It will be further understood that the terms “comprises,” “comprising,” “includes” and/or “including,” when used herein, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof.

[0024] FIG. 1A is an exemplary illustration of a unified computing system 100 including two processors, a CPU 102 and an APD 104. CPU 102 can include one or more single or multi core CPUs. In one embodiment of the present invention, the system 100 is formed on a single silicon die or package, combining CPU 102 and APD 104 to provide a unified programming and execution environment. This environment enables the APD104 to be used as fluidly as the CPU 102 for some programming tasks. However, it is not an absolute requirement of this invention that the CPU 102 and APD 104 be formed on a single silicon die. In some embodiments, it is possible for them to be formed separately and mounted on the same or different substrates.

[0025] In one example, system 100 also includes a memory 106, an OS 108, and a communication infrastructure 109. The OS 108 and the communication infrastructure 109 are discussed in greater detail below.

[0026] The system 100 also includes a kernel mode driver (KMD) 110, a software scheduler (SWS) 112, and a memory management unit 116, such as input/output memory management unit (IOMMU). Components of system 100 can be implemented as hardware, firmware, software, or any combination thereof. A person of ordinary skill in the art will appreciate that system 100 may include one or more software, hardware, and firmware components in addition to, or different from, that shown in the embodiment shown in FIG. 1A.

[0027] In one example, a driver, such as KMD 110, typically communicates with a device through a computer bus or communications subsystem to which the hardware connects. When a calling program invokes a routine in the driver, the driver issues instructions to the device. Once the device sends data back to the driver, the driver may invoke routines in the original calling program. In one example, drivers are hardware-dependent and operating-system-specific. They usually provide the interrupt handling required for any necessary asynchronous time-dependent hardware interface. Device drivers, particularly on modern Microsoft Windows® platforms, can run in kernel-mode (Ring 0) or in user-mode (Ring 3).

- [0028] A benefit of running a driver in user mode is improved stability, since a poorly written user mode device driver cannot crash the system by overwriting kernel memory. On the other hand, user/kernel-mode transitions usually impose a considerable performance overhead, thereby prohibiting user mode-drivers for low latency and high throughput requirements. Kernel space can be accessed by user modules only through the use of system calls. End user programs like the UNIX shell or other GUI based applications are part of the user space. These applications interact with hardware through kernel supported functions.
- [0029] CPU 102 can include (not shown) one or more of a control processor, field programmable gate array (FPGA), application specific integrated circuit (ASIC), or digital signal processor (DSP). CPU 102, for example, executes the control logic, including the OS 108, KMD 110, SWS 112, and applications 111, that control the operation of computing system 100. In this illustrative embodiment, CPU 102, according to one embodiment, initiates and controls the execution of applications 111 by, for example, distributing the processing associated with that application across the CPU 102 and other processing resources, such as the APD 104.
- [0030] APD 104, among other things, executes instructions and programs for selected functions, such as graphics operations and other operations that may be, for example, particularly suited for parallel processing. In general, APD 104 can be frequently used for executing graphics pipeline operations, such as pixel operations, geometric computations, and rendering an image to a display. In various embodiments of the present invention, APD 104 can also execute compute processing operations (e.g., those operations unrelated to graphics such as, for example, video operations, physics simulations, computational fluid dynamics, etc.), based on commands or instructions received from CPU 102.
- [0031] For example, commands can be considered as special instructions that are not typically defined in the instruction set architecture (ISA). A command may be executed by a special processor such as a dispatch processor, command processor, or network controller. On the other hand, instructions can be considered as, for example, a single operation of a processor within a computer architecture. In one example, when using two sets of ISAs, some instructions are used to execute x86 programs and some instructions are used to execute kernels on APD compute unit.

[0032] In an illustrative embodiment, CPU 102 transmits selected commands and/or other instructions to APD 104. These selected instructions can include graphics instructions and other commands amenable to parallel execution. These selected instructions, that can also include compute processing instructions, can be executed substantially independently from CPU 102.

[0033] APD104 can include its own compute units (not shown), such as, but not limited to, one or more single instruction multiple data (SIMD) processing cores. As referred to herein, a SIMD is a pipeline, or programming model, where a kernel is executed concurrently on multiple processing elements each with its own data and a shared program counter. All processing elements execute an identical set of instructions. The use of predication enables work-items to participate or not for each issued command.

[0034] In one example, each APD 104 compute unit can include one or more scalar and/or vector floating-point units and/or arithmetic and logic units (ALUs). The APD compute unit can also include special purpose processing units (not shown), such as inverse-square root units and sine/cosine units. In one example, the APD compute units are referred to herein collectively as shader core 122.

[0035] Having one or more SIMDs, in general, makes APD 104 ideally suited for execution of data-parallel tasks such as those that are common in graphics processing.

[0036] Some graphics pipeline operations, such as pixel processing, and other parallel computation operations, can require that the same instruction stream or compute kernel be performed on streams or collections of input data elements. Respective instantiations of the same compute kernel can be executed concurrently on multiple compute units in shader core 122 in order to process such data elements in parallel. As referred to herein, for example, a compute kernel is a function containing instructions declared in a program and executed on an APD compute unit. This function is also referred to as a kernel, a shader, a shader program, or a program.

[0037] In one illustrative embodiment, each compute unit (e.g., SIMD processing core) can execute a respective instantiation of a particular work-item to process incoming data. A work-item is one of a collection of parallel executions of a kernel invoked on a device by an instruction. A work-item can be executed by one or more processing elements as part of a work-group executing on a compute unit.

[0038] A work-item is distinguished from other executions within the collection by its global ID and local ID. In one example, a subset of work-items in a workgroup that

execute simultaneously together on a SIMD can be referred to as a wavefront 136. The width of a wavefront is a characteristic of the hardware of the compute unit (e.g., SIMD processing core). As referred to herein, a workgroup is a collection of related work-items that execute on a single compute unit. The work-items in the group execute the same kernel and share local memory and work-group barriers.

[0039] In the exemplary embodiment, all wavefronts from a workgroup are processed on the same SIMD processing core. Instructions across a wavefront are issued one at a time, and when all work-items follow the same control flow, each work-item executes the same program. Wavefronts can also be referred to as warps, vectors or threads.

[0040] An execution mask and work-item predication are used to enable divergent control flow within a wavefront, where each individual work-item can take a unique code path through the kernel. Partially populated wavefronts can be processed when a full set of work-items is not available at wavefront start time. For example, shader core 122 can simultaneously execute a predetermined number of wavefronts 136, each wavefront 136 comprising a multiple work-items.

[0041] Within the system 100, APD 104 includes its own memory, such as graphics memory 130 (although memory 130 is not limited to graphics only use). Graphics memory 130 provides a local memory for use during computations in APD 104. Individual compute units (not shown) within shader core 122 can have their own local data store (not shown). In one embodiment, APD 104 includes access to local graphics memory 130, as well as access to the memory 106. In another embodiment, APD 104 can include access to dynamic random access memory (DRAM) or other such memories (not shown) attached directly to the APD104 and separately from memory 106.

[0042] In the example shown, APD 104 also includes one or "n" number of command processors (CPs) 124. CP 124 controls the processing within APD 104. CP 124 also retrieves instructions to be executed from command buffers 125 in memory 106 and coordinates the execution of those instructions on APD 104.

[0043] In one example, CPU 102 inputs instructions based on applications 111 into appropriate command buffers 125. As referred to herein, an application is the combination of the program parts that will execute on the compute units within the CPU and APD.

[0044] A plurality of command buffers 125 can be maintained with each process scheduled for execution on the APD 104.

- [0045] CP 124 can be implemented in hardware, firmware, or software, or a combination thereof. In one embodiment, CP 124 is implemented as a reduced instruction set computer (RISC) engine with microcode for implementing logic including scheduling logic.
- [0046] APD 104 also includes one or "n" number of dispatch controllers (DCs) 126. In the present application, the term dispatch refers to a instruction executed by a dispatch controller that uses the context state to initiate the start of the execution of a kernel for a set of work groups on a set of compute units. DC 126 includes logic to initiate workgroups in the shader core 122. In some embodiments, DC 126 can be implemented as part of CP 124.
- [0047] System 100 also includes a hardware scheduler (HWS) 128 for selecting a process from a run list 150 for execution on APD 104. HWS 128 can select processes from run list 150 using round robin methodology, priority level, or based on other scheduling policies. The priority level, for example, can be dynamically determined. HWS 128 can also include functionality to manage the run list 150, for example, by adding new processes and by deleting existing processes from run-list 150. The run list management logic of HWS 128 is sometimes referred to as a run list controller (RLC).
- [0048] In various embodiments of the present invention, when HWS 128 initiates the execution of a process from run list 150, CP 124 begins retrieving and executing instructions from the corresponding command buffer 125. In some instances, CP 124 can generate one or more instructions to be executed within APD 104, which correspond with instructions received from CPU 102. In one embodiment, CP 124, together with other components, implements a prioritizing and scheduling of instructions on APD 104 in a manner that improves or maximizes the utilization of the resources of APD 104 and/or system 100.
- [0049] APD 104 can have access to, or may include, an interrupt generator 146. Interrupt generator 146 can be configured by APD 104 to interrupt the OS 108 when interrupt events, such as page faults, are encountered by APD 104. For example, APD 104 can rely on interrupt generation logic within IOMMU 116 to create the page fault interrupts noted above.
- [0050] APD 104 can also include preemption and context switch logic 120 for preempting a process currently running within shader core 122. Context switch logic 120,

for example, includes functionality to stop the process and save its current state (e.g., shader core 122 state, and CP 124 state).

[0051] As referred to herein, the term state can include an initial state, an intermediate state, and/or a final state. An initial state is a starting point for a machine to process an input data set according to a programming in order to create an output set of data. There is an intermediate state, for example, that needs to be stored at several points to enable the processing to make forward progress. This intermediate state is sometimes stored to allow a continuation of execution at a later time when interrupted by some other process. There is also final state that can be recorded as part of the output data set

[0052] Preemption and context switch logic 120 can also include logic to context switch another process into the APD 104. The functionality to context switch another process into running on the APD 104 may include instantiating the process, for example, through the CP 124 and DC 126 to run on APD 104, restoring any previously saved state for that process, and starting its execution.

[0053] Memory 106 can include non-persistent memory such as DRAM (not shown). Memory 106 can store, e.g., processing logic instructions, constant values, and variable values during execution of portions of applications or other processing logic. For example, in one embodiment, parts of control logic to perform one or more operations on CPU 102 can reside within memory 106 during execution of the respective portions of the operation by CPU 102.

[0054] During execution, respective applications, OS functions, processing logic instructions, and system software can reside in memory 106. Control logic instructions fundamental to OS 108 will generally reside in memory 106 during execution. Other software instructions, including, for example, kernel mode driver 110 and software scheduler 112 can also reside in memory 106 during execution of system 100. Software instructions, prior to being loaded into system memory 106 for execution, can reside on a persistent storage computer readable media 151 such as a hard disk, flash memory, and the like.

[0055] In this example, memory 106 includes command buffers 125 that are used by CPU 102 to send instructions to APD 104. Memory 106 also contains process lists and process information (e.g., active list 152 and process control blocks 154). These lists, as well as the information, are used by scheduling software executing on CPU 102 to communicate scheduling information to APD 104 and/or related scheduling hardware.

Access to memory 106 can be managed by a memory controller 140, which is coupled to memory 106. For example, requests from CPU 102, or from other devices, for reading from or for writing to memory 106 are managed by the memory controller 140.

[0056] Referring back to other aspects of system 100, IOMMU 116 is a multi-context memory management unit.

[0057] As used herein, context can be considered the environment within which the kernels execute and the domain in which synchronization and memory management is defined. The context can include a set of devices, the memory accessible to those devices, the corresponding memory properties and one or more command-queues used to schedule execution of a kernel(s) or operations on memory objects.

[0058] Referring back to the example shown in FIG. 1A, IOMMU 116 includes logic to perform virtual to physical address translation for memory page access for devices including APD 104. IOMMU 116 may also include logic to generate interrupts, for example, when a page access by a device such as APD 104 results in a page fault. IOMMU 116 may also include, or have access to, a translation lookaside buffer (TLB) 118. TLB 118, as an example, can be implemented in a content addressable memory (CAM) to accelerate translation of logical (i.e., virtual) memory addresses to physical memory addresses for requests made by APD 104 for data in memory 106.

[0059] In the example shown, communication infrastructure 109 interconnects the components of system 100 as needed. Communication infrastructure 109 can include (not shown) one or more of a peripheral component interconnect (PCI) bus, extended PCI (PCI-E) bus, advanced microcontroller bus architecture (AMBA) bus, accelerated graphics port (AGP), or other such communication infrastructure. Communications infrastructure 109 can also include an Ethernet, or similar network, or any suitable physical communications infrastructure that satisfies an application's data transfer rate requirements. Communication infrastructure 109 includes the functionality to interconnect components including components of computing system 100.

[0060] In this example, OS 108 includes functionality to manage the hardware components of system 100 and to provide common services. In various embodiments, OS 108 can execute on CPU 102 and provide common services. These common services can include, for example, scheduling applications for execution within CPU 102, fault management, interrupt service, as well as processing the input and output of other applications.

- [0061] In some embodiments, based on interrupts generated by an interrupt controller, such as interrupt controller 148, OS 108 invokes an appropriate interrupt handling routine. For example, upon detecting a page fault interrupt, OS 108 may invoke an interrupt handler to initiate loading of the relevant page into memory 106 and to update corresponding page tables.
- [0062] OS 108 may also include functionality to protect system 100 by ensuring that access to hardware components is mediated through OS managed kernel functionality. In effect, OS 108 ensures that applications, such as applications 111, run on CPU 102 in user space. OS 108 also ensures that applications 111 invoke kernel functionality provided by the OS to access hardware and/or input/output functionality.
- [0063] According to an embodiment, the operating system includes an OS memory manager 153 and an OS scheduler 155. OS memory manager 153 may operate to manage memory objects such as, but not limited to, page tables 157 and page event queues 156. Page tables 157 include tables that indicate the location of pages currently loaded in memory 106. Page event queue 156 includes a queue in which page related events, such as page fault events, are enqueued by other devices, such as IOMMU 116, in order to communicate page related information to the OS. OS scheduler 155 may operate to determine the status of page faults and to determine if a APD context switch should be initiated in response to a page fault. Exception statistics 159 may collect statistics related to page faults and other exceptions.
- [0064] By way of example, applications 111 include various programs or instructions to perform user computations that are also executed on CPU 102. CPU 102 can seamlessly send selected instructions for processing on the APD 104.
- [0065] In one example, KMD 110 implements an application program interface (API) through which CPU 102, or applications executing on CPU 102 or other logic, can invoke APD 104 functionality. For example, KMD 110 can enqueue instructions from CPU 102 to command buffers 125 from which APD 104 will subsequently retrieve the instructions. Additionally, KMD 110 can, together with SWS 112, perform scheduling of processes to be executed on APD 104. SWS 112, for example, can include logic to maintain a prioritized list of processes to be executed on the APD.
- [0066] In other embodiments of the present invention, applications executing on CPU 102 can entirely bypass KMD 110 when enqueueing instructions.

[0067] In some embodiments, SWS 112 maintains an active list 152 in memory 106 of processes to be executed on APD 104. SWS 112 also selects a subset of the processes in active list 152 to be managed by HWS 128 in the hardware. Information relevant for running each process on APD 104 is communicated from CPU 102 to APD104 through process control blocks (PCB) 154.

[0068] Processing logic for applications, OS, and system software can include commands and/or other instructions specified in a programming language such as C and/or in a hardware description language such as Verilog, RTL, or netlists, to enable ultimately configuring a manufacturing process through the generation of maskworks/photomasks to generate a hardware device embodying aspects of the invention described herein.

[0069] A person of skill in the art will understand, upon reading this description, that computing system 100 can include more or fewer components than shown in FIG. 1A. For example, computing system 100 can include one or more input interfaces, non-volatile storage, one or more output interfaces, network interfaces, and one or more displays or display interfaces.

[0070] FIG. 1B is an embodiment showing a more detailed illustration of APD 104 shown in FIG. 1A. In FIG. 1B, CP 124 can include CP pipelines 124a, 124b, and 124c. CP 124 can be configured to process the command lists that are provided as inputs from command buffers 125, shown in FIG. 1A. In the exemplary operation of FIG. 1B, CP input 0 (124a) is responsible for driving commands into a graphics pipeline 162. CP inputs 1 and 2 (124b and 124c) forward instructions to a compute pipeline 160. Also provided is a controller mechanism 166 for controlling operation of HWS 128.

[0071] In FIG. 1B, graphics pipeline 162 can include a set of blocks, referred to herein as ordered pipeline 164. As an example, ordered pipeline 164 includes a vertex group translator (VGT) 164a, a primitive assembler (PA) 164b, a scan converter (SC) 164c, and a shader-export, render-back unit (SX/RB) 176. Each block within ordered pipeline 164 may represent a different stage of graphics processing within graphics pipeline 162. Ordered pipeline 164 can be a fixed function hardware pipeline. Other implementations can be used that would also be within the spirit and scope of the present invention.

[0072] Although only a small amount of data may be provided as an input to graphics pipeline 162, this data will be amplified by the time it is provided as an output from graphics pipeline 162. Graphics pipeline 162 also includes DC 166 for counting through

ranges within work-item groups received from CP pipeline 124a. Compute work submitted through DC 166 is semi-synchronous with graphics pipeline 162.

[0073] Compute pipeline 160 includes shader DCs 168 and 170. Each of the DCs 168 and 170 is configured to count through compute ranges within work groups received from CP pipelines 124b and 124c.

[0074] The DCs 166, 168, and 170, illustrated in FIG. 1B, receive the input ranges, break the ranges down into workgroups, and then forward the workgroups to shader core 122.

[0075] Since graphics pipeline 162 is generally a fixed function pipeline, it is difficult to save and restore its state, and as a result, the graphics pipeline 162 is difficult to context switch. Therefore, in most cases context switching, as discussed herein, does not pertain to context switching among graphics processes. An exception is for graphics work in shader core 122, which can be context switched.

[0076] After the processing of work within graphics pipeline 162 has been completed, the completed work is processed through a render back unit 176, which does depth and color calculations, and then writes its final results to memory 130.

[0077] Shader core 122 can be shared by graphics pipeline 162 and compute pipeline 160. Shader core 122 can be a general processor configured to run wavefronts. In one example, all work within compute pipeline 160 is processed within shader core 122. Shader core 122 runs programmable software code and includes various forms of data, such as state data.

[0078] FIG. 2 illustrates a flowchart of a method 200 for APD context switching, according to an embodiment of the present invention. For example, method 200 may run on system 100 shown in Figures 1A and 1B. With method 200, an APD can detect a memory exception, e.g., a page fault, and is able to implement a context switch of processes. The steps shown in method 200 may not occur in the order shown, or require all of the steps.

[0079] In step 202, the APD runs a first process. Running of a process which could include the HWS or SWS-selecting the process from a run-list of processes and maps it to the APD which will process it on the appropriate processing pipeline. Graphics processing utilizes the graphics pipeline of the APD, and the compute pipeline utilizes the compute pipeline. Both types of processes can utilize a shader core of the APD for processing, e.g., as discussed above.

[0080] In step 204, the first process running on the APD requests data. According to an embodiment, the request from the first process running on the APD is intercepted by, or directed to, a memory management unit. The memory management unit can be an IOMMU communicatively coupled to the APD, e.g., as discussed above with regards to system 100. The IOMMU may be incorporated in the APD, may be incorporated in another memory management unit, such as a memory controller, or may be implemented separately. The IOMMU can include the functionality to translate between the virtual memory address space as seen by the APD and the system memory physical address space.

[0081] In step 206, a page fault or other exception is detected. According to an embodiment, the IOMMU receives the request for data from the APD and attempts to perform the translation of the requested data from the APD's virtual address space to the physical address space. The IOMMU then attempts to retrieve the data from memory based upon the determined physical address of the data. The IOMMU attempts to retrieve the requested data from, for example, system memory, such as system memory 106.

[0082] In an example, if the requested data is not in system memory, a memory exception or page fault is triggered. The page fault can be triggered by the IOMMU or other hardware or software component associated with the IOMMU access to memory. A page fault, as used herein, indicates that a requested memory object, such as a page of in-memory data, does not exist in physical memory that is accessible to the requesting entity. A page fault can be triggered, for example, if an entry corresponding to the virtual address is not present in the page table. A page fault can also be triggered for other reasons, such as when a page table entry for the requested virtual address exists but the corresponding page is not accessible to the requesting process. For example, a page may not be accessible to the requesting process due to synchronization or mapping issues between the page table accessible to the IOMMU and system memory, due to memory protection errors such as when the requesting device or process is not permitted to access the area of memory in which the requested address is present, and the like.

[0083] In the exemplary embodiment, upon receiving from the APD a request for data, the IOMMU accesses a TLB with the request for data. The TLB can be implemented in the IOMMU, the APD, or separately. The TLB is a cache, typically implemented in a content addressable memory (CAM), which performs translation between the system

memory physical address space and a virtual address space in a more efficient manner than by using page table lookup. The TLB is looked up using a virtual address as seen by the process executing on the APD. If the TLB currently has an entry corresponding to that virtual address, then the corresponding physical memory address is returned to the IOMMU. The IOMMU can then attempt to retrieve the corresponding page from the memory.

[0084] If the TLB currently does not have an entry corresponding to that virtual address, then a TLB miss occurs. Upon a TLB miss, the IOMMU may lookup the requested virtual address in the page table. The process of the IOMMU looking up the page table for a virtual address is sometimes referred to as a page table walk. The page table walk, in general, is more time consuming than looking up that address using a TLB.

[0085] In step 208, notification of the page fault issues (in the exemplary embodiment it is issued to the OS). According to an embodiment, the OS is notified by an interrupt generated by the IOMMU. The IOMMU may generate the interrupt upon detecting that the requested virtual address does not exist in the TLB and in the page table accessible to the IOMMU. According to another embodiment, the interrupt can be generated by the IOMMU upon receiving a signal from the APD requesting the generation of the interrupt.

[0086] In step 210, a determination is made if there should be a context switch or a stall in response to the detected page fault. In some embodiments, this determination may be made by software such as the OS. According to an embodiment, an interrupt handler of the OS traps the interrupt generated by the IOMMU to indicate the page fault. The interrupt handler can invoke the logic to determine if, based on the detected page fault, a context switch or stall should be implemented in the APD. The OS functionality to determine if a context switch should be attempted upon a page fault can be implemented by an OS scheduler, for example.

[0087] The decision to select between a context switch or stall can be based upon the type of the page fault or an estimate of the time required to recover from such page fault. If the page fault is due to a page that is not loaded into memory from a virtual memory, such page faults can be considered as requiring substantial time to recover.

[0088] According to other embodiments, even in cases where the requested page is required to be loaded from virtual memory to system memory, the required recovery time to load the page and update the corresponding tables (such as page tables) can be

considered short or long based on factors such as the location in virtual memory and current utilization of system memory. If the page fault is due to a page table entry not being updated after a page is loaded into memory, then recovery from the page fault may require only a short time.

[0089] In the exemplary embodiment, the OS can decide to initiate a context switch in cases where the page fault is considered a long page fault (i.e., expected recovery time is substantial), and can decide to stall the process on the APD if the page fault is a short page fault (i.e., expected recovery time is short or insubstantial). The functionality to load the page into memory and update page tables may be included in an OS memory manager.

[0090] According to yet one other embodiment, the OS maintains statistics corresponding to the various page faults. For example, the maintained statistics can include, but are not limited to, the type and count of each type of page fault, page faults per active process, page faults for pages from predetermined areas/segments of memory, and the like. The statistics can also include response time related information, such as, but not limited to, the time that was required for loading the pages in response to particular types of page faults in a predetermined interval. The statistics can be updated, for example, by the interrupt handler servicing the page fault interrupt and/or the OS memory manager. The decision to initiate a context switch based upon a detected page fault can, in some embodiments, be based upon such statistics available to the OS. For example, if the statistics maintained by the OS is indicative of the requesting process having page faulted beyond a threshold number of times, then the OS may determine that the process should be context switch.

[0091] The OS can also consider other factors, such as, the number of processes and/or the priority of processes that are ready to be run on the APD when determining whether to initiate a context switch. For example, the OS can consider processes in the active list and/or run-list of processes queued for running by the APD. According to an embodiment, the OS can more frequently determine that lower priority processes are to be context switched when high priority processes are enqueued for running on the APD.

[0092] If, in step 210, it is determined that no context switch is required and that a stall should be adequate, then in step 212 other operations can be performed in response to the stall. A stall does not require the OS to communicate that determination to the APD or other component. According to another embodiment, the OS can update any statistics

maintained, such as, for example the number of stalls and/or the number of stalls caused by the requesting process.

[0093] According to another embodiment, the OS can communicate the determination to stall to the APD and/or IOMMU. For example, a message can be queued in a buffer monitored by the IOMMU indicating the determination to stall. The IOMMU can, for example, initiate or continue attempts to perform the address translation based on the communication from the OS.

[0094] If, in step 210, it is determined that a context switch is to be initiated, then in step 214, initiation of the preemption of the current process from the APD occurs.

[0095] Subsequent to initiating the preemption of the current process, in step 216 other processes can be initiated, e.g., a the running of a second process in the APD. The initiation of the preemption of an APD process and the initiation of the context switch can be performed by, for example, OS scheduler. OS aided preemption of APD processes and context switching is further described below in relation to FIG. 5.

[0096] FIG. 3 is a flowchart illustrating a method 206 for detection of a memory exception or page fault, according to an embodiment of the present invention. For example, method 206 may run on system 100 in Figures 1A and 1B. According to an embodiment, steps 302-308 may be used in performing the functionality of step 206 discussed above. The steps shown in method 206 may not occur in the order shown, or require all of the steps.

[0097] In step 302, the IOMMU receives a request for data from the APD. The request can include the virtual address of a single data element or a block of data.

[0098] In step 304, the IOMMU initiates a TLB lookup for the requested data. The TLB lookup can be performed using the virtual address provided by the APD, or a virtual address derived based upon the received virtual address. For example, the IOMMU can derive an address of a block of data based upon the address provided by the APD.

[0099] The TLB lookup, if successful, results in the TLB returning the physical address corresponding to the virtual address on which the lookup was based.

[0100] If the TLB lookup is not successful (i.e., TLB miss), an indication of the miss is returned to the IOMMU and/or APD. Accordingly, the IOMMU and/or APD, upon receiving the TLB miss indication, can choose to initiate a retry sequence of performing the TLB lookup. For example, the retry sequence may involve continually retrying the TLB lookup at predetermined intervals.

[0101] In step 306, according to another embodiment, upon receiving the TLB miss indication, the APD or IOMMU attempts to locate the address in the page tables. According to an embodiment, the IOMMU receives the TLB miss indication and initiates the page table walk to locate the address in the page tables. In another embodiment, the APD is notified of the miss, for example, by the IOMMU which initially receives the TLB request miss notification, or by receiving the miss notification directly in the APD from the TLB. The APD may then request the IOMMU to translate the address and/or retrieve the data corresponding to the requested address using the page tables. According to an embodiment, the APD can transmit an Address Translation Service (ATS) Request to the IOMMU to initiate the page walk for the requested virtual address.

[0102] If an entry corresponding to the virtual address is located in the page table, then the corresponding request for the data can be sent to memory to retrieve data as appropriate. If no entry corresponding to the requested virtual address is located in the page table, then the IOMMU signals a page fault to the APD. According to an embodiment, the APD is notified of the page fault using an ATS Response.

[0103] In step 308, the APD requests fault handling of the page fault from the OS. According to an embodiment, the APD makes the request through the IOMMU. The fault handling request to the OS can be generated by an interrupt and/or message (such as a Page Request Interface Request) that the APD exchanges with the IOMMU.

[0104] FIG. 4 illustrates a flowchart of a method 208 to notify the OS about a page fault, according to an embodiment of the present invention. For example, method 208 may run on system 100 in Figures 1A and 1B. According to an embodiment, steps 402-08 can be performed in order to implement at least some of the functionality of step 208 described above. The steps shown in method 208 may not occur in the order shown, or require all of the steps.

[0105] In step 402, an interrupt corresponding to page faults is generated. According to an embodiment, the interrupt is generated by the IOMMU on behalf of the APD. According to another embodiment, the APD can directly cause the generation of the interrupt without going through the IOMMU.

[0106] In step 404, a page fault event is enqueued in a page fault event queue in system memory. The enqueued page fault event can include information necessary for the OS to service the page fault and to maintain statistics about page faults. According to an embodiment, the page fault event includes the type of page fault, the time the page fault

is generated, virtual address upon which the page fault was generated, the process requesting the virtual address upon which the page fault was generated, and the identity of the device causing the page fault.

[0107] In step 406, the OS allocates the page for which the page fault was generated, and updates corresponding page table entries to indicate that the page was allocated. According to an embodiment, the OS becomes aware of the page fault, for example, by being notified by an interrupt service routine that traps the interrupt generated by the IOMMU and/or APD in step 402. Upon receiving the interrupt, the OS can retrieve the corresponding page fault event from a page fault event buffer. The page fault event provides the OS with detailed information regarding the page fault. According to an embodiment, if the page fault is due to a page not being in memory, the OS attempts to load the page into memory and update the page tables correspondingly. If the page fault is due to the page tables not being correctly updated with information regarding pages already in memory, the OS updates the corresponding entry in the page table.

[0108] In step 408, the OS notifies the IOMMU and/or APD that the page was loaded. According to an embodiment, this notification is delivered using an interrupt which is trapped by the IOMMU. The OS may, according to embodiments, issue this notification upon the initiating or upon the completing of the loading of the corresponding page and/or updating of the corresponding page tables.

[0109] It should be noted that the OS, upon being notified of the page fault in step 404, also initiates the determination whether to preemptively context switch the process on the APD. This is described, for example, above in relation to step 210 and also below in relation to FIG. 5.

[0110] FIG. 5 is a flowchart illustrating a method 210, according to an embodiment of the present invention. For example, method 210 includes steps 502-506 for determining if an APD should be context switched, according to an embodiment of the present invention. The method 502-506 can be performed, for example, in implementing step 210 described above on, for example, system 100 shown in Figures 1A and 1B. The steps shown in method 210 may not occur in the order shown, or require all of the steps.

[0111] In step 502, the OS accesses information regarding the page fault and other page fault statistics that are accessible to the OS. For example, information about the current page fault can be available based upon the corresponding page fault event enqueued by

the IOMMU. Statistics of page faults can be maintained by the OS based upon each of the page faults of which it is notified. The statistics can be stored in system memory.

[0112] In step 504, the determination whether to preemptively context switch or stall is made by the OS, for example, by an OS memory scheduler.

[0113] If it is determined to context switch, then, in step 506, the context switch is initiated by the OS. According to an embodiment, the OS signals a software scheduler (SWS) for the APD to initiate the context switch. For example, OS scheduler 155 can signal SWS 112. The SWS can be implemented as either a part of a kernel mode driver (KMD) for the APD or as a separate module that communicates with the APD through the KMD. The KMD, SWS and OS execute on the CPU.

[0114] According to an embodiment, upon being signaled by the OS to initiate a context switch on the APD, the SWS performs the scheduling of processes to run on the APD. The SWS can maintain a list of processes from which the processes to be run on the APD are selected. The list of runnable processes can be maintained as a single or multi-level list. According to an embodiment, the list of runnable processes is maintained as a two-level list. At the higher level, the SWS enqueues the runnable processes to an active-list, such as active list 152 maintained in system memory. The active-list includes an entry for each process that the SWS has scheduled to be run on the APD. Each entry in the active-list can include, or can point to, information regarding the process that may be needed for the execution of the process on the APD. For example, each entry in the active-list can point to a corresponding entry in the list of process control blocks 154 in system memory. The process control blocks can include information regarding, for example, process state, program counter, and the like. The SWS can select some processes from the active-list and enqueue them in a second level list of runnable processes referred to herein as the run-list. According to an embodiment, the run-list may include a plurality of processes selected to be run on the APD by the SWS. The run-list can be implemented in the hardware or firmware, and can be managed by the APD or an associated hardware scheduler (HWS). According to an embodiment, whereas the SWS selects the processes to be input to the run-list, the HWS can select the process to be run on the APD from those in the run-list. The selection of the next process to be run on the APD can be based upon a round-robin or other selection discipline.

- [0115] Upon being signaled by the OS to initiate a context switch on the APD, the SWS first signals the APD to preempt the current process which caused the page fault. According to an embodiment, the SWS first signals the APD to stop running the current process. The SWS next signals the APD to remove the current process from the run-list, and to save the context of the current process to system memory. The SWS may also provide an address in system memory to which the context of the current process is to be saved. Stopping the current process from running on the shader core, removing it from the run-list, and saving its context completes the preemption of the current process from running on APD. The SWS instructions directing the preemption are received and acted upon by the CP to preempt the current process from running on the shader processor.
- [0116] Having preempted the current process from running on the APD, the SWS selects a second process to run on the APD. According to an embodiment, the SWS selects the second process from the active-list or as a new process to be added to the list of runnable processes. If the second process already has stored context, for example, from a previous execution, then the SWS signals the APD to restore the context for the second process. The SWS can then signal the APD to add the second process to the run-list managed by the HWS. When the HWS selects the second process to run from the run-list, the CP will dispatch the second process and restore any context necessary for the execution of the second process on the APD.
- [0117] The Summary and Abstract sections may set forth one or more but not all exemplary embodiments of the present invention as contemplated by the inventor(s), and thus, are not intended to limit the present invention and the appended claims in any way.
- [0118] The present invention has been described above with the aid of functional building blocks illustrating the implementation of specified functions and relationships thereof. The boundaries of these functional building blocks have been arbitrarily defined herein for the convenience of the description. Alternate boundaries can be defined so long as the specified functions and relationships thereof are appropriately performed.
- [0119] The foregoing description of the specific embodiments will so fully reveal the general nature of the invention that others can, by applying knowledge within the skill of the art, readily modify and/or adapt for various applications such specific embodiments, without undue experimentation, without departing from the general concept of the present invention. Therefore, such adaptations and modifications are intended to be within the meaning and range of equivalents of the disclosed embodiments, based on the

teaching and guidance presented herein. It is to be understood that the phraseology or terminology herein is for the purpose of description and not of limitation, such that the terminology or phraseology of the present specification is to be interpreted by the skilled artisan in light of the teachings and guidance.

[0120] The breadth and scope of the present invention should not be limited by any of the above-described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents.

WHAT IS CLAIMED IS:

1. A method, comprising:
responsive to an exception upon access to a memory by a process running on a accelerated processing device, determining whether to preempt the process based on the exception; and
preempting, based upon the determining, the process from running on the accelerated processing device.
2. The method of claim 1, wherein the determining whether to preempt the process is performed by an operating system.
3. The method of claim 1, further comprising:
requesting, using a memory management unit, data from the memory;
determining, using the memory management unit, that the data is absent from an accessible area of the memory;
receiving notification of the absence; and
generating an interrupt associated with the absence.
4. The method of claim 3, further comprising:
queuing an event indicating the exception to the operating system.
5. The method of claim 4, further comprising:
requesting, by the accelerated processing device, fault handling associated with the exception from the memory management unit.
6. The method of claim 5, further comprising:
receiving, from an operating system, a signal indicating a status regarding the queued event.
7. The method of claim 1, wherein the determining comprises:
determining a type of the exception; and
selecting to preempt or stall the process based upon the determined type.

8. The method of claim 7, wherein:
the determining further comprises accessing statistics associated with memory exceptions; and
the selecting to preempt or stall is further based upon the accessed statistics.
9. The method of claim 8, wherein the accessed statistics include aggregate statistics for the determined type.
10. The method of claim 3, wherein the determining comprises:
notifying a kernel mode driver associated with the accelerated processing device about the absence; and
determining by the kernel mode driver whether to preempt or stall the process.
11. The method of claim 10, wherein the determining by the kernel mode driver comprises:
accessing operating system maintained statistics associated with exceptions; and
selecting between stalling or preempting the process based upon the accessed statistics.
12. A system comprising:
a memory;
an accelerated processing device coupled to the memory and configured to:
run at least one process on the accelerated processing device; and
responsive to one or more instructions, preempt the at least one process from running on the accelerated processing device.
13. The system of claim 12, further comprising:
an operating system configured to:
responsive to an exception upon access to the memory by the at least one process, determine whether to preempt the at least one process based on the detected exception; and
generate the one or more instructions to preempt the at least one process.
14. The system of claim 13, further comprising:

a memory management unit configured to:

- receive a request for data;
- determine that the data is absent from the memory; and
- generate an interrupt associated with the absence.

15. The system of claim 14, further comprising:
 - a translation lookahead buffer (TLB) coupled to the memory management unit and configured to return a TLB miss if the requested data is not present in the memory.
16. The system of claim 14, further comprising:
 - at least one central processing unit coupled to the memory and configured to run one or more processes to initiate the process on the accelerated processing device.
17. The system of claim 16, further comprising:
 - a scheduling module configured to:
 - run on the central processing unit;
 - send instructions from the central processing unit to the accelerated processing device; and
 - receive notification from the operating system to preempt or context-switch the process running on the accelerated processing device.
18. The system of claim 17, wherein the scheduling module includes at least one of a kernel mode driver, a hardware scheduler (HWS), and a software scheduler (SWS).
19. The system of claim 17, wherein the scheduling module is further configured to:
 - determine whether to preempt or stall the process based upon memory exception information maintained by the operating system.
20. The system of claim 19, wherein the scheduling module is further configured to:
 - initiate preemption of the process from the accelerated processing device.
21. The system of claim 20, wherein the scheduling module is further configured to:
 - initiate context-switching a second process on the accelerated processing device.

22. A computer readable medium storing instructions, wherein the instructions, if executed, cause a method comprising:
responsive to an exception upon access to a memory by a process running on a accelerated processing device, determining by an operating system whether to preempt the process; and
preempting the process from running on the accelerated processing device based upon the determining.
23. The computer readable medium of claim 22, wherein the determining by the operating system comprises:
determining a type of the exception; and
selecting to preempt or stall the process based upon the determined type.
24. A method comprising:

receiving a signal in an operating system indicating a memory exception associated with a process running on an accelerated processing device; and

transmitting, by the operating system, one or more instructions to preempt the process from running on the accelerated processing device.
25. A method, comprising:
detecting an exception upon access to a memory by a process running on a accelerated processing device;

receiving one or more instructions from an operating system in response to the detected exception; and

preempting the process from running on the accelerated processing device based on the received one or more instructions.

1/6




FIG. 1A

2/6




FIG. 1B

3/6


200

FIG. 2

4/6




FIG. 3




FIG. 4




FIG. 5

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2011/064550

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F9/48
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 practical, search terms used)

EP0-Internal, IBM-TDB, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2003/005251 A1 (WILSON KENNETH M [US] ET AL) 2 January 2003 (2003-01-02) paragraph [0015] - paragraph [0019] paragraph [0024] - paragraph [0026] -----	1-25
X	US 6 016 542 A (GOTTLIEB ROBERT STEVEN [US] ET AL) 18 January 2000 (2000-01-18) column 1, line 28 - column 2, line 2 column 3, line 44 - line 60 -----	1-25
X	US 6 049 867 A (EICKEMEYER RICHARD JAMES [US] ET AL) 11 April 2000 (2000-04-11) column 3, line 6 - line 10 column 4, line 27 - line 55 -----	1-25

☐

Further documents are listed in the continuation of Box C.

☒

See patent family annex.

* Special categories of cited documents :

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

"E" earlier document 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

5 April 2012

Date of mailing of the international search report

07/05/2012

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

Milasinovic, Goran

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2011/064550

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2003005251 A1	02-01-2003	US 2003005251 A1	02-01-2003
		US 2003005252 A1	02-01-2003
		US 2003005257 A1	02-01-2003

US 6016542 A	18-01-2000	NONE	

US 6049867 A	11-04-2000	EP 0747816 A2	11-12-1996
		JP 9006633 A	10-01-1997
		US 6049867 A	11-04-2000
