

(51) International Patent Classification:
G06F 12/10 (2006.01)(21) International Application Number:
PCT/US2010/043169(22) International Filing Date:
24 July 2010 (24.07.2010)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/508,882 24 July 2009 (24.07.2009) US(71) Applicant (for all designated States except US): **ADVANCED MICRO DEVICES, INC.** [US/US]; One AMD Place, P.O. Box 3453, Sunnyvale, California 94088 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **KEGEL, Andrew, G.** [US/US]; 17011 NE 136th Place, Redmond, Washington 98052 (US). **HUMMEL, Mark, D.** [US/US]; 68 Steward St., Franklin, Massachusetts 02038 (US).(74) Agent: **MEYERTONS, HOOD, KIVLIN, KOWERT & GOETZEL, P.C.**; Kivlin, B. Noel, P.O. Box 398, Austin, Texas 78767-0398 (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, RO, RS, RU, 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, 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, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

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

(54) Title: I/O MEMORY MANAGEMENT UNIT INCLUDING MULTILEVEL ADDRESS TRANSLATION FOR I/O AND COMPUTATION OFFLOAD

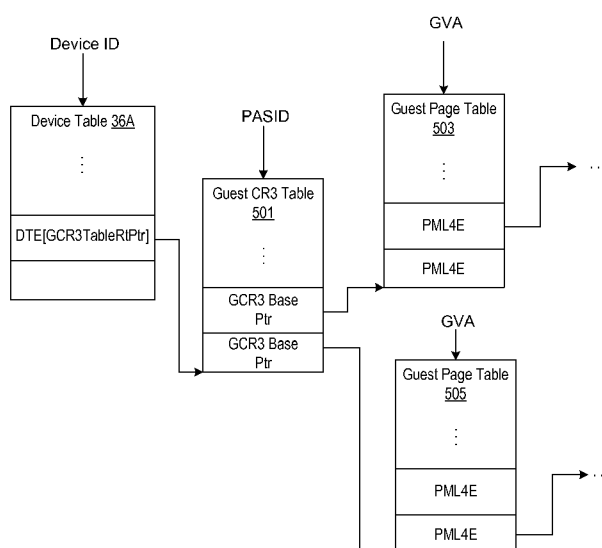


FIG. 5

(57) Abstract: An input/output memory management unit (IOMMU) configured to control requests by an I/O device to a system memory includes control logic that may perform a two-level guest translation to translate an address associated with an I/O device-generated request using translation data stored in the system memory. The translation data includes a device table having a number of entries. The control logic may select the device table entry for a given request by the using a device identifier that corresponds to the I/O device that generates the request. The translation data may also include a first set of I/O page tables including a set of guest page tables and a set of nested page tables. The selected device table entry for the given request may include a pointer to the set of guest translation tables, and a last guest translation table includes a pointer to the set of nested page tables

**TITLE: I/O MEMORY MANAGEMENT UNIT INCLUDING MULTILEVEL
ADDRESS TRANSLATION FOR I/O AND COMPUTATION OFFLOAD**

BACKGROUND OF THE INVENTION

5

Field of the Invention

[0001] This invention is related to the field of computer systems and, more particularly, to memory management units for input/output (I/O) devices.

10 Description of the Related Art

[0002] Computer systems of various types are ubiquitous in modern society, including personal computers (PCs), workstations, servers, various personal digital assistant (PDA) devices, etc. Most, if not all, of these computer systems have implemented memory management functionality for processor accesses to memory. Generally, the memory management
15 functionality has included translating addresses from a virtual address space used by each process to a physical address space that spans the actual system memory, along with various memory protections (e.g. read only, read/write, privilege level requirements, etc.). The memory management functionality has a variety of uses: protecting the memory used by each process from unauthorized access by other processes; permitting large virtual spaces to be used by
20 processes even if the physical memory system is not that large; relocation of virtual addresses to available physical memory without the participation of the process; etc.

[0003] While the processor addresses are frequently translated, addresses used by input/output (I/O) devices in computer systems are generally not translated. That is, the I/O devices use physical addresses to access memory. In a single operating system (OS) computer system, such
25 as most PCs, the OS controls access to the I/O devices by other processes (applications and OS services). Accordingly, the OS can control which process has access to a given device at any given point in time, and can at least somewhat control the addresses accessed by the device. However, such mechanisms become more complicated and cumbersome in virtual machine systems, which may have multiple guest OSs running on a virtual machine monitor.

30 Additionally, devices' use of physical addresses reduces the overall security of the system, since a rogue device (or a device programmed by a malicious software agent) can access memory unimpeded.

[0004] In a virtualized system, many applications may run in a virtual memory while the I/O devices access physical memory. In some systems, a virtual machine manager may intercept and
35 translate addresses that will be used by an I/O device. However, this approach may be

cumbersome and instruction intensive causing the system performance to be degraded, or at least reducing any benefits obtained by using an I/O device such as a computational offload engine.

SUMMARY

5 [0005] Various embodiments of an IOMMU including multilevel address translation are disclosed. In one embodiment, the IOMMU is configured to control requests by an I/O device to a system memory. The IOMMU includes control logic that may perform a two-level guest translation to translate an address associated with an I/O device-generated request using translation data stored in the system memory. The translation data may include a device table
10 having a number of entries. The device table entry for a given request may be selected by the control logic using a device identifier that corresponds to the I/O device that generates the request. The translation data may also include a first set of I/O page tables including a set of guest page tables and a set of nested page tables. The guest page tables include entries that may be similar to control register CR3 entries, for example. The selected device table entry for the
15 given request may include a pointer to the set of guest translation tables, and a last guest translation table includes a pointer to the set of nested page tables. The IOMMU may also include a cache memory that may be configured to store completed translations that may be used for future requests.

[0006] In one specific implementation, the translation data may also include a second set of
20 I/O page tables including a set of host translation page tables. The control logic may perform the two-level translation using the first set of I/O page tables if the I/O request includes a process identifier, and perform a one-level translation using the second set of I/O page tables if the I/O request does not include a process identifier.

25 BRIEF DESCRIPTION OF THE DRAWINGS

[0007] FIG. 1 is a block diagram of a high level view of one embodiment of a computer system.

[0008] FIG. 2 is a block diagram of another embodiment of a computer system.

[0009] FIG. 3 is a diagram depicting one embodiment of a set of I/O translation tables of FIG.
30 1.

[0010] FIG. 4 is a diagram of one embodiment of a device table entry for a device table shown in FIG. 3.

[0011] FIG. 5 is a diagram depicting more detailed aspects of the guest virtual address translation mechanism.

[0012] FIG. 6 is a flow diagram describing the operation of the embodiments of the computer system shown in FIG. 1 through FIG. 5.

[0013] While the invention is susceptible to various modifications and alternative forms, specific embodiments thereof are shown by way of example in the drawings and will herein be described in detail. It should be understood, however, that the drawings and detailed description thereto are not intended to limit the invention to the particular form disclosed, but on the contrary, the intention is to cover all modifications, equivalents, and alternatives falling within the spirit and scope of the present invention as defined by the appended claims. It is noted that the word “may” is used throughout this application in a permissive sense (i.e., having the potential to, being able to), not a mandatory sense (i.e., must).

DETAILED DESCRIPTION

[0014] Turning now to FIG. 1, a block diagram illustrating a simplified, high-level view of one embodiment of a computer system 10 is shown. In the illustrated embodiment, the system 10 includes one or more processors 12, a memory management unit 14 comprising one or more translation lookaside buffers (TLBs) 16, a memory controller (MC) 18, a memory 20, one or more I/O devices 22 which may comprise one or more I/O TLBs (IOTLBs) 24, and an I/O MMU (IOMMU) 26 which may comprise a table walker 28, a cache 30, control registers 32, and control logic 34. The processors 12 are coupled to the MMU 14, which is coupled to the memory controller 18. The I/O devices 22 are coupled to the IOMMU 26, which is coupled to the memory controller 18. Within the IOMMU 26, the table walker 28, the CACHE 30, the control registers 32, and the control unit 34 are coupled together.

[0015] As described further below, the IOMMU 26 may include various features to simplify virtualization in the system 10. The description below will refer to a virtual machine monitor (VMM) that manages the virtual machines (scheduling their execution on the underlying hardware), controls access to various system resources, etc. It is noted that VMMs are also sometimes referred to as hypervisors. In the illustrated embodiment, processor(s) 12 is executing software in a virtualized environment. Accordingly, three virtual machines 100A, 100B, and 100C (e.g., VM guest 1 - 3) and a VMM 106 are shown. The number of virtual machines in a given embodiment may vary, and may dynamically change during use as virtual machines are started and stopped by a user. In the illustrated embodiment, the virtual machine 100A includes one or more guest applications 102 and a guest operating system (OS) 104. The OS 104 is referred to as a “guest” OS, since the OS 104 controls the virtual machine created for it by the VMM 106, rather than the physical hardware of the system 10. Similarly, the VM 100B and VM 100C may also each include one or more guest applications and a guest OS.

[0016] Generally, the applications in the virtual machines use a guest virtual address space and thus, guest virtual addresses (GVA). The guest OS in each virtual machine may manage mappings of the GVA to guest “physical” addresses (GPA) in the virtual machine. If the guest OS were running directly on the system 10 hardware, with no VMM, the physical addresses generated by the guest OS would indeed be the system physical addresses (SPA) of the memory locations in the system 10. However, in the virtual machine environment, the VMM 106 may manage the mappings from GPA to SPA. Thus, when processor 12 performs memory requests, the guest OS 104 may manage mappings of GVA to GPA (which may be further mapped to SPA by the VMM 106).

[0017] As illustrated in FIG. 1, the path from the I/O devices 22 to the memory 20 is at least partially separate from the path of the processors 12 to the memory 20. Specifically, the path from the I/O devices 22 to memory 20 does not pass through the MMU 14, but instead goes through the IOMMU 26. Accordingly, the MMU 14 may not provide memory management for the memory requests sourced from the I/O devices 22. Generally, memory management may comprise address translation from one type of virtual address (i.e., an address that may be used by software) to a physical address (i.e., an address that may be used by the memory controller) and memory protection. Memory protection may control read and/or write access to the memory at some level of granularity (e.g. a page), along with various other attributes such as privilege level requirements, cacheability and cache controls (e.g. writethrough or writeback), coherency, etc. Any set of memory protections may be implemented in various embodiments. In some embodiments, the memory protections implemented by the IOMMU 26 may differ from the memory protections implemented by the MMU 14, in at least some respects. In one embodiment, the memory protections implemented by the IOMMU 26 may be defined so that the translation tables storing the translation data used by the IOMMU 26 and the MMU 14 may be shared (although shown separately in FIG. 1 for ease of discussion). Other embodiments may not share translation tables between the IOMMU 26 and the MMU 14, as desired.

[0018] Generally, the I/O devices 22 may be configured to issue memory requests, such as memory read and write requests, to access memory locations in the memory 20 and in some cases, translation requests. The memory requests may be part of a direct memory access (DMA) read or write operation, for example. The DMA operations may be initiated by software executed by the processors 12, programming the I/O devices 22 directly or indirectly to perform the DMA operations. Depending on the address space in which the software executing on the processor is running, the I/O devices 22 may be provided with addresses corresponding to that address space to access the memory 20. For example, a guest application (e.g., App 102) executing on processor 12 may provide an I/O device 22 with GVAs, while a guest OS executing

on processor 12 (e.g., OS 104) may provide GPAs to the I/O devices 22. In either case, when the I/O device 22 requests a memory access, the guest addresses may be translated by the IOMMU 26 to corresponding system physical addresses (SPA) to access the memory, and the system physical addresses may be provided to the memory controller 18 for access. That is, the IOMMU 26 may modify the memory requests sourced by the I/O devices 22 to change (i.e., translate) the received address in the request to an SPA, and the memory request may be forwarded to the memory controller 18 to access the memory 20.

[0019] In various embodiments, the IOMMU 26 may provide one-level, two-level, or no translations depending on the type of address it receives from the I/O device. More particularly, the IOMMU 26 may perform one-level nested translations or two-level guest translations. That is to say, the IOMMU 26 may provide both GPA to SPA translations (one-level), and GVA to SPA translations (two-level). Thus, as mentioned above a guest application may provide GVA addresses directly to an I/O device when requesting memory accesses, thereby making conventional VMM interception and translation unnecessary. This functionality may allow advanced computation architectures such as compute offload, user-level I/O, and accelerated I/O devices to be used more seamlessly in virtualized systems. It is noted that although one-level, two-level, or no translations are described, it is contemplated that in other embodiments, additional address spaces may be used. In such embodiments, additional levels of translation (i.e., multilevel translations) may be performed by IOMMU 26 to accommodate the additional address spaces.

[0020] The IOMMU 26 may use a number of data structures such as one or more sets of I/O translation tables 36 stored in the memory 20 to translate the addresses of memory and translation requests from the I/O devices 22. Generally, translation tables may be tables of translation data that can be used to translate addresses from one type to another. The translation tables may store the translation data in any fashion. For example, in one embodiment, the I/O translation tables 36 may include page tables similar to those defined in the x86 and AMD64TM instruction set architectures. Depending on the translation level, various subsets of the guest virtual address bits or guest physical address may be used to index levels of the tables, and each level may either be the end of translation (i.e. storing a real page number for the translation) or may point to another table (indexed by another set of address bits). The page may be the unit of translation (i.e. each address in the virtual page translates to the same physical page). Pages may have varying sizes, from 4 kilobytes up to Megabytes or even Gigabytes.

[0021] Additionally, the I/O translation tables 36 may include a device table (shown in FIG. 3) that maps I/O devices to sets of page tables (e.g., by device identifiers). The device identifier (ID) may be defined in a variety of ways, and may be dependent on the peripheral interconnect to

which the device is attached. For example, Peripheral Component Interconnect (PCI) devices may form a device ID from the bus number, device number and function number (BDF).

HyperTransport™ (HT) devices may use a bus number and unit ID to form a device ID. As described further below, the device table may include a plurality of entries indexed by the device ID, and each entry may include a pointer to a set of page tables used by the device having the corresponding device ID. In addition, in situations where an I/O device is assigned directly to a process or may run computations in the same address space as a user process, the process address space is identified and provided to the IOMMU 26 to enforce memory isolation protections. In some embodiments, the device table may further include a pointer to an interrupt remapping table (shown in FIG. 3) to remap the device's interrupts. Thus, in general, a translation from a guest virtual address or a guest physical address to a system physical address may be stored in one or more entries in one or more translation tables, and some of the entries may be shared with other translations. Traversing or “walking” the tables from entry to entry may be part of identifying the translation for the virtual address. In one embodiment, the translation tables 36 may include the interrupt remapping table mentioned above.

[0022] Specifically, the IOMMU 26 illustrated in FIG. 1 may include the table walker 28 to search the I/O translation tables 36 for a translation for a given memory request. The table walker 28 may generate memory requests, e.g. read memory requests, to read the translation data from the translation tables 36. The translation table reads are illustrated by dotted arrows 38 and 40 in FIG. 1.

[0023] To facilitate more rapid translations, the IOMMU 26 may cache some translation data. For example, the cache 30 may be a form of cache similar to a TLB, which caches the result of previous translations, mapping guest virtual and guest physical page numbers to system physical page numbers and corresponding translation data. If a translation is not found in the cache 30 for the given memory request, the table walker 28 may be invoked. In various embodiments, the table walker 28 may be implemented in hardware, or in a microcontroller or other processor and corresponding executable code (e.g. in a read-only memory (ROM) in the IOMMU 26). Additionally, other caches may be included to cache page tables, or portions thereof, and/or device tables, or portions thereof, as part of cache 30. Accordingly, the IOMMU 26 may include one or more memories to store translation data that is read from, or derived from, translation data stored in the memory 20.

[0024] The control logic 34 may be configured to access the cache 30 to detect a hit/miss of the translation for a given memory request, and may invoke the table walker 28. The control logic 34 may also be configured to modify the memory request from the I/O device with the translated address, and to forward the request upstream toward the memory controller 18.

Additionally, the control logic 34 may control various functionality in the IOMMU 26 as programmed into the control registers 32. For example, the control registers 32 may define an area of memory to be a command queue 42 for memory management software to communicate control commands to the IOMMU 26, in this embodiment. The control logic 34 may be configured to read the control commands from the command queue 42 and execute the control commands. Similarly, the control registers 32 may define another area of memory to be an event log buffer 44. The control logic 34 may detect various events and write them to the event log buffer 44. The events may include various errors detected by the control logic 34 with respect to translations and/or other functions of the IOMMU 26. The control logic 34 may also implement other features of the IOMMU 26.

[0025] The I/O devices 22 may comprise any devices that communicate between the computer system 10 and other devices, provide human interface to the computer system 10, provide storage (e.g. disk drives, compact disc (CD) or digital video disc (DVD) drives, solid state storage, etc.), and/or provide enhanced functionality to the computer system 10. For example, the I/O devices 22 may comprise one or more of: network interface cards, integrated network interface functionality, modems, video accelerators, audio cards or integrated audio hardware, hard or floppy disk drives or drive controllers, hardware interfacing to user input devices such as keyboard, mouse, tablet, etc., video controllers for video displays, printer interface hardware, bridges to one or more peripheral interfaces such as PCI, PCI express (PCIe), PCI-X, USB, firewire, SCSI (Small Computer Systems Interface), etc., sound cards, and a variety of data acquisition cards such as GPIB or field bus interface cards, etc. The term "peripheral device" may also be used to describe some I/O devices.

[0026] In some cases, one or more of the I/O devices 22 may comprise an IOTLB, such as IOTLBs 24 and/or an MMU (e.g., Private MMU 23D of FIG. 2). These IOTLBs may be referred to as "remote IOTLBs", since they are external to the IOMMU 26. In such cases, the addresses that have already been translated may be marked in some fashion so that the IOMMU 26 does not attempt to translate the memory request again. In one embodiment, the translated addresses may simply be marked as "pretranslated."

[0027] The memory controller 18 may comprise any circuitry designed to interface between the memory 20 and the rest of the system 10. The memory 20 may comprise any semiconductor memory, such as one or more RAMBUS DRAMs (RDRAMs), synchronous DRAMs (SDRAMs), DDR SDRAM, static RAM, etc. The memory 20 may be distributed in a system, and thus there may be multiple memory controllers 18.

[0028] The MMU 14 may comprise a memory management unit for memory requests sourced by a processor 12. The MMU may include TLBs 16, as well as table walk functionality. When a

translation is performed by the MMU 14, the MMU 14 may generate translation memory requests (e.g. shown as dotted arrows 46 and 48 in FIG. 1) to the CPU translation tables 50. The CPU translation tables 50 may store translation data as defined in the instruction set architecture implemented by the processors 12.

[0029] The processors 12 may comprise any processor hardware, implementing any desired instruction set architecture. In one embodiment, the processors 12 implement the x86 architecture, and more particularly the AMD64TM architecture. Various embodiments may be superpipelined and/or superscalar. Embodiments including more than one processor 12 may be implemented discretely, or as chip multiprocessors (CMP) and/or chip multithreaded (CMT).

[0030] The system 10 illustrates high level functionality of the system, and the actual physical implementation may take many forms. For example, the MMU 14 is commonly integrated into each processor 12. While a memory 20 is shown, the memory system may be a distributed memory system, in some embodiments, in which the memory address space is mapped to multiple, physically separate memories coupled to physically separate memory controllers. The IOMMU 26 may be placed anywhere along the path between I/O-sourced memory requests and the memory 20, and there may be more than one IOMMU. Still further, IOMMUs may be located at different points in different parts of the system.

[0031] FIG. 2 is one example of an embodiment of a system that uses HyperTransportTM (HT) coherent links 61 between processor nodes and HT I/O links (e.g., 70A and 70B) between processor nodes and I/O devices. More particularly, in one embodiment, the HT I/O links (e.g., 71A) may be used between I/O hubs and I/O devices that bridge to other peripheral interconnects. In addition, any other coherent interconnect may be used between processor nodes and/or any other I/O interconnect may be used between processor nodes and the I/O devices. For example, I/O links 71B, 71C and 71D may be PCIe links in various embodiments. Furthermore, another example may include processors coupled to a Northbridge, which is further coupled to memory and one or more I/O interconnects, in a traditional PC design.

[0032] Referring to the embodiment shown in FIG. 2, the system 10a comprises processing nodes 60A-60B, which respectively comprise processors 12A-12B further comprising MMUs 14A-14B. The processor nodes 60A-60B also include memory controllers 18A-18B. Each of processors 12A-12B may be an instance of a processor 12 as mentioned above. Similarly, each of MMUs 14A-14B and memory controllers 18A-18B may be instances of the MMU 14 and memory controller 18 shown in FIG. 1. In the illustrated embodiment, the MMU functionality is incorporated into the processor.

[0033] The system 10a includes a distributed memory system, comprising memories 20A-20B. The system physical address space may be distributed over the memories 20A-20B.

Accordingly, a given memory request specifying a given address is routed to the memory controller 18A or 18B coupled to the memory 20A or 20B to which that given address is assigned.

[0034] Memory requests from the I/O devices (e.g. I/O devices 22A-22D, coupled to I/O Hubs 62A-62B as illustrated in FIG. 2) may not all take the same path to arrive at the memory controller 18A-18B that will service the request. For example, the I/O devices 22A-22B may transmit memory requests to the I/O hub 62A, which transmits the requests to the processing node 60A. If the address of a given memory request is assigned to the memory 20B, the processing node 60A may transmit the given memory request to the processing node 60B, so that the memory controller 18B may receive and process the request. The I/O devices 22C-22D may transmit memory requests to the I/O Hub 62B, which may transmit the requests to the processing node 60B. If the address of a given memory request is assigned to the memory 20A, the processing node 60B may transmit the given memory request to the processing node 60A.

[0035] The IOMMU may be placed anywhere along the path between I/O-sourced memory requests and the memory 20. In the illustrated embodiment, IOMMUs 26A-26B are included in the I/O hubs 62A-62B. Thus, any memory requests sourced by an I/O device coupled to the corresponding hub may be translated by the IOMMU in the I/O hub. Other embodiments may locate the IOMMU in different places, from IOTLBs in the I/O devices to IOMMUs within the processing nodes 60A-60B, or even IOMMUs at the memory controllers 18A-18B. Still further, IOMMUs may be located at different points in different parts of the system. For example, in some embodiments the I/O devices may be devices that may include local memory 24D as shown. Accordingly, I/O device 22D also includes a local/private MMU 23D.

[0036] Turning to FIG. 3, a diagram illustrating one embodiment of the I/O translation tables 36 is shown. Specifically, the I/O translation tables 36 may include a device table 36A, an interrupt remapping table 36B, and two sets of I/O page tables (e.g., 36C and 36D). Also shown in FIG. 3 is one of the control registers 32 (control register 32A). The control register 32A may store a base address of the device table 36A.

[0037] The device table 36A includes a plurality of entries, indexed by a device ID assigned to the device. Thus, a given device corresponds to one of the entries in the device table 36A (unless the device has multiple device IDs, or unless the device has its traffic aggregated with others at a bridge device, and the traffic is transmitted under the bridge's device ID). The device table entry may include a variety of data. Portions of an exemplary device table entry 325 are shown in FIG. 4 and described in more detail below.

[0038] A given device table entry 325 may include a pointer (represented by arrows 301 and 302) to the I/O page tables 36C and 36D. The pointer to may point to a page table that is the

starting point for translation searching in the I/O page tables 36C or 36D. As shown, the starting page table may include pointers to other page tables, in a hierarchical fashion, as mentioned above. As described further below, some tables may be indexed by a process identifier, while other tables may be indexed using various bits of either the guest virtual address or the guest physical address to be translated, according to the implemented translation process.

[0039] The device table entry 325 may also include a pointer to the interrupt remapping table 36B (represented by arrow 304). The interrupt remapping data may be used when an interrupt request is transmitted by a device, and may be indexed by an interrupt ID. The interrupt ID may comprise data that identifies the requested interrupt, and may vary based on the mechanism used to transmit the interrupt request.

[0040] It is noted that, while one device table 36A is shown, multiple device tables may be maintained if desired. The device table base address in the control register 32A may be changed to point to other device tables. Furthermore, device tables may be hierarchical, if desired, similar to the page tables described above. Similarly, while one interrupt remapping table 36B is shown, there may be multiple interrupt mapping tables, e.g. up to one per entry in the device table 36A. There may also be multiple sets of page tables, e.g. up to one per entry in the device table 36A. It is noted that although independent of one another, other embodiments may implement I/O translation functionality without interrupt remapping functionality.

[0041] In the illustrated embodiment, the I/O translation tables 36C and 36D may support various combinations of one-level and two-level guest translations. As illustrated by the dotted lines, when a translation is a guest (i.e., two-level) translation (represented by dashed lines), the path beginning with arrow 301 is followed. However, when guest translations are not valid, then just the solid line path is followed beginning with arrow 302.

[0042] More particularly, when a memory access request or a translation request including a GVA is received from an I/O device, the device table entry corresponding to the device ID of the request is accessed. As described in greater detail below in conjunction with the description of FIG. 5, an SPA pointer (e.g., 301) to the base address of a GCR3 table 307 is used, and a process ID (e.g., PASID) may be used to index into the GCR3 table 307. The entry within the GCR3 table 307 may include a GPA pointer to the base of the GCR3 page table 311. However, this GPA and all subsequent GPA pointers in the tables may be translated to an SPA using translation logic and a set of page tables such as tables 333, and 335, for example, to access the next page table in memory. This translation mechanism is represented in FIG. 3 by the GPA/SPA Trans units 309A, 309B, 309C, and 309D. Some number of bits of the GVA of the request may be used to index into the GCR3 page table 311. In one embodiment, the corresponding entry in 311 may include a GPA pointer to the base of a page map level four (PML4) page table 313. The

subsequent addresses in each of the nested page tables may also be GPAs that point to the base address of the next table, or in the case of the last table (e.g., 315) the entry may include a GPA pointer to the physical memory page. Each of these tables (e.g., 313, 315) may also be indexed by the GVA of the I/O request. The final translation (SPA) may be obtained by concatenating a portion of the GVA (e.g., the last several bits) and the SPA of the physical page in memory 317.

[0043] Thus, both guest and nested translation tables may be used in the two-level guest translation. However, if a GPA is provided by the I/O device, then a one-level translation may be performed by using the pointer 302 which provides the base address of the host translation page tables 36D. More particularly, the guest information in the device table entry is ignored, and other page table routing entry information (shown in FIG. 4) may be used to access the corresponding host translation page tables. In various levels of translation, different subsets of bits of the GPA may be used to index into each successive page table. Specifically, the pointer 302 may be an SPA pointer to the base of the PML4 page table 333. In contrast to the guest tables 36C, each pointer in the host page tables 36D may be an SPA which points to a base of the next table. The entry indexed by the GPA within PM4L page table 333 includes an SPA pointer to the base of PML3/2 page table 335. Depending on the actual page table level (i.e., 3 or 2), the pointer in the entry of page table 335 may be an SPA pointer to the next table, or an SPA of the physical memory page 337. Accordingly, the final translation (SPA) may be obtained by concatenating a portion of the GPA (e.g., the last several bits) and the SPA of the physical page in memory 337.

[0044] To preserve memory isolation, in one embodiment, when accessing each subsequent page table, permissions may be checked, and if there is a permission discrepancy or the page is not present, the IOMMU 26 may abort the translation and or lookup and internally issue a page fault and log the error. In other embodiments, at each stage, the IOMMU may simply accumulate any errors and at the end of the translation, discard the translation, issue a page fault and log the error, or some combination thereof.

[0045] Turning to FIG. 4, a diagram of one embodiment of an exemplary device table entry is shown. It is noted that the fields shown in the embodiment of FIG. 4 may be related to interrupt remapping and/or address translation, and other fields may be provided for other purposes in various embodiments. In the illustrated embodiment, the device table entry 325 is a 256-bit entry. However, other embodiments may implement supersets of the fields and other fields, subsets of the fields, or subsets in combination with other fields. Consequently the device table entry 325 may be any number of bits, as desired. A field may comprise one or more bits, the encoding of which are assigned particular meanings when interpreted by the IOMMU 26. If a field is a single bit, for this embodiment, it may be referred to as a bit. However, multi-bit fields

for similar purposes may be used in other embodiments. Multi-bit fields are shown with bit ranges in FIG. 4, for this embodiment.

[0046] Various specific interrupt control bits may be provided in the Reserved and Other field to control whether certain specific interrupts are blocked or passed unmodified. The interrupt table pointer field (IntTablePtr) may store the base address of an interrupt remapping table mentioned above. The interrupt table length (IntTableLen) specifies the extent of the interrupt table. The interrupt table length field may be encoded for several possible lengths (e.g. 1-2048 entries in powers of two, for one embodiment). The IG bit may be used to indicate whether or not event log entries in the event log buffer 44 are to be created if an I/O page fault is detected for an interrupt message. The interrupt valid (IV) bit may indicate whether or not the interrupt-related fields are valid. If the fields are not valid, the IOMMU 26 may pass all interrupts unmodified.

[0047] The guest CR3 (GCR3) table root pointer (GCR3TableRtPtr) fields may include the SPA of the guest CR3 table for the I/O device. The peripheral page service request (PPR) bit indicates whether a peripheral may issue peripheral page service request that the IOMMU converts to peripheral page service request log entries or whether the request is handled as an error. It is noted that although the GCR3TableRtPtr is broken up into multiple fields in the illustrated embodiment, in other embodiments, the GCR3TableRtPtr may be one contiguous field, or broken up differently as desired.

[0048] The SysMgt field may be encoded to provide further control of communications in the system management range. Specifically, in one embodiment, the SysMgt field may be encoded to: block requests in the range; forward requests in the range unmodified (posted writes only); forward requests that map to INTx messages unmodified (posted writes only); or translate requests using the I/O page tables. The IoCtl field may be encoded to provide further control of communications in the I/O space range. Specifically, in one embodiment, the IoCtl field may be encoded to: block requests in the range; forward the requests unmodified; or translate the requests using the I/O page tables. However, in one embodiment, some of these fields (e.g., SysMgt, IoCtl) may be relevant only to GPA-SPA translations but not applied to GVA-SPA translations

[0049] The Domain ID is used to tag cache entries and any other cache entries in the IOMMU 26 so that different devices differentiate their translation data. If devices share translation tables, they may have the same Domain ID to share cache entries. The Domain ID is completely under the control of software, and thus may permit flexibility for controlling software (e.g. a virtual machine monitor, or an operating system in non-virtual machine implementations) to group I/O devices into a domain to share translation data, or to separate the devices. For

example, devices assigned to a given virtual machine may have the same Domain ID, and different Domain IDs may be used for different virtual machines. Any combination of separated devices and grouped devices may thus be created.

[0050] In one embodiment, the guest CR3 level (GLX) bit specifies the type of GCR3 lookup (i.e., one-level or two-level GCR3 tables) will be performed by the IOMMU. The guest translation valid (GV) bit determines whether two-level (i.e., nested and guest level) translation or one-level (e.g., host) translation will be performed. In one embodiment, when the GV bit is clear, the GLX bit and GCR3TableRtPtr fields are ignored.

[0051] The page table pointer (PageTablePtr) is the pointer to the host translation page tables 36D. This pointer may be used when I/O devices provide a GPA. The mode field (Mode) may be coded to indicate the depth of the device's I/O page tables and whether translation is disabled entirely. For example, for devices providing an SPA and which mark the requests as pre-translated, the corresponding device entry may have this field encoded to disable translations. Other encodings may indicate the level in the page table hierarchy at which translation is to begin for requests mapped to this entry. The TV bit indicates whether or not the page translation data are valid, and the V bit indicates if the entry 80 is valid.

[0052] Referring to FIG. 5, a diagram depicting more detailed aspects of the guest virtual address translation mechanism is shown. As described above, the device table 36A may include a number of entries (DTE), each including a GCR3TableRtPtr field. If the guest translation is valid for that entry, the field will include an SPA pointer to the guest CR3 table 501 in memory. As shown, the GCR3 table 501 may be indexed by the PASID value associated with the device's request. As mentioned above, the device PASID may be obtained in a variety of ways dependent upon the type of I/O interconnect that is used. As shown, GCR3 table 501 has two entries, each including a GCR3 base pointer, which as described above, may be a GPA pointer to guest page tables 503 and 505. In the illustrated embodiment, guest page tables 503 and 505 may be indexed by the GVA of the I/O request. Further, each valid entry in the guest page tables 503 and 505 may include a GPA pointer to the base address of a respective page map level four table.

[0053] In one embodiment, the GCR3 table 501 must be contiguous in memory. Thus in systems having a large number of PASIDs, the memory management may get cumbersome.

Accordingly, in an alternative embodiment, the GCR3 tables may be implemented in a hierarchical manner. For example, one or more 2nd level GCR3 tables (not shown) may be used. In such an embodiment, the 1st level GCR3 table may be a GCR3 base table and be indexed using a first subset of the PASID bits. A given entry in the 1st level GCR3 table may include a GPA pointer to a respective 2nd level GCR3 table. The 2nd level GCR3 tables may be indexed using a second subset of the PASID. Each entry of the 2nd level GCR3 tables may include a GPA pointer

to a guest page table such as table 503 or 505. Thus the choice of whether to implement one or two-level GCR3 tables may be dependent upon the number of PASIDs in use. It is noted that in various embodiments, the subsets of PASID bits may be non-overlapping or they may be overlapping, as desired.

5 **[0054]** Turning to FIG. 6, a flow diagram describing operational aspects of the embodiments of the system shown in FIG 1 through FIG. 5 is shown. Referring collectively to FIG. 1-FIG. 6 and beginning in block 601 of FIG. 6, an I/O request is received by IOMMU 26. For example, the request may be an I/O memory access request such as a DMA access, for example, on behalf of one of the VM applications. Alternatively, the request may be a translation request by an I/O
10 device with a remote IOTLB, for example. If the request is not a memory request, operation may proceed as shown in block 604. However, if the request is a memory request, IOMMU 26 may determine if the request is marked as pre-translated (block 602), as in the case of an I/O device having an onboard MMU, for example. In that case the address may be an SPA, and not need any translation. If the request is marked pre-translated, IOMMU 26 may provide the request to
15 the memory controller 18 (block 603).

[0055] Referring back to block 602, if the request is not marked as pre-translated or if the request is a translation request, IOMMU 26 may do a lookup within cache 30 for the translation (block 604). If the translation is present, the IOMMU 26 may provide the translation back to the requestor, or provide the translation along with the request to the memory controller 18 (block
20 603).

[0056] However, if the translation is not present in the cache 30 (block 604), table walker 28 in conjunction with control logic 34 may access the I/O translation tables 36. More particularly, table walker 34 may access the device table base register 32A to obtain the device table base address in memory 20. Table walker 28 may use the device ID to index into the device table
25 36A (block 605). If the request includes a PASID value (block 606), the table walker 28 may perform a two-level guest translation using the GCR3 control bits and the GCR3TableRtPtr address in the device table entry to find the GCR3 tables in memory (block 607). As described above, the GCR3TableRtPtr address may be an SPA, in one embodiment. The table walker 28 may use the GPA pointers to locate the corresponding GCR3 and guest tables and the PASID to
30 index the GCR3 tables. As shown in block 614, and described above, each GPA pointer may be translated by the IOMMU 26 using, for example, the host translation page tables 36D. Some subset of the GVA bits of the request may be used to index the guest page table(s). The last guest page table may include a GPA pointer to the base address of the host page map level four tables. As each successive guest table is accessed the permissions are checked, and any failed
35 permissions may generate a page fault or other errors, in one embodiment (block 608).

[0057] The table walker 28 may then successively find the host translation tables 36D using the GPA pointers within each entry of each preceding table. At various levels of the page tables, different subsets of the GVA bits of the request may be used to index into each host translation page table. As described above, once the last page table is accessed, the corresponding entry may contain the GPA of the physical page in memory. The control logic 34 may translate that GPA into an SPA (block 609). The control logic 34 may then concatenate the SPA of the physical memory page with the last few bits of the GVA to obtain the final translation address (block 610). The control logic 34 may provide the translation address to the memory controller 18 along with the request, if the request is a memory request. Alternatively, the control logic 34 may provide the translation to the requesting I/O device if the request was a translation request. The control logic 34 may also store the translation within the cache 30 for use in future translation lookups (block 611).

[0058] Referring back to block 606, if the request does not include a PASID, then the address in the request is a GPA. In this case, the control logic 34 may perform a one-level translation using the PageTableRtPtr address in the device table entry to find the host translation page tables 36D in memory (block 612). In a way similar to that described above, table walker 28 may use the SPA pointer to the base address of the host page map level four tables for the I/O device. In contrast to the guest translations, in the host translation page tables, each pointer to a next page table/ memory physical page is an SPA. Thus, the table walker 28 may successively find the host translation tables using the SPA pointers within each entry of each preceding table. At each level of the host page tables, different subsets of the GPA bits of the request may be used to index into each host translation page table. As described above, once the last page table is accessed, the corresponding entry may contain the SPA of the physical page in memory (block 613). The control logic 34 may concatenate the SPA of the physical memory page with the last few bits of the GPA of the request to obtain the final translation address (block 610). Operation may proceed as described above in conjunction with block 611.

[0059] It is noted that the above embodiments may be implemented in hardware, software, or a combination of both. In addition, various ones of the above hardware components may first be designed using software written in a hardware definition language. Accordingly, any of the instructions used to implement the above embodiments stored on various types of storage. Thus various embodiments may further include storing instructions and/or data implemented in accordance with the foregoing description upon a computer readable storage medium. Generally speaking, a computer readable storage medium may include storage media or memory media such as magnetic or optical media, e.g., disk, CD-ROM, DVD, volatile or non-volatile media such as RAM (e.g. SDRAM, DDR SDRAM, RDRAM, SRAM, etc.), ROM, and the like.

[0060] Although the embodiments above have been described in considerable detail, numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.

WHAT IS CLAIMED IS:

1. An input/output (I/O) memory management unit (IOMMU) configured to control requests by an I/O device to a memory, the IOMMU comprising:

control logic configured to perform a two-level guest translation to translate an address associated with an I/O device-generated request using translation data stored in the system memory;

wherein the translation data includes:

one or more device table entries in a device table, wherein the device table entry for a given request is selected based upon a device identifier and a first set of page tables including a set of guest page tables and a set of nested page tables, wherein a selected device table entry for a given request includes a pointer to the set of guest translation tables, and wherein a guest page table includes a pointer to the set of nested page tables.

2. The IOMMU as recited in claim 1, wherein the pointer to the set of guest translation tables comprises a system physical address (SPA) corresponding to an address space mapped by a virtual memory monitor (VMM) executing in on a processor of the computer system, wherein a given SPA is used to access a physical memory page in the memory.

3. The IOMMU as recited in claim 2, wherein the pointer to the set of nested page tables comprises a guest physical address (GPA) corresponding to an address space mapped by a guest operating system executing in a virtual machine (VM) on a processor of the computer system, wherein the control logic is further configured to translate the GPA to an SPA for accessing a physical memory page in the memory.

4. The IOMMU as recited in claim 2, wherein the address associated with the I/O request comprises a guest virtual address (GVA) corresponding to an address space mapped by a guest application executing in a virtual machine (VM) on a processor of the computer system, wherein the control logic is configured to translate the GVA to an SPA using the two-level guest translation.

5. The IOMMU as recited in claim 1, wherein at least some of the guest page tables include at least one entry that includes a pointer to a respective next guest page table.

6. The IOMMU as recited in claim 1, wherein a last guest page table of the set of guest page tables includes at least one entry that includes a pointer to a respective first page table of the set of nested page tables.

7. The IOMMU as recited in claim 3, wherein a last page table of the set of nested page tables includes at least one entry that includes a pointer having an address corresponding to a respective physical page in the memory.

8. The IOMMU as recited in claim 7, wherein the pointer having an address corresponding to a respective physical page in the memory comprises a GPA.

9. The IOMMU as recited in claim 8, wherein the control logic is further configured to translate the GPA pointer corresponding to a respective physical page in the memory to an SPA of the respective physical page.

10. The IOMMU as recited in claim 9, wherein the control logic is further configured to concatenate a subset of bits of the address associated with the I/O request with the SPA of the respective physical page to provide a final translation address.

11. The IOMMU as recited in claim 1, wherein the device table entry for the given request further comprises a second pointer to a second set of I/O page tables including a set of host translation page tables.

12. The IOMMU as recited in claim 11, wherein at least some of the host translation tables include entries having a pointer to a next successive host translation table, wherein each pointer comprises a system physical address (SPA) corresponding to an address space mapped by a virtual memory monitor (VMM) executing in on a processor of the computer system, wherein a given SPA is used to access a physical memory page in the memory.

13. A system comprising:
a processor;
a memory coupled to the processor and configured to store translation data, wherein the translation data includes:
a device table including one or more entries;

a set of page tables including a set of guest page tables and a set of nested page tables;

a least one input/output (I/O) device configured to generate a request for accessing the system memory;

5 an I/O memory management unit (IOMMU) coupled to the I/O device and to the memory, wherein the IOMMU is configured to perform a two-level guest translation to translate an address associated with the I/O device-generated request using the translation data;

wherein the IOMMU is configured to:

10 select a device table entry for a given request based upon a device identifier;
use a pointer from a selected device table entry to access the set of guest translation data structures, and
use a pointer from a guest page table to access the set of nested page tables.

15 14. The system as recited in claim 13, wherein the IOMMU is configured to use the device table and the set of guest page tables in a first guest translation level, and to use the set of nested page tables in a second guest translation level to obtain a final translation address.

20 15. The system as recited in claim 13, wherein the memory is mapped into a plurality of address spaces including a system physical address (SPA) corresponding to an address space mapped by a virtual memory monitor (VMM) executing in on the processor, a guest physical address (GPA) corresponding to an address space mapped by a guest operating system executing in a virtual machine (VM) on the processor, and a guest virtual address (GVA) corresponding to an address space mapped by a guest application executing in a virtual machine (VM) on the
25 processor.

16. The system as recited in claim 15, wherein the pointer to the set of guest translation tables comprises an SPA, wherein a given SPA is used to access a physical memory page in the
30 memory.

17. The system as recited in claim 15, wherein the pointer to the set of nested page tables comprises a GPA, wherein the IOMMU is further configured to translate the GPA to an SPA for accessing a physical memory page in the memory.

18. The system as recited in claim 14, wherein the address associated with the I/O request comprises a GVA, wherein the IOMMU is configured to translate the GVA to an SPA using the two-level guest translation.

19. The system as recited in claim 15, wherein the IOMMU is further configured to: index into at least a portion of the set of guest page tables using a process identifier associated with the I/O request, and index into another portion of the set of guest page tables using the GVA associated with the I/O request.

20. The system as recited in claim 19, wherein in response to the I/O device-generated request including the process identifier, the IOMMU is configured to perform the two-level guest translation to translate the GVA associated with the I/O device-generated request.

21. The system as recited in claim 19, wherein in response to the I/O device-generated request not including the process identifier, the IOMMU is configured to perform a one-level translation to translate the address associated with the I/O device-generated request, wherein the address comprises a GPA.

22. A method for controlling input/output I/O requests to a memory of a computer system, the method comprising:

storing translation data in the memory of a computer system, wherein the translation data includes:

one or more device table entries in a device table; and

a first set of page tables including a set of guest page tables and a set of nested page tables;

control logic performing a first translation level of a two-level guest translation on an address received in an I/O request by:

selecting a device table entry for a given request based upon a device identifier; and

using a pointer from a selected device table entry for accessing the set of guest translation tables; and

the control logic performing a second translation level of the two-level guest translation by using a pointer from a guest page table for accessing the set of nested page data structures.

23. The method as recited in claim 22, the control logic performing the two-level guest translation in response to receiving a process identifier in the I/O request.

24. The method as recited in claim 22, wherein the address received in the I/O request comprises a guest virtual address (GVA) corresponding to an address space mapped by a guest application executing in a virtual machine (VM) on a processor of the computer system.

25. The method as recited in claim 23, the control logic performing a one-level translation in response to receiving an I/O request that does not include a process identifier.

26. The method as recited in claim 25, wherein the address received in the I/O request that does not include a process identifier comprises a guest physical address (GPA) corresponding to an address space mapped by a guest operating system executing in a virtual machine (VM) on a processor of the computer system.

27. A computer readable storage medium comprising a data structure which is operated upon by a program executable on a computer system, the program operating on the data structure to perform a portion of a process to fabricate an integrated circuit including circuitry described by the data structure, the circuitry described in the data structure including:

control logic configured to perform a two-level guest translation to translate an address associated with an I/O device-generated request using translation data stored in the memory;

wherein the translation data includes:

one or more device table entries in a device table, wherein the device table entry

for a given request is selected based upon a device identifier; and

a first set of page tables including a set of guest page tables and a set of nested page tables, wherein a selected device table entry for a given request includes a pointer to the set of guest translation tables, and wherein a guest page table includes a pointer to the set of nested page tables.

28. The IOMMU as recited in claim 1, further comprising a cache memory coupled to the control logic and configured to store completed translations.

29. The method as recited in claim 24, further comprising the control logic storing completed translations within a cache memory.

30. An input/output (I/O) memory management unit (IOMMU) comprising:
control logic configured to perform a multi-level hierarchical translation to translate a
guest address associated with an I/O device-generated request using translation
data stored in a memory, wherein the guest address corresponds to a guest virtual
address space mapped by a guest application executing in a guest virtual machine
on a processor;
wherein the translation data includes:
one or more device table entries in a device data structure; and
a first set of page data structures including a set of guest page data structures and a
set of nested page data structures;
wherein the control logic is further configured to:
select a device data structure entry for a given request based upon a device
identifier that corresponds to the I/O device that generates the request;
use a pointer from a selected device data structure entry to access the set of guest
translation data structures, and
use a pointer from a guest page data structure to access the set of nested page data
structures.

1/6

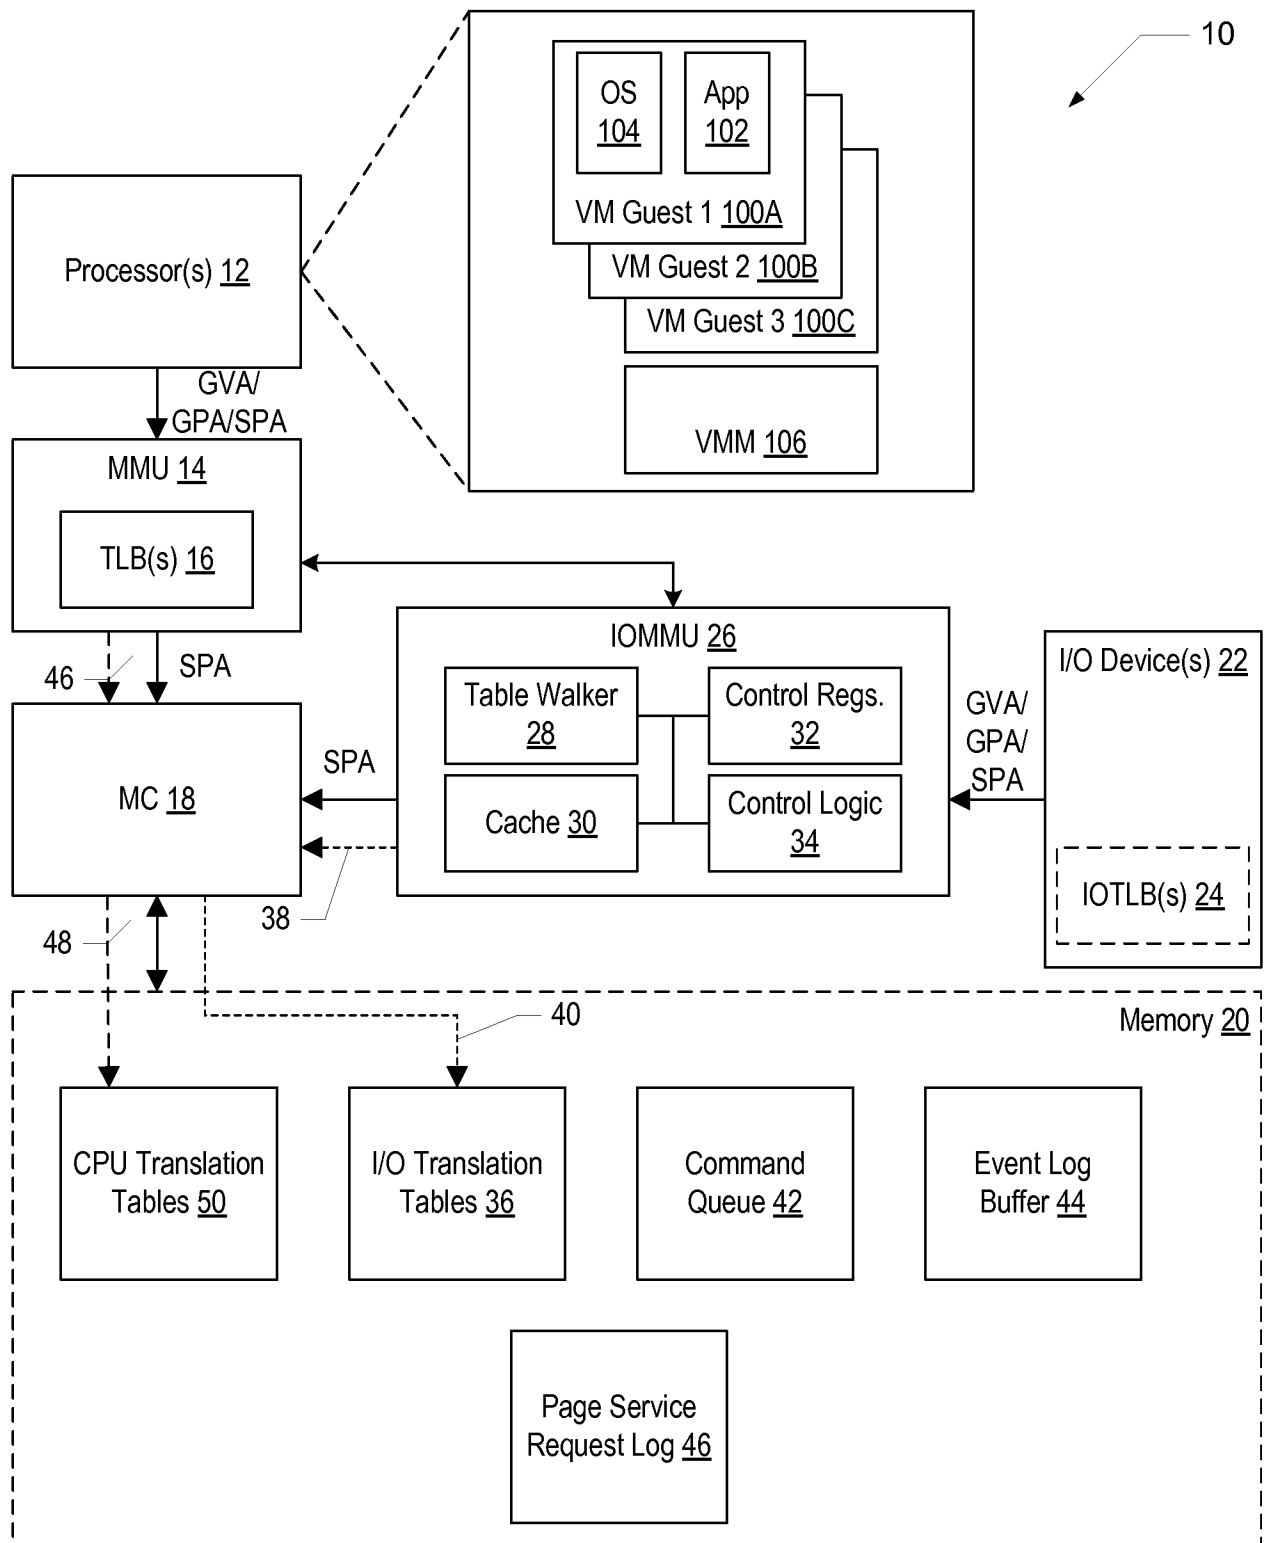


FIG. 1

2/6





FIG. 2

3/6



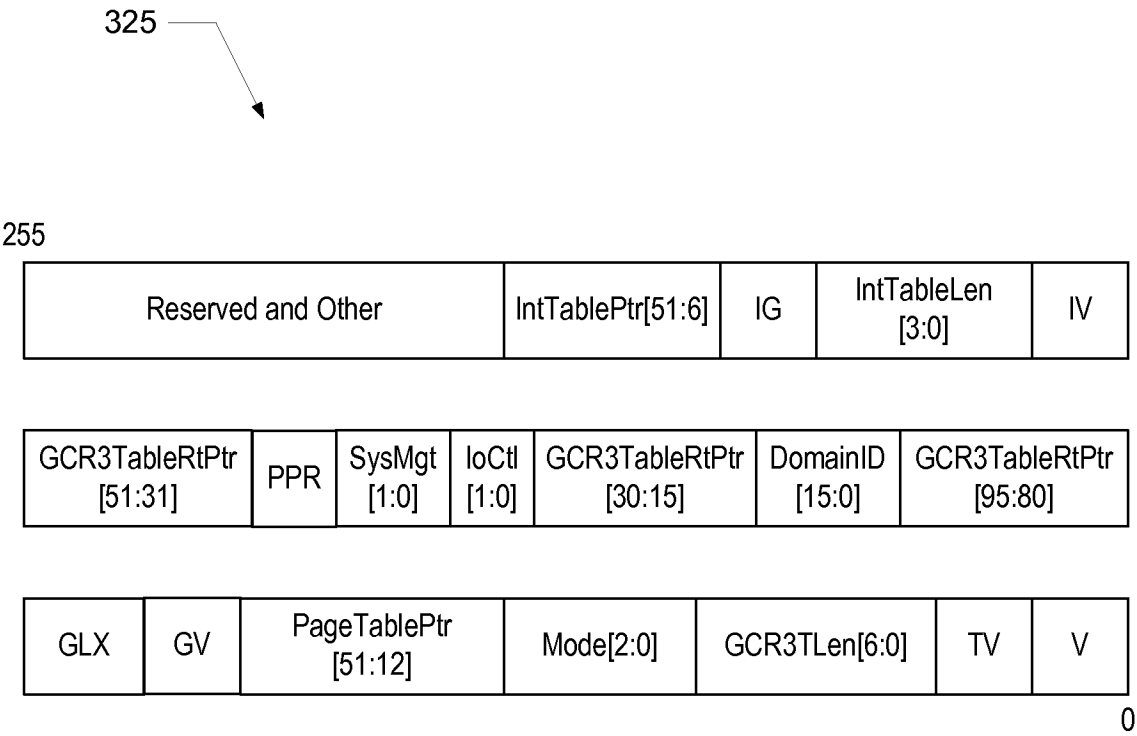


FIG. 4

5/6

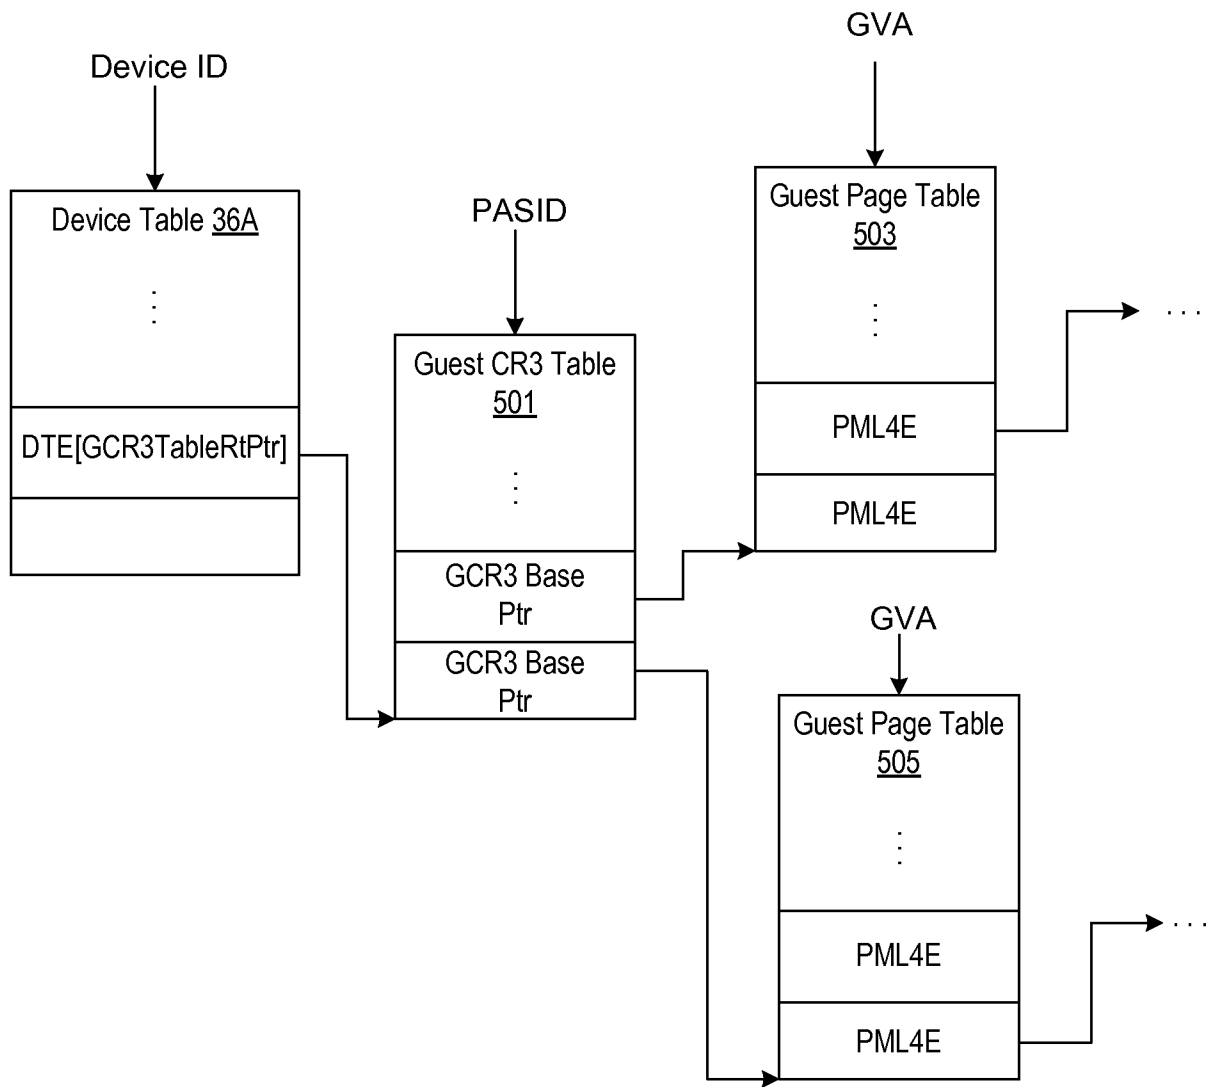


FIG. 5

6/6

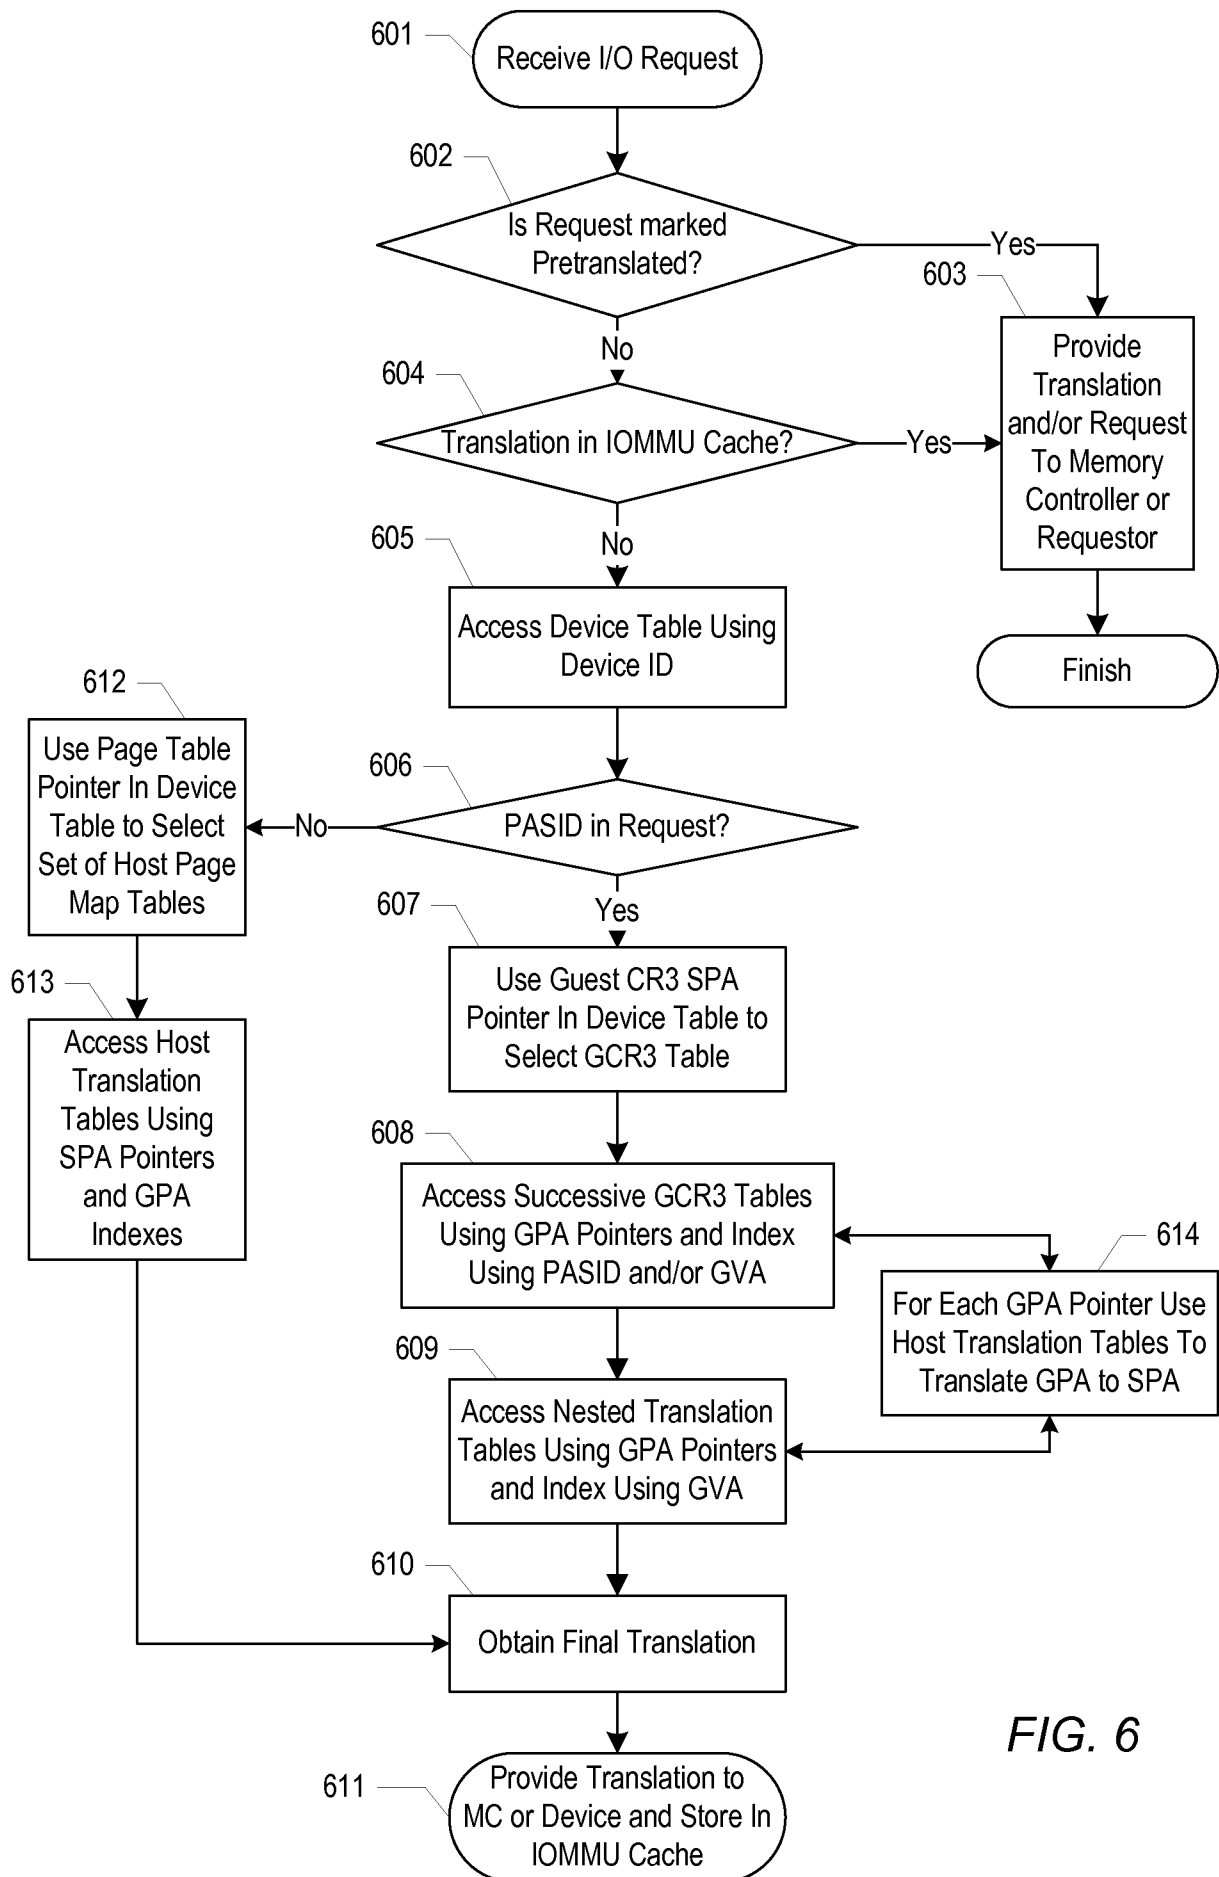


FIG. 6

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2010/043169

A. CLASSIFICATION OF SUBJECT MATTER

INV. G06F12/10

ADD.

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F

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

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

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	<p>US 2007/168643 A1 (HUMMEL [US] ET AL) 19 July 2007 (2007-07-19)</p> <p>paragraph [0006] - paragraph [0009] paragraph [0023] - paragraph [0029] paragraph [0041] - paragraph [0046] paragraph [0057] - paragraph [0076] figures 1-10</p> <p style="text-align: center;">----- -/--</p>	<p>1, 5, 6, 11, 13-15, 17-30</p>



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

1 November 2010

Date of mailing of the international search report

05/11/2010

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

McDonagh, Fintan

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2010/043169

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	<p>US 2009/187697 A1 (SEREBRIN [US]) 23 July 2009 (2009-07-23)</p> <p>paragraph [0038] - paragraph [0043] paragraph [0049] - paragraph [0050] paragraph [0071] - paragraph [0072] paragraph [0082] - paragraph [0083] paragraph [0107] - paragraph [0110] figure 15</p> <p>-----</p>	<p>1, 5, 6, 11, 13-15, 17-30</p>
A	<p>US 2006/069899 A1 (SCHOINAS [US] ET AL) 30 March 2006 (2006-03-30) paragraph [0039] - paragraph [0048] paragraph [0066] - paragraph [0070] paragraph [0119] - paragraph [0131] figures 4, 5</p> <p>-----</p>	<p>1-30</p>
A	<p>US 2006/206687 A1 (VEGA [US]) 14 September 2006 (2006-09-14) paragraph [0013] - paragraph [0036] figures 2, 3, 5</p> <p>-----</p>	<p>1-30</p>

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2010/043169

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2007168643 A1	19-07-2007	US 2010095085 A1	15-04-2010
US 2009187697 A1	23-07-2009	US 2009187726 A1	23-07-2009
		US 2009187902 A1	23-07-2009
		US 2009187698 A1	23-07-2009
		US 2009187729 A1	23-07-2009
		US 2009187904 A1	23-07-2009
		WO 2009094163 A2	30-07-2009
US 2006069899 A1	30-03-2006	CN 101027653 A	29-08-2007
		DE 112005002298 T5	27-09-2007
		GB 2432443 A	23-05-2007
		JP 2008515093 T	08-05-2008
		US 2010011187 A1	14-01-2010
		WO 2006039057 A2	13-04-2006
US 2006206687 A1	14-09-2006	NONE	