




- (51) **International Patent Classification:**
G06F 9/455 (2006.01)
- (21) **International Application Number:**
PCT/US2016/066375
- (22) **International Filing Date:**
13 December 2016 (13.12.2016)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
14/994,393 13 January 2016 (13.01.2016) US
- (71) **Applicant:** INTEL CORPORATION [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95054 (US).
- (72) **Inventors:** SANKARAN, Rajesh M.; 15132 NW Vance Drive, Portland, Oregon 97229 (US). CAMPBELL, Randolph L.; 904 Halidon Way, Folsom, California 95630 (US). SETHI, Prashant; 235 Silberhorn Drive, Folsom, California 95630 (US). HARRIMAN, David J.; 2845 NW Cumberland Rd., Portland, Oregon 97210 (US).
- (74) **Agent:** LANE, Thomas R.; Nicholson De Vos Webster & Elliott LLP, 217 High Street, Palo Alto, California 94301 (US).

- (81) **Designated States** (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, 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).

Published:
— with international search report (Art. 21(3))

(54) **Title:** ADDRESS TRANSLATION FOR SCALABLE VIRTUALIZATION OF INPUT/OUTPUT DEVICES

FIGURE 3



(57) **Abstract:** Embodiments of an invention for address translation for scalable I/O device virtualization are disclosed. In one embodiment, an apparatus includes PASID table lookup circuitry. The PASID table lookup circuitry is to find a PASID-entry in a PASID table. The PASID-entry is to include a PASID processing mode (PPM) indicator and a first pointer to a first translation structure. The PPM indicator is to specify one of a plurality of translation types, the one of the plurality of translation types to use the first translation structure.

ADDRESS TRANSLATION FOR SCALABLE VIRTUALIZATION OF INPUT/OUTPUT DEVICES

BACKGROUND

1. Field

[0001] The present disclosure pertains to the field of information processing, and more particularly, to the field of virtualization in information processing systems.

2. Description of Related Art

[0002] Generally, the concept of virtualization in information processing systems allows multiple instances of one or more operating systems (each, an OS) to run on a single information processing system, even though each OS is designed to have complete, direct control over the system and its resources. Virtualization is typically implemented by using software (e.g., a virtual machine monitor (VMM) or hypervisor) to present to each OS a virtual machine (VM) having virtual resources, including one or more virtual processors, that the OS may completely and directly control, while the VMM maintains a system environment for implementing virtualization policies such as sharing and/or allocating the physical resources among the VMs (the virtual environment).

BRIEF DESCRIPTION OF THE FIGURES

[0003] The present invention is illustrated by way of example and not limitation in the accompanying figures.

[0004] Figure 1 illustrates a system including address translation for scalable virtualization of input/output devices according to an embodiment of the present invention.

[0005] Figure 2 illustrates a processor in the system of Figure 1.

[0006] Figure 3 illustrates an input/output memory management unit according to an embodiment of the present invention.

[0007] Figure 4 illustrates a PASID-entry format according to an embodiment of the present invention.

[0008] Figure 5 illustrates a method for address translation for scalable virtualization of input/output devices according to an embodiment of the present invention.

DETAILED DESCRIPTION

[0009] Embodiments of an invention for address translation for scalable virtualization of input/output devices are described. In this description, numerous specific details, such as component and system configurations, may be set forth in order to provide a more thorough

understanding of the present invention. It will be appreciated, however, by one skilled in the art, that the invention may be practiced without such specific details. Additionally, some well-known structures, circuits, and other features have not been shown in detail, to avoid unnecessarily obscuring the present invention.

[0010] In the following description, references to “one embodiment,” “an embodiment,” “example embodiment,” “various embodiments,” etc., indicate that the embodiment(s) of the invention so described may include particular features, structures, or characteristics, but more than one embodiment may and not every embodiment necessarily does include the particular features, structures, or characteristics. Further, some embodiments may have some, all, or none of the features described for other embodiments.

[0011] As used in this description and the claims and unless otherwise specified, the use of the ordinal adjectives “first,” “second,” “third,” etc. to describe an element merely indicate that a particular instance of an element or different instances of like elements are being referred to, and is not intended to imply that the elements so described must be in a particular sequence, either temporally, spatially, in ranking, or in any other manner.

[0012] Also, the terms “bit,” “flag,” “field,” “entry,” “indicator,” etc., may be used to describe any type or of or content of a storage location in a register, table, database, or other data structure, whether implemented in hardware or software, but are not meant to limit embodiments of the invention to any particular type of storage location or number of bits or other elements within any particular storage location. The term “clear” may be used to indicate storing or otherwise causing the logical value of zero to be stored in a storage location, and the term “set” may be used to indicate storing or otherwise causing the logical value of one, all ones, or some other specified value to be stored in a storage location; however, these terms are not meant to limit embodiments of the present invention to any particular logical convention, as any logical convention may be used within embodiments of the present invention.

[0013] Also, as used in descriptions of embodiments of the present invention, a “/” character between terms may mean that an embodiment may include or be implemented using, with, and/or according to the first term and/or the second term (and/or any other additional terms).

[0014] As described in the background section, information processing systems may provide for virtualization of resources. These resources may include input/output (I/O) devices, where the definition of I/O devices may include controllers (such as bus,

interconnect, and/or fabric controllers, network interface controllers (NICs), etc.), processors/accelerators (such as graphics processors, cryptographic and/or compression accelerators, etc.), and any other resources that may be addressed in the system. Various approaches to and usages of virtualization have been and continue to be developed, including those with many (hundreds to thousands or more) of VMs or other containers (e.g., OS-managed separate and/or isolated execution environments), where any number of I/O devices may be used and/or shared at various times by any number of VMs or containers. Therefore, embodiments of the present invention may be desired to provide for efficient and scalable virtualization of I/O devices.

[0015] Figure 1 illustrates system 100, an information processing system including an embodiment of the present invention for address translation for scalable virtualization of I/O devices. System 100 may represent any type of information processing system, such as a server, a desktop computer, a portable computer, a set-top box, a hand-held device such as a tablet or a smart phone, or an embedded control system. System 100 includes processor 112, memory controller 114, host fabric controller 116, I/O controller 140, I/O memory management unit (IOMMU) 142, system memory 120, graphics processor 130, and hardware accelerator 150. Systems embodying the present invention may include any number of each of these components and any other components or other elements, such as additional peripherals and/or I/O devices. Any or all of the components or other elements in this or any system embodiment may be connected, coupled, or otherwise in communication with each other through any number of buses, point-to-point, or other wired or wireless interfaces or connections, unless specified otherwise. Any components or other portions of system 100, whether shown in Figure 1 or not shown in Figure 1, may be integrated or otherwise included on or in a single chip (a system-on-a-chip or SOC), die, substrate, or package, such as SOC 110.

[0016] System memory 120 may be dynamic random access memory (DRAM) or any other type of medium readable by processor 112. Memory controller 114 may represent any circuitry or component for accessing, maintaining, and/or otherwise controlling system memory 120. Host fabric controller 116 may represent any circuitry or component for controlling an interconnect network or fabric through which processors and/or other system components may communicate. Graphics processor 130 may include any processor or other component for processing graphics data for display 132. Hardware accelerator 150 may represent any cryptographic, compression, or other accelerator to which

a processor may offload functionality such as the hardware acceleration of encryption or compression algorithms.

[0017] I/O controller 140 may represent any circuitry or component, such as a chipset component, including or through which peripheral, input/output (I/O), or other components or devices, such as I/O device 144 (e.g., a touchscreen, keyboard, microphone, speaker, other audio device, camera, video or other media device, motion or other sensor, receiver for global positioning or other information, etc.), NIC 146, and/or information storage device 148, may be connected or coupled to processor 112. Information storage device 148 may represent any one or more components including any one more types of persistent or non-volatile memory or storage, such as a flash memory and/or a solid state, magnetic, or optical disk drive, and may include its own controller 149.

[0018] Processor 112 may represent all or part of a hardware component including one or more processors or processor cores integrated on a single substrate or packaged within a single package, each of which may include multiple execution threads and/or multiple execution cores, in any combination. Each processor represented as or in processor 112 may be any type of processor, including a general purpose microprocessor, such as a processor in the Intel® Core® Processor Family or other processor family from Intel® Corporation or another company, a special purpose processor or microcontroller, or any other device or component in an information processing system in which an embodiment of the present invention may be implemented. Processor 112 may be architected and designed to operate according to any instruction set architecture (ISA), with or without being controlled by microcode.

[0019] Processor 112 may support virtualization according to any approach. For example, processor 112 may operate in two modes – a first mode in which software runs directly on the hardware, outside of any virtualization environment, and a second mode in which software runs at its intended privilege level, but within a virtual environment hosted by a VMM running in the first mode. In the virtual environment, certain events, operations, and situations, such as interrupts, exceptions, and attempts to access privileged registers or resources, may be intercepted, i.e., cause the processor to exit the virtual environment (a VM exit) so that the VMM may operate, for example, to implement virtualization policies. The processor may support instructions for establishing, entering (a VM entry), exiting, and maintaining a virtual environment, and may include register bits or other structures that indicate or control virtualization capabilities of the processor.

[0020] Figure 2 illustrates processor 200, which may represent an embodiment of processor 112 in Figure 1 or an execution core of a multicore processor embodiment of processor 112 in Figure 1. Processor 200 may include storage unit 210, instruction unit 220, execution unit 230, control unit 240, and memory management unit (MMU) 250. Processor 200 may also include any other circuitry, structures, or logic not shown in Figure 1.

[0021] Storage unit 210 may include any combination of any type of storage usable for any purpose within processor 200; for example, it may include any number of readable, writable, and/or read-writable registers, buffers, and/or caches, implemented using any memory or storage technology, in which to store capability information, configuration information, control information, status information, performance information, instructions, data, and any other information usable in the operation of processor 200, as well as circuitry usable to access such storage and/or to cause or support various operations and/or configurations associated with access to such storage.

[0022] Instruction unit 220 may include any circuitry, logic, structures, and/or other hardware, such as an instruction decoder, to fetch, receive, decode, interpret, schedule, and/or handle instructions to be executed by processor 200. Any instruction format may be used within the scope of the present invention; for example, an instruction may include an opcode and one or more operands, where the opcode may be decoded into one or more micro-instructions or micro-operations for execution by execution unit 230. Operands or other parameters may be associated with an instruction implicitly, directly, indirectly, or according to any other approach.

[0023] Execution unit 230 may include any circuitry, logic, structures, and/or other hardware, such as arithmetic units, logic units, floating point units, shifters, etc., to process data and execute instructions, micro-instructions, and/or micro-operations. Execution unit 230 may represent any one or more physically or logically distinct execution units.

[0024] Control unit 240 may include any microcode, firmware, circuitry, logic, structures, and/or hardware to control the operation of the units and other elements of processor 200 and the transfer of data within, into, and out of processor 200. Control unit 240 may cause processor 200 to perform or participate in the performance of method embodiments of the present invention, such as the method embodiments described below, for example, by causing processor 200, using execution unit 230 and/or any other resources, to execute instructions received by instruction unit 220 and micro-instructions or micro-operations derived from instructions received by instruction unit 220. The execution of

instructions by execution 230 may vary based on control and/or configuration information stored in storage unit 210.

[0025] MMU 250 may include any circuitry, logic, structures, and/or other hardware to manage system memory, such as providing for the virtualization of physical memory according to any desired approach and the protection of system memory. In an embodiment, MMU 250 supports the use of virtual memory to provide software, including software running in a VM, with an address space for storing and accessing code and data that is larger than the address space of the physical memory in the system, e.g., system memory 120. The virtual memory space of processor 200 may be limited only by the number of address bits available to software running on the processor, while the physical memory space of processor 200 may be limited to the size of system memory 120. MMU 250 supports a memory management scheme, such as paging, to swap the executing software's code and data in and out of system memory 120 on an as-needed basis. As part of this scheme, the software may access the virtual memory space of the processor with an un-translated address that is translated by the processor to a translated address that the processor may use to access the physical memory space of the processor.

[0026] Accordingly, MMU 250 may include translation lookaside buffer 252 in which to store translations of a virtual, logical, linear, or other un-translated address to a physical or other translated address, according to any known memory management technique, such as paging. To perform these address translations, MMU 250 may include page-walk hardware 254 to refer to one or more data structures stored in processor 200, system memory 130, storage locations in system 100 not shown in Figure 1, and/or any combination of these locations. These data structures may include page directories, page tables, and other paging data structures according to any known paging architecture. Each such paging data structure, as well as TLB 252, may include (or have associated with individual or groups of entries) one or more bits or other indicators to be used to indicate and enforce various permissions (e.g., read, write, or execute) that may define or restrict access to pages (or other regions) of memory.

[0027] The virtualization capabilities of processor along with MMU 250 may provide for various approaches to creating and maintaining containers, where a container may be any execution or processing environment, created and maintained by a hypervisor, VMM, OS, or any other system or host software. Any platform, system, or machine, including the "bare metal" platform shown as system 100 in Figure 1 as well as any VM or other container

abstracted from a bare metal platform, from which one or more containers are abstracted may be referred to as a host or host machine, and each VM or other such container abstracted from a host machine may be referred to as a guest or guest machine. Accordingly, the term “host software” may mean any hypervisor, VMM, OS, or any other software that may run, execute, or otherwise operate on a host machine and create, maintain, and/or otherwise manage one or more containers, and the term “guest software” may mean any OS, system, application, user, or other software that may run, execute, or otherwise operate on a guest machine. Note that in a layered container architecture, software may be both host software and guest software. For example, a first VMM running on a bare metal platform may create a first VM, in which a second VMM may run and create a second VM abstracted from the first VM, in which the case the second VMM is both host software and guest software.

[0028] For convenience, the use of the term “container process” may mean any context, task, application, software, privileged process, unprivileged process, kernel-mode process, supervisor-mode process, user-mode process, or any other process running or runnable within a container. A container may have an address space (a container address space or a guest address space) that is different from the system address space (e.g., the address space of system memory 120) or the host address space (e.g. the address space of the host machine). An address with which the system address space may be directly accessed (i.e., without translation) may be referred to as a host physical address (HPA). For isolation, protection, or any other purpose, any container address space may be different from any other container address space. Therefore, each container process may access memory using addresses that are to be translated, filtered, or otherwise processed to HPAs differently than they are translated, filtered, or otherwise processed for any other container. The difference in translation/processing of container addresses may be due to virtualization and isolation of container address spaces (e.g., guest software may use guest virtual addresses (GVAs) that are translated to guest physical address spaces (GPAs) that are translated to HPAs) and may also be due to the use of a variety of different types of containers (e.g., VMs, OS-managed containers, etc.) and/or different container architectures (e.g., layered architectures including VMs hosting multiple VMs, VMs hosting multiple OS-managed containers, etc.).

[0029] An address used by a container process to access memory (a container address) may be any of many different types of addresses, including an HPA, a virtual address, a guest physical address (GPA), a guest virtual address (GVA), a direct memory access (DMA) address, etc., and may go through one or more of any of a variety of techniques, types, levels,

layers, rounds, and/or steps of translation, filtering, and/or processing, in any combination, using any of a variety of data structures (e.g., page tables, extended page table, nested page tables, DMA translation tables, memory access filters, memory type filters, memory permission filters, etc.) to result in an HPA and/or in a fault, error, or any other type of determination that a requested access is not allowed. Various approaches may include layering and/or nesting of containers (e.g., a VMM hosting a VM running a guest OS, the guest OS supporting multiple containers; a VMM hosting multiple VMs each running a guest OS, etc.), involving various combinations of address translation techniques.

[0030] Each physical function within an I/O device in system 100 may become usable and/or shareable by one or more clients (e.g., containers, container processes, host processes, etc.) by reporting to system software the number of “command portals” that it supports, where a command portal is a command interface through which a client may submit a work request to the I/O device (e.g., through a portal driver). For example, a command portal for a NIC (e.g., NIC 146) may be a paired transmit queue and receive queue; a command portal for an InfiniBand, remote DMA (RDMA), or other host fabric controller (e.g., host fabric controller 116) may be a Queue Pair; a command portal for a Non-Volatile Memory Express (NVMe) or other storage device controller (e.g., storage device controller 149) may be a Command Queue; a command portal for a graphics processing unit (GPU), general purpose computing on GPU (GPGPU), or other accelerator (e.g., hardware accelerator 150) may be a schedulable context through which work may be submitted; etc. A command portal may be distinguished from an “admin portal” as being an interface for a client to submit work, whereas an admin portal is an interface through which a container host sets up or configures the command portals.

[0031] An I/O device may report to host software that it supports one or more command portals for use according to embodiments of the present invention, as well as how many command portals it supports, through capability/attribute information that it provides according to a system bus or interconnect specification (e.g., through a new capability added to the Peripheral Component Interconnect Express (PCIe) specification), by a device driver for the physical function, or according to any other known technique for reporting physical function capabilities/attributes.

[0032] Then, the host software may use the I/O device’s admin portal to allocate, map, and/or assign each command portal to a client. This assignment includes assigning a process address space identifier (PASID) to the command portal, where the PASID corresponds to

the address space associated with the client. In an embodiment, a PASID may be a 20-bit tag defined by the PCIe specification and carried by the translation layer packet (TLP) prefix header in transactions generated by the I/O device. After the assignment of a command portal has been completed, clients may submit work requests to it according to any known approach.

[0033] Figure 3 illustrates IOMMU 300 for address translation for scalable virtualization of I/O devices according to an embodiment of the present invention. An IOMMU such as IOMMU 300 may represent an embodiment of IOMMU 142 in Figure 1 and/or may be integrated into an I/O controller such as I/O controller 140, a processor such as processor 112, and/or an SOC such as SOC 110. An IOMMU according to embodiments of the present invention supports the efficient and scalable assignment of a command portal of any type of I/O device to any type of client by providing for any of a variety of types of client address to HPA address translations to be performed, where the type of translation to be performed is determined based on a PASID associated with an I/O transaction, as further described below.

[0034] IOMMU 300 may include root table lookup circuitry 310 to find an entry in a root table, such as root table 314. Root table 314 may be a data structure in local memory of IOMMU 300, in system memory 120, and/or elsewhere in system 100, in which one or more root-entries (e.g., root-entry 316) may be stored. Root table 314 may have a root table base address stored in a root table pointer storage location 312. Root table lookup circuitry 310 may include circuitry to add a root table index to the root table base address, where the root table index may correspond to a bus number in an embodiment in which an identifier of an I/O device may include a bus number. For example, an I/O device in a system including a PCIe bus may be assigned a bus number, device number, and function number (BDF) by system configuration software or firmware.

[0035] IOMMU 300 may also include context table lookup circuitry 320 to find an entry in a context table, such as context table 324. Context table 324 may be a data structure in local memory of IOMMU 300, in system memory 120, and/or elsewhere in system 100, in which one or more context-entries (e.g., context-entry 326) may be stored. Embodiments of the present invention may include any number of context tables, each with any number of context-entry locations. Context table 324 may have a context table base address stored in a root-entry location (e.g., root-entry 316). Context table lookup circuitry 320 may include circuitry to add a context table index to the context table base address, where the context

table index may correspond to a device and function number in an embodiment in which an identifier of an I/O device may include a device and function number (e.g., a PCIe BDF).

[0036] IOMMU 300 may also include PASID table lookup circuitry 330 to find an entry in a PASID table, such as PASID table 334. PASID table 334 may be a data structure in local memory of IOMMU 300, in system memory 120, and/or elsewhere in system 100, in which one or more PASID-entries (e.g., PASID-entry 336) may be stored. Embodiments of the present invention may include any number of PASID tables, each with any number of PASID-entry locations. PASID table 334 may have a PASID table base address stored in a context-entry location (e.g., context-entry 326). PASID table lookup circuitry 330 may include circuitry to add a PASID table index to the PASID table base address, where the PASID table index may correspond to a PASID.

[0037] Each PASID-entry may include any number of fields to facilitate translation of any of a variety of address types that may be carried by or otherwise associated with a transaction initiated by a command portal. Figure 4 illustrates PASID-entry format 400 according to an embodiment of the present invention.

[0038] PASID-entry 400 includes first-level translation pointer 410, second-level translation pointer 420, and PASID processing mode (PPM) field 430. First-level translation pointer 410 may be a pointer to a first-level translation data structure, e.g., a first hierarchy of page tables stored in system memory 120. Second-level translation pointer 420 may be a pointer to a second-level translation data structure, e.g., a second hierarchy of page tables stored in system memory 120. PPM field may store a PPM indicator to specify which of a variety of modes of translation is to be performed. For example, a first mode of translation may use only the first-level translation data structure, a second mode of translation may use only the second-level translation data structure, and a third mode may use both the first-level and the second-level translation data structures according to a nested or extended translation technique. Therefore, host software may program the PPM field for each PASID-entry, such that a variety of translation techniques may be used on a per-PASID granularity, and any command portal may be assigned to any of variety of types of client, each client using any of a variety of types of addresses (e.g., restricted HPA, GPA, virtual, GVA, DMA, etc.)

[0039] Therefore, a PASID carried or otherwise associated with a transaction by an I/O device in system 100 may be used by IOMMU 300 to initiate any type of a variety of types of address translation. The I/O transaction may be a memory read, a memory write, an atomic memory read/write, or any other transaction that includes accessing system memory.

[0040] As a first example, an I/O device may be a user-mode capable device allocated to operate in shared virtual memory space with host applications, the I/O device having a command portal to which a first PASID has been assigned. The command portal may initiate an I/O transaction including a memory address to be accessed, an identifier of the I/O device (e.g., a BDF) and the first PASID, from which the IOMMU may find a first PASID-entry. The first PASID-entry may include a pointer to a first-level translation structure that is a page table hierarchy to translate virtual addresses (in the virtual address space of processor 112) to HPAs. The first PASID-entry may also include a PPM specifying that only a first-level translation is to be performed. Therefore, the memory address included in or associated with the I/O transaction may be a virtual address.

[0041] As a second example, an I/O device may have a command portal allocated to a VM and to which a second PASID has been assigned. The command portal may initiate an I/O transaction including a memory address to be accessed, an identifier of the I/O device (e.g., a BDF) and the second PASID, from which the IOMMU may find a second PASID-entry. The second PASID-entry may include a pointer to a second-level translation structure that is a page table hierarchy to translate GPAs (in the physical address space of the VM) to HPAs. The second PASID-entry may also include a PPM specifying that only a second-level translation is to be performed. Therefore, the memory address included in or associated with the I/O transaction may be a GPA.

[0042] As a third example, an I/O device may have a command portal allocated to an OS-managed container in a VM and to which a third PASID has been assigned. The command portal may initiate an I/O transaction including a memory address to be accessed, an identifier of the I/O device (e.g., a BDF) and the third PASID, from which the IOMMU may find a third PASID-entry. The third PASID-entry may include a first pointer to a first-level translation structure that is a page table hierarchy to translate virtual addresses (in the virtual address space of the container) to GPAs (in the physical address space of the VM) and second pointer to a second-level translation structure that is a page table hierarchy to translate GPAs (in the physical address space of the VM) to HPAs. The third PASID-entry may also include a PPM specifying that an extended or nested translation, using both the first-level and second-level translation structures, is to be performed. Therefore, the memory address included in or associated with the I/O transaction may be a GVA.

[0043] As a fourth example, an I/O device may have a command portal allocated to operate within a restricted HPA space, and to which a fourth PASID has been assigned. The

command portal may initiate an I/O transaction including a memory address to be accessed, an identifier of the I/O device (e.g., a BDF) and the fourth PASID, from which the IOMMU may find a fourth PASID-entry. The fourth PASID-entry may include a pointer to a translation structure that is a filter to restrict access to a range or ranges of HPAs. The fourth PASID-entry may also include a PPM specifying that a translation is to be performed using the HPA filter. Therefore, the memory address included in or associated with the I/O transaction may be an HPA.

[0044] As a fifth example, an I/O device may have a command portal allocated to operate within an OS-based container having a DMA address space mapped to the HPA address space, and to which a fifth PASID has been assigned. The command portal may initiate an I/O transaction including a memory address to be accessed, an identifier of the I/O device (e.g., a BDF) and the fifth PASID, from which the IOMMU may find a fifth PASID-entry. The fifth PASID-entry may include a pointer to a translation structure that is a DMA address to HPA mapping structure. The fifth PASID-entry may also include a PPM specifying that a translation is to be performed using the DMA mapping structure. Therefore, the memory address included in or associated with the I/O transaction may be a DMA address.

[0045] Embodiments of the present invention may support backwards compatibility by with a translation-type field in context-entries, where the translation-type field may be configured by host software to specify a translation type. A first translation type (e.g., scalable I/O virtualization) may be used to enable the IOMMU to operate according to embodiments of the present invention, for example by blocking I/O transactions that do not carry or include a PASID. A second translation type (e.g., single-root I/O virtualization or SR-IOV) may be used to enable the IOMMU to operate according to a different architecture (e.g., PCIe SR-IOV), for example, by supporting (according to a previously known technique) I/O transactions that do not carry or include a PASID.

[0046] Various embodiments of the present invention may also include, in an IOMMU, support for scalable interrupt resource allocation for each command portal. For example, configuration of a command portal by host software may include assignment of an interrupt message address/data value pair to the command portal. An I/O device may use local or implementation specific storage to store the interrupt message address/data value for each of its command portals. For example, for a GPU context, the interrupt message address/data value may be saved in memory as part of the context supervisory state that is loaded by the GPU in connection with scheduling the GPU context.

[0047] Then, in connection with generated an interrupt (or other event), the command portal may perform a posted write using the address/data value pair, for which host software has configured the IOMMU to generate the appropriate type of interrupt. Interrupt types may include host kernel (kernel interrupt targeting a logical processor), virtual kernel (kernel interrupt targeting a virtual processor), host user (user interrupt targeting a host application thread), virtual user (user interrupt targeting a guest application thread), etc.

[0048] Alternatively, a physical function of an I/O device may support its own table structure in which each entry may be configured as an interrupt resource for one of multiple command portals (i.e., each command portal may use one or more entries in the table). Host software may program each entry with an interrupt message address/data value pair to specify the type of interrupt.

[0049] Software to support embodiments of the present invention may include VMMs or other system software, scalable I/O managers, device drivers, and command portal drivers. A VMM or other system software may create, maintain, and/or manage containers; map I/O device resources to clients; and program IOMMU resources (including translation structures and interrupt resources) according to requests from a scalable I/O manager.

[0050] A scalable I/O manager may operate as a layer or other part of the system software to detect, allocate, and configure scalable I/O resources, according to embodiments of the present invention, and allocate command portals to clients. Its interactions with devices are abstracted through interfaces with device drivers.

[0051] A device driver may be associated with each physical function of an I/O device and may be loaded/executed as part of the host OS or VMM I/O software stack. It may accept requests from system software or a scalable I/O manager software stack to help configure the device to support scalable I/O. For example, it may execute requests for enumeration and/or configuration of command portals by programming the device through an admin portal. For example, it may execute, through the admin portal, a request to assign a PASID and an interrupt message address/data pair value to a command portal.

[0052] A command portal driver may be implemented as a kernel-mode driver for a kernel clients or a user-mode driver for user clients to which a command portal may be allocated and mapped. It operates with awareness of the specific hardware/software interface of the command portal, the work request completion schemes supported by the device (e.g., kernel interrupts, user interrupts, semaphores, etc.), any shared work queues, and/or any queuing instructions supported by the device.

[0053] Figure 5 illustrates method 500 for address translation for scalable virtualization of I/O devices according to an embodiment of the present invention. Although method embodiments of the invention are not limited in this respect, reference may be made to elements of Figures 1, 2, 3, and 4 to help describe the method embodiment of Figure 5. Various portions of method 500 may be performed by hardware, firmware, software, and/or a user of a system such as system 100.

[0054] In box 510 of method 500, host software may create one or more containers on an information processing system (e.g., system 100).

[0055] In box 520, any number of command portals supported by any number of I/O devices in the system may be configured, including, in box 522, assigning a PASID to each command portal and, in box 524, assigned an interrupt message address/data value pair to each command portal.

[0056] In box 530, an IOMMU (e.g., IOMMU 300) may be configured, including, in box 532, programming any number of root-entries; in box 534, programming any number of context-entries; in box 536, programming any number of PASID-entries. In box 538, execution of any one or more processes on system 100 continues.

[0057] In box 540, a first memory access transaction may be initiated by a first command portal. In box 542, the IOMMU may use a first requester identifier and a first PASID from the first transaction to find a first PASID-entry. In box 544, the IOMMU may use a first PPM in the first PASID-entry to determine that a first address provided by the first transaction is to be translated using a first translation structure pointed to by a first pointer in the first PASID-entry, where the first address is a virtual address and the first translation structure is a page table hierarchy to translate virtual addresses to HPAs. In box 546, the translation is performed. In box 548, execution of any one or more processes on system 100 continues.

[0058] In box 550, a second memory access transaction may be initiated by a second command portal. In box 552, the IOMMU may use a second requester identifier and a second PASID from the second transaction to find a second PASID-entry. In box 554, the IOMMU may use a second PPM in the second PASID-entry to determine that a second address provided by the second transaction is to be translated using a second translation structure pointed to by a second pointer in the second PASID-entry, where the second address is a GPA and the second translation structure is an extended page table hierarchy to

translate GPAs to HPAs. In box 556, the translation is performed. In box 558, execution of any one or more processes on system 100 continues.

[0059] In box 560, a third memory access transaction may be initiated by a third command portal. In box 562, the IOMMU may use a third requester identifier and a third PASID from the third transaction to find a third PASID-entry. In box 564, the IOMMU may use a third PPM in the third PASID-entry to determine that a third address provided by the third transaction is to be translated according to a nested approach, using a third translation structure and a fourth translation structure, pointed to by a third pointer and a fourth pointer, respectively, in the third PASID-entry, where the third address is a GVA and the third translation structure is a page table hierarchy to translate GVAs to GPAs and the fourth translation structure is a nested or extended page table hierarchy to translate GPAs to HPAs. In box 566, the translation is performed. In box 568, execution of any one or more processes on system 100 continues.

[0060] In box 570, a fourth memory access transaction may be initiated by a fourth command portal. In box 572, the IOMMU may use a fourth requester identifier and a fourth PASID from the fourth transaction to find a fourth PASID-entry. In box 574, the IOMMU may use a fourth PPM in the fourth PASID-entry to determine that a fourth address provided by the fourth transaction is to be translated using a fifth translation structure pointed to by a fifth pointer in the fourth PASID-entry, where the fourth address is an HPA and the fifth translation structure is an HPA filter to restrict access to a range of HPA addresses. In box 576, the filtering is performed. In box 578, execution of any one or more processes on system 100 continues.

[0061] In box 580, a fifth memory access transaction may be initiated by a fifth command portal. In box 582, the IOMMU may use a fifth requester identifier and a fifth PASID from the fifth transaction to find a fifth PASID-entry. In box 584, the IOMMU may use a fifth PPM in the fifth PASID-entry to determine that a fifth address provided by the fifth transaction is to be translated using a sixth translation structure pointed to by a sixth pointer in the fifth PASID entry, where the fifth address is a DMA address and the sixth translation structure is a DMA map to translate DMA addresses to HPAs. In box 586, the translation is performed. In box 588, execution of any one or more processes on system 100 continues.

[0062] In various embodiments of the present invention, the method illustrated in Figure 4 may be performed in a different order, with illustrated boxes combined or omitted, with additional boxes added, or with a combination of reordered, combined, omitted, or additional

boxes. Furthermore, method embodiments of the present invention are not limited to method 400 or variations thereof. Many other method embodiments (as well as apparatus, system, and other embodiments) not described herein are possible within the scope of the present invention.

[0063] Embodiments or portions of embodiments of the present invention, as described above, may be stored on any form of a machine-readable medium. For example, all or part of method 400 may be embodied in software or firmware instructions that are stored on a medium readable by a processor, which when executed by a processor, cause the processor to execute an embodiment of the present invention. Also, aspects of the present invention may be embodied in data stored on a machine-readable medium, where the data represents a design or other information usable to fabricate all or part of a processor or other component.

[0064] Thus, embodiments of an invention for address translation for scalable I/O device virtualization have been described. While certain embodiments have been described, and shown in the accompanying drawings, it is to be understood that such embodiments are merely illustrative and not restrictive of the broad invention, and that this invention not be limited to the specific constructions and arrangements shown and described, since various other modifications may occur to those ordinarily skilled in the art upon studying this disclosure. In an area of technology such as this, where growth is fast and further advancements are not easily foreseen, the disclosed embodiments may be readily modifiable in arrangement and detail as facilitated by enabling technological advancements without departing from the principles of the present disclosure or the scope of the accompanying claims.

CLAIMS

What is claimed is:

1. An apparatus comprising:
PASID table lookup circuitry to find a PASID-entry in a PASID table,
the PASID-entry to include a PASID processing mode (PPM) indicator and a first pointer to a first translation structure;
wherein the PPM indicator is to specify one of a plurality of translation types,
the one of the plurality of translation types to use the first translation structure.
2. The apparatus of claim 1, further comprising context table lookup circuitry to find a context-entry in a context table, the context-entry to point to the PASID table.
3. The apparatus of claim 2, further comprising root table lookup circuitry to find a root-entry in a root table, the root-entry to point to the context table.
4. The apparatus of claim 1, wherein the PASID table lookup logic is to use a PASID from a memory access transaction to find the PASID entry.
5. The apparatus of claim 4, wherein the one of the plurality of translation types is to be performed on an address from the memory access transaction.
6. The apparatus of claim 1, wherein the one of the plurality of translation types is a virtual address to a host physical address translation.
7. The apparatus of claim 1, wherein the one of the plurality of translation types is a guest physical address to a host physical address translation.
8. The apparatus of claim 1, wherein the one of the plurality of translation types is a host physical address filtering.
9. The apparatus of claim 1, wherein the one of the plurality of translation types is a direct memory access address to a host physical address translation.
10. The apparatus of claim 1, wherein:
the PASID-entry is also to include a second pointer to a second translation structure;
and the one of the plurality of translation types to use the first translation structure and the second translation structure.
11. The apparatus of claim 10, wherein the one of the plurality of translation types is a guest virtual address to a host physical address translation.

12. A method comprising:

finding a PASID-entry in a PASID table, the PASID-entry to include a PASID processing mode (PPM) indicator and a first pointer to a first translation structure, wherein the PPM indicator is to specify one of a plurality of translation types,

performing the one of the plurality of translation types using the first translation structure.

13. The method of claim 12, further comprising finding a context-entry in a context table, the context-entry to point to the PASID table.

14. The method of claim 13, further comprising finding a root-entry in a root table, the root-entry to point to the context table.

15. The method of claim 12, wherein the one of the plurality of translation types is a virtual address to a host physical address translation.

16. The method of claim 12, wherein the one of the plurality of translation types is a guest physical address to a host physical address translation.

17. The method of claim 12, wherein the one of the plurality of translation types is a host physical address filtering.

18. The method of claim 12, wherein the one of the plurality of translation types is a direct memory access address to a host physical address translation.

19. The method of claim 12, wherein:

the PASID-entry is also to include a second pointer to a second translation structure;

the one of the plurality of translation types to use the first translation structure and the second translation structure; and

the one of the plurality of translation types is a guest virtual address to a host physical address translation.

20. A system comprising:

a system memory;

an input/output device to initiate a transaction to the system memory; and

an input/output memory management unit, including PASID table lookup circuitry to find, using a PASID from the transaction, a PASID-entry in a PASID table, the PASID-entry to include a PASID processing mode (PPM) indicator and a first pointer to a first translation structure, wherein the PPM indicator is to specify one of a plurality of translation types, the one of the plurality of translation types to use the first translation structure.

FIGURE 1




FIGURE 2




FIGURE 3





FIGURE 4

PASID-entry 400

first-level pointer 410	second-level pointer 420	PPD 430
-------------------------	--------------------------	---------

FIGURE 5
METHOD 500



INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2016/066375**A. CLASSIFICATION OF SUBJECT MATTER****G06F 9/455(2006.01)i**

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 9/455; G06F 12/02; G06F 17/30; G06F 12/10; G06F 12/00; G06F 12/08

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: address, translation, virtualization, input/output device, PASID, table lookup, entry, indicator, pointer, types, structure

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2011-0022818 A1 (ANDREW G. KEGEL et al.) 27 January 2011 See paragraphs [0007], [0018], [0043]-[0067]; and figure 5.	1-20
Y	US 2006-0075146 A1 (IOANNIS SCHOINAS et al.) 06 April 2006 See paragraphs [0054], [0062]-[0066], [0101], [0120]; and figure 5.	1-20
A	US 2012-0246381 A1 (ANDY KEGEL et al.) 27 September 2012 See paragraphs [0029]-[0115]; and figures 1A-7.	1-20
A	US 2013-0080714 A1 (ANDREW G. KEGEL et al.) 28 March 2013 See paragraphs [0020]-[0078]; and figures 1-5.	1-20
A	US 2006-0069899 A1 (IOANNIS SCHOINAS et al.) 30 March 2006 See paragraphs [0025]-[0199]; and figures 1-11B.	1-20



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

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

"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

28 March 2017 (28.03.2017)

Date of mailing of the international search report

29 March 2017 (29.03.2017)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea



Facsimile No. +82-42-481-8578

Authorized officer

CHIN, Sang Bum

Telephone No. +82-42-481-8398



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2016/066375

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2011-0022818 A1	27/01/2011	CN 102498478 A	13/06/2012
		CN 102498478 B	29/04/2015
		EP 2457165 A1	30/05/2012
		JP 05680642 B2	04/03/2015
		JP 2013-500524 A	07/01/2013
		KR 10-1575827 B1	08/12/2015
		KR 10-2012-0044370 A	07/05/2012
		US 9535849 B2	03/01/2017
		WO 2011-011768 A1	27/01/2011
US 2006-0075146 A1	06/04/2006	CN 101031893 A	05/09/2007
		CN 101031893 B	12/05/2010
		DE 112005002304 T5	27/09/2007
		GB 2431757 A	02/05/2007
		GB 2431757 B	10/09/2008
		KR 10-2007-0052334 A	21/05/2007
		TW 200622620 A	01/07/2006
		TW I285315 B	11/08/2007
		US 7444493 B2	28/10/2008
		WO 2006-039051 A1	13/04/2006
US 2012-0246381 A1	27/09/2012	CN 103270502 A	28/08/2013
		EP 2652626 A1	23/10/2013
		JP 2013-546104 A	26/12/2013
		KR 10-2013-0142166 A	27/12/2013
		WO 2012-082864 A1	21/06/2012
US 2013-0080714 A1	28/03/2013	None	
US 2006-0069899 A1	30/03/2006	CN 100593160 C	03/03/2010
		CN 101027653 A	29/08/2007
		DE 112005002298 T5	27/09/2007
		GB 2432443 A	23/05/2007
		GB 2432443 B	02/04/2008
		JP 04772795 B2	14/09/2011
		JP 2008-515093 A	08/05/2008
		TW I321751 B	11/03/2010
		US 2010-0011187 A1	14/01/2010
		US 8843727 B2	23/09/2014
		WO 2006-039057 A2	13/04/2006
		WO 2006-039057 A3	29/06/2006