

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
3 August 2006 (03.08.2006)

PCT

(10) International Publication Number
WO 2006/081582 A2

(51) International Patent Classification:
G06F 9/455 (2006.01)

(21) International Application Number:
PCT/US2006/003587

(22) International Filing Date: 27 January 2006 (27.01.2006)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/045,524 28 January 2005 (28.01.2005) US

(71) Applicant (for all designated States except US): **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95052 (US).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **ANDERSON, Andrew** [US/US]; 677 SE 68th Ave, Hillsboro, Oregon 97123 (US).

(74) Agents: **VINCENT, Lester J.** et al.; **BLAKELY SOKOLOFF TAYLOR & ZAFMAN**, 12400 Wilshire Boulevard, 7th Floor, Los Angeles, California 90025 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: A METHOD AND APPARATUS FOR SUPPORTING ADDRESS TRANSLATION IN A VIRTUAL MACHINE ENVIRONMENT

(57) Abstract: In one embodiment, a method includes receiving control transitioned from a virtual machine (VM) due to a privileged event pertaining to a translation-lookaside buffer (TLB), and determining which entries in a guest translation data structure were modified by the VM. The determination is made based on metadata extracted from a shadow translation data structure maintained by a virtual machine monitor (VMM) and attributes associated with entries in the shadow translation data structure. The method further includes synchronizing entries in the shadow translation data structure that correspond to the modified entries in the guest translation data structure with the modified entries in the guest translation data structure.



WO 2006/081582 A2

A METHOD AND APPARATUS FOR SUPPORTING ADDRESS TRANSLATION IN A VIRTUAL MACHINE ENVIRONMENT

Field

5 **[0001]** Embodiments of the invention relate generally to virtual machines, and more specifically to supporting address translation in a virtual machine environment.

Background

10 **[0002]** A conventional virtual-machine monitor (VMM) typically runs on a computer and presents to other software the abstraction of one or more virtual machines. Each virtual machine may function as a self-contained platform, running its own “guest operating system” (i.e., an operating system (OS) hosted by the VMM) and other software, collectively referred to as guest software. The guest software expects to operate as if it were running on a dedicated computer rather than a virtual machine. That is, the guest
15 software expects to control various events and have access to hardware resources such as physical memory and memory-mapped input/output (I/O) devices. For example, the guest software expects to maintain control over address-translation operations and have the ability to allocate physical memory, provide protection from and between guest applications, use a variety of paging techniques, etc. However, in a virtual-machine
20 environment, the VMM should be able to have ultimate control over the computer’s resources to provide protection from and between virtual machines.

Brief Description of the Drawings

25 **[0003]** The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

[0004] **Figure 1** illustrates one embodiment of a virtual-machine environment, in which the present invention may operate;

30 **[0005]** **Figure 2** illustrates operation of a virtual TLB, according to one embodiment of the present invention;

[0006] **Figures 3A and 3B** illustrate a process of creating and maintaining metadata for a shadow PT hierarchy, according to two alternative embodiments of the present invention;

[0007] **Figure 4** is a flow diagram of one embodiment of a process for synchronizing guest translation data structure and shadow translation data structure;

[0008] **Figure 5** is a flow diagram of one embodiment of a process for maintaining metadata for a shadow translation data structure;

5 [0009] **Figure 6** is a flow diagram of one embodiment of a process for facilitating a change of an address space;

[0010] **Figure 7** is a flow diagram of one embodiment of a process for synchronizing entries of two translation data structures for a specified address;

10 [0011] **Figure 8** is a flow diagram of one embodiment of a process for removing a shadow PT hierarchy from a working set of shadow PT hierarchies maintained by the VMM;

[0012] **Figure 9** is a flow diagram of one embodiment of a process for adding an entry to a PD of a shadow PT hierarchy;

15 [0013] **Figure 10** is a flow diagram of one embodiment of a process for removing an entry from a PD of a shadow PT hierarchy;

[0014] **Figure 11** is a flow diagram of one embodiment of a process for adding an entry to a PT of a shadow PT hierarchy;

[0015] **Figure 12** is a flow diagram of one embodiment of a process for removing an entry from a PT of a shadow PT hierarchy;

20 [0016] **Figure 13** is a flow diagram of one embodiment of a process for monitoring a PTE of a shadow PT hierarchy; and

[0017] **Figure 14** is a flow diagram of one embodiment of a process for removing monitoring from a PTE of a shadow PT hierarchy.

25 Description of Embodiments

[0018] A method and apparatus for supporting address translation in a virtual machine environment is described. In the following description, for purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art that the present invention can be practiced without these specific details.

30 [0019] Some portions of the detailed descriptions that follow are presented in terms of algorithms and symbolic representations of operations on data bits within a computer system's registers or memory. These algorithmic descriptions and

representations are the means used by those skilled in the data processing arts to convey most effectively the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of operations leading to a desired result. The operations are those requiring physical manipulations of physical
5 quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

10 **[0020]** It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussions, it is appreciated that throughout the present invention, discussions utilizing terms such as "processing" or "computing" or "calculating" or "determining" or
15 the like, may refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer-system memories or registers or other such information storage, transmission or display devices.

20 **[0021]** In the following detailed description of the embodiments, reference is made to the accompanying drawings that show, by way of illustration, specific embodiments in which the invention may be practiced. In the drawings, like numerals describe substantially similar components throughout the several views. These embodiments are described in sufficient detail to enable those skilled in the art to practice the invention.

25 Other embodiments may be utilized and structural, logical, and electrical changes may be made without departing from the scope of the present invention. Moreover, it is to be understood that the various embodiments of the invention, although different, are not necessarily mutually exclusive. For example, a particular feature, structure, or characteristic described in one embodiment may be included within other embodiments.

30 The following detailed description is, therefore, not to be taken in a limiting sense, and the scope of the present invention is defined only by the appended claims, along with the full scope of equivalents to which such claims are entitled.

[0022] Although the below examples may describe providing support for address translation in a virtual machine environment in the context of execution units and logic circuits, other embodiments of the present invention can be accomplished by way of software. For example, in some embodiments, the present invention may be provided as a computer program product or software which may include a machine or computer-readable medium having stored thereon instructions which may be used to program a computer (or other electronic devices) to perform a process according to the present invention. In other embodiments, processes of the present invention might be performed by specific hardware components that contain hardwired logic for performing the processes, or by any combination of programmed computer components and custom hardware components.

[0023] Thus, a machine-readable medium may include any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer), but is not limited to, floppy diskettes, optical disks, Compact Disc, Read-Only Memories (CD-ROMs), and magneto-optical disks, Read-Only Memories (ROMs), Random Access Memories (RAMs), Erasable Programmable Read-Only Memories (EPROMs), Electrically Erasable Programmable Read-Only Memories (EEPROMs), magnetic or optical cards, flash memories, a transmission over the Internet, electrical, optical, acoustical or other forms of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.) or the like.

[0024] Further, a design may go through various stages, from creation to simulation to fabrication. Data representing a design may represent the design in a number of manners. First, as is useful in simulations, the hardware may be represented using a hardware description language or another functional description language. Additionally, a circuit level model with logic and/or transistor gates may be produced at some stages of the design process. Furthermore, most designs, at some stage, reach a level of data representing the physical placement of various devices in the hardware model. In the case where conventional semiconductor fabrication techniques are used, data representing a hardware model may be the data specifying the presence or absence of various features on different mask layers for masks used to produce the integrated circuit. In any representation of the design, the data may be stored in any form of a machine-readable medium. An optical or electrical wave modulated or otherwise generated to transmit such information, a memory, or a magnetic or optical storage such as a disc may

be the machine readable medium. Any of these mediums may “carry” or “indicate” the design or software information. When an electrical carrier wave indicating or carrying the code or design is transmitted, to the extent that copying, buffering, or re-transmission of the electrical signal is performed, a new copy is made. Thus, a communication provider or a network provider may make copies of an article (a carrier wave) embodying techniques of the present invention.

[0025] **Figure 1** illustrates one embodiment of a virtual-machine environment 100, in which the present invention may operate. In this embodiment, bare platform hardware 116 comprises a computing platform, which may be capable, for example, of executing a standard operating system (OS) or a virtual-machine monitor (VMM), such as a VMM 112.

[0026] The VMM 112, typically implemented in software, may emulate and export a bare machine interface to higher level software. Such higher level software may comprise a standard or real-time OS, may be a highly stripped-down operating environment with limited operating system functionality, may not include traditional OS facilities, etc. Alternatively, for example, the VMM 112 may be run within, or on top of, another VMM. VMMs may be implemented, for example, in hardware, software, firmware or by a combination of various techniques.

[0027] The platform hardware 116 can be of a personal computer (PC), mainframe, handheld device, portable computer, set-top box, or any other computing system. The platform hardware 116 includes a processor 118 and memory 120.

[0028] Processor 118 can be any type of processor capable of executing software, such as a microprocessor, digital signal processor, microcontroller, or the like. The processor 118 may include microcode, programmable logic or hardcoded logic for performing the execution of method embodiments of the present invention. Although **Figure 1** shows only one such processor 118, there may be one or more processors in the system.

[0029] Memory 120 can be a hard disk, a floppy disk, random access memory (RAM) (e.g., dynamic RAM (DRAM) or static RAM (SRAM)), read only memory (ROM), flash memory, any combination of the above devices, or any other type of machine medium readable by processor 118. Memory 120 may store instructions and/or data for performing the execution of method embodiments of the present invention.

[0030] The VMM 112 presents to other software (i.e., “guest” software) the abstraction of one or more virtual machines (VMs), which may provide the same or different abstractions to the various guests. **Figure 1** shows two VMs, 102 and 114. The guest software running on each VM may include a guest OS such as a guest OS 104 or 106 and various guest software applications 108 and 110. Each of the guest OSs 104 and 106 expects to access physical resources (e.g., processor registers, memory and I/O devices) within the VMs 102 and 114 on which the guest OS 104 or 106 is running and to perform other functions. For example, the guest OS 104 or 106 expects to have access to all registers, caches, structures, I/O devices, memory and the like, according to the architecture of the processor and platform presented in the VM 102 and 114. The resources that can be accessed by the guest software may either be classified as “privileged” or “non-privileged.” For privileged resources, the VMM 112 facilitates functionality desired by guest software while retaining ultimate control over these privileged resources. Non-privileged resources do not need to be controlled by the VMM 112 and can be accessed directly by guest software.

[0031] Further, each guest OS expects to handle various fault events such as exceptions (e.g., page faults, general protection faults, etc.), interrupts (e.g., hardware interrupts, software interrupts), and platform events (e.g., initialization (INIT) and system management interrupts (SMIs)). Some of these fault events are “privileged” because they must be handled by the VMM 112 to ensure proper operation of VMs 102 and 114 and for protection from and among guest software.

[0032] When a privileged fault event occurs or guest software attempts to access a privileged resource, control may be transferred to the VMM 112. The transfer of control from guest software to the VMM 112 is referred to herein as a VM exit. After facilitating the resource access or handling the event appropriately, the VMM 112 may return control to guest software. The transfer of control from the VMM 112 to guest software is referred to as a VM entry.

[0033] In one embodiment, the processor 118 controls the operation of the VMs 102 and 114 in accordance with data stored in a virtual machine control structure (VMCS) 125. The VMCS 125 is a structure that may contain state of guest software, state of the VMM 112, execution control information indicating how the VMM 112 wishes to control operation of guest software, information controlling transitions between the VMM 112 and a VM, etc. The processor 118 reads information from the VMCS 125 to determine the

execution environment of the VM and to constrain its behavior. In one embodiment, the VMCS is stored in memory 120. In some embodiments, multiple VMCS structures are used to support multiple VMs.

5 [0034] During address translation operations, the VM 102 or 114 expects to allocate physical memory, provide protection from and between guest software applications (e.g., applications 108 or 110), use a variety of paging techniques, etc. In a non-virtual machine environment, an address translation mechanism expected by an OS may be based on a translation lookaside buffer (TLB) 122 controlled by the processor 118 and a translation data structure, such as a page-table (PT) hierarchy, controlled by the OS
10 and used to translate virtual memory addresses into physical memory addresses when paging is enabled.

[0035] The architecture of the Intel® Pentium® 4 Processor supports a number of paging modes. The most commonly used paging mode supports a 32-bit linear address space using a two-level hierarchical paging structure (referred to herein as a two-level
15 hierarchy paging mode). Embodiments of the invention are not limited to this paging mode, but instead may be employed by one skilled in the art to virtualize other paging modes (e.g., Physical Address Extension (PAE) mode, Intel® Extended Memory 64 Technology (EM64T) mode, etc.) and implementations (e.g., hashed page tables). In one embodiment based on a TLB, translation of a virtual memory address into a physical
20 memory address begins with searching the TLB 122 using either the upper 20 bits (for a 4 KB page frame) or the upper 10 bits (for a 4MB page frame) of the virtual address. If a match is found (a TLB hit), the upper bits of a physical page frame that are contained in the TLB 122 are conjoined with the lower bits of the virtual address to form a physical address. The TLB also contains access and permission attributes associated with the
25 mapping. If no match is found (a TLB miss), the processor consults the PT hierarchy to determine the virtual-to-physical translation, which is then cached in the TLB 122. Entries in the PT hierarchy may include some attributes that are automatically set by the processor on certain accesses.

[0036] If the PT hierarchy is modified, the TLB 122 may become inconsistent with
30 the PT hierarchy if a corresponding address translation exists in the TLB 122. The OS may expect to be able to resolve such an inconsistency by issuing an instruction to the processor 118. For example, in the instruction set architecture (ISA) of the Intel® Pentium® 4 (referred to herein as the IA-32 ISA), a processor allows software to invalidate

cached translations in the TLB by issuing the INVLPG instruction. In addition, the OS may expect to request the processor 118 to change the address space completely, which should result in the removal of all translations from the TLB 122. For example, in the IA-32 ISA, an OS may use a MOV instruction or a task switch to request a processor to load
5 CR3 (which contains the base address of the PT hierarchy), thereby removing all translations from the TLB. Different levels of the page table hierarchy may have different names based upon mode and implementation. In the two-level hierarchy paging mode, there are two levels of paging structures. The CR3 register points to the base of the page directory page. Entries in the page directory may either specify a mapping to a large-size
10 page (e.g., a 4MB superpage, a 2MB superpage, 1GB superpage, etc.), or a reference to a page table. The page table in turn may contain mappings to small-size pages.

[0037] As discussed above, in the virtual-machine environment, the VMM 112 should be able to have ultimate control over physical resources including the TLB 122. Embodiments of the present invention address the conflict between the expectations of the
15 VMs 102 and 114 and the role of the VMM 112 by using a virtual TLB that emulates the functionality of the processor's physical TLB.

[0038] The virtual TLB includes the TLB 122 and a set of shadow PT hierarchies controlled by the VMM 112. The set of shadow PT hierarchies derive its format and content from guest PT hierarchies that may be currently used or not used by the VM 102
20 or 114. If the VM 102 or 114 modifies the content of the guest PT hierarchies, this content becomes inconsistent with the content of the shadow PT hierarchies. The inconsistencies between the guest PT hierarchies and the shadow PT hierarchies are resolved using techniques analogous to those employed by the processor 118 in managing the TLB 122. Some of these techniques force the VM 102 or 114 to issue an event
25 indicating an attempt to manipulate the TLB (e.g., INVLPG, page fault, and load CR3). Such events are privileged and, therefore, result in a VM exit to the VMM 112. The VMM then evaluates the event and synchronizes all maintained shadow PT hierarchies with the current guest state if needed. We will refer to the set of maintained shadow PT hierarchies as the working set. As multiple processes may use the same guest page table, it
30 is possible for the same shadow PT to be a part of multiple guest PT hierarchies. The corresponding shadow PT will in turn be a member of multiple shadow PT hierarchies.

[0039] Note that synchronization performed by the VMM may update shadow page table or page directory entries for a shadow PT hierarchy that is not currently in-use.

Likewise synchronization may be required to guest pages that are not part of the in-use guest PT hierarchy.

[0040] In one embodiment, the VMM 112 includes an address translation module 126 that is responsible for creating and maintaining a working set of shadow PT hierarchies for each of the VM 102 and 114 in a virtual TLB (VTLB) data store 124. The working set of shadow PT hierarchies is maintained for corresponding active processes of the VM 102 or 114 (i.e., processes that are likely to be activated in the near future by the VM 102 or 114). With the IA32 ISA, the only explicitly defined guest hierarchy is that defined by the currently used paging structures. In practice there is a high deal of temporal locality for guest processes and their address spaces. The VMM may employ heuristics or explicit information to determine a set of active process.

[0041] When the VM 102 or 114 enables a guest PT hierarchy for one of the active processes of the VM 102 or 114, the address translation module 126 identifies a corresponding shadow PT hierarchy in the working set and requests the processor 118 to load its base address. When applicable, the address translation module 126 can then reuse previously computed mappings that are stored in the shadow PT hierarchies.

[0042] If the VM 102 or 114 activates a new process, the address translation module 126 derives a new shadow PT hierarchy from a corresponding guest PT hierarchy and adds it to the working set. Alternatively, if the VM 102 or 114 de-activates an existing process, the address translation module 126 removes information corresponding to the guest PT hierarchy from the working set.

[0043] In one embodiment, the address translation module 126 is responsible for extracting metadata from each new shadow PT hierarchy, storing the metadata in the VTLB data store 124, and updating the metadata when the shadow PT hierarchy is modified. In one embodiment, the metadata includes a PT vector (PTV), a PD vector (PDV), an active PTE list, and an active PDE list.

[0044] The PTV and PDV track the guest frames that are used as PTs and PDs. In one embodiment, this information is encoded in bit vectors. The PTV may be indexed by page frame number (PFN), with each entry bit being set if a corresponding PFN is a PT. The PDV may be indexed by a page frame number (PFN), with each entry bit being set if a corresponding PFN is a PD.

[0045] The active PTE list is a list of PT entries (PTEs) in the shadow PT hierarchy that point to frames holding PTs and PD. The active PDE list identifies PD

entries (PDEs) in the shadow PT hierarchy that point to PTs containing PT entries identified in the active PTE list.

[0046] In one embodiment, active PDE and PTE lists contain additional metadata describing whether the mapping is to a PD or PT frame.

5 [0047] One skilled in the art will understand that embodiments of this invention may use a variety of data structures which may be more or less space or time efficient than those described herein. One skilled in the art will also recognize the extension of tracking structures to support additional paging modes. For example, an EM64T paging mode maps a 64-bit virtual address to a physical address through a four-level hierarchical paging
10 structure. The actual number of bits supported in the virtual or physical address spaces may be implementation dependent and may be less than 64 bits in a particular implementation. As will be discussed in more detail below, an EM64T implementation may require additions of a page-map level 4 (PML4) page vector and a page directory pointer (PDP) page vector to track the additional page tables used in the EM64T paging
15 structure. Likewise, one skilled in the art will recognize that the active PTE list will be extended to include entries which map any page used within the paging structures (e.g., PML4 or PDP pages for EM64T).

[0048] In one embodiment, active PTE/PDE list metadata is maintained to track the number of PD and PT frames that are mapped through a page table. When the number
20 of mappings per page is incremented from 0, then PDEs which map the PT must be added to the active PDE list, and when the number of mappings is decreased to zero, then PDEs that map this PT must be removed from the active PDE list.

[0049] In one embodiment, the address translation module 126 is responsible for synchronizing a current shadow PT hierarchy with a current guest PT hierarchy when such
25 synchronization is needed. The address translation module 126 performs the synchronization by determining which entries in the guest PT hierarchy have recently been modified and then updating corresponding entries in the shadow PT hierarchy accordingly. The address translation module 126 determines which entries in the guest PT hierarchy have recently been modified based on the metadata extracted from the shadow PT
30 hierarchy and attributes associated with the entries of the shadow PT hierarchy. In one embodiment, the attributes include access attributes associated with PD entries in the shadow PT hierarchy and update attributes associated with PT entries in the shadow PT hierarchy.

[0050] Figure 2 illustrates operation of a virtual TLB 204, according to one embodiment of the present invention. Virtual TLB 204 includes a shadow translation data structure represented by a shadow PT hierarchy 206 and a physical TLB 208. The shadow PT hierarchy 206 derives its structure and content from a guest translation data structure represented by a guest PT hierarchy 202. In one embodiment, the VMM maintains a working set of shadow PT hierarchies for active processes of the VM.

[0051] In one embodiment, when the VM requests the processor to enable a different guest PT hierarchy (e.g., by issuing MOV to CR3 or task switch in the IA-32 ISA), control transitions to the VMM, which instructs the processor to load the base address 214 of a shadow PT hierarchy 206 corresponding to the requested guest PT hierarchy 202. In some embodiments, this shadow PT hierarchy 206 is synchronized with the guest PT hierarchy 202 using relevant metadata and attributes, as will be discussed in greater detail below.

[0052] In one embodiment, the virtual TLB maintains access and update attributes in the entries of the shadow PD and PTs. These attributes are also referred to as an accessed (A) bit and a dirty (D) bit. In one embodiment, when a page frame is accessed by guest software for the first time, the processor sets the accessed (A) attribute in the corresponding PT entry or PD entry in the shadow PT hierarchy 206. If guest software attempts to write a page frame, the processor sets the dirty (D) attribute in the corresponding shadow PT entry.

[0053] Guest software is allowed to freely modify the guest PT hierarchy 202 including changing virtual-to-physical mapping, permissions, etc. Accordingly, the shadow PT hierarchy 206 may not be always consistent with the guest PT hierarchy 202. When a problem arises from an inconsistency between the hierarchies 202 and 206, the guest OS, which treats the virtual TLB 204 as a physical TLB, attempts to change the virtual TLB 204 by requesting a processor to perform an operation defined by a relevant ISA. For example, in the IA-32 ISA, such operations include the INVLPG instruction, CR3 loads, paging activation (modification of CR0.PG), modification of global paging (toggling of the CR4.PGE bit), etc. The operations attempting to change the virtual TLB 204 are configured by the VMM as privileged (e.g., using corresponding execution controls stored in the VMCS), and, therefore, result in a VM exit to the VMM. The VMM then determines the cause of the VM exit and modifies the content of the shadow PT hierarchy 206 if necessary. For example, if the VM exit occurs due to a page fault that

should be handled by the guest OS (e.g., a page fault caused by an access not permitted by the guest PT hierarchy 202), the page fault is injected to the guest OS for handling.

Alternatively, if the VM exit occurs due to a page fault (or any other operations such as INVLPG) resulting from an inconsistency between the entries of the hierarchies 202 and 206, the VMM may need to remove stale entries, add new entries, or modify existing entries, as will be discussed in more detail below. Page faults caused by the guest PT hierarchy are referred to herein as 'real' page faults, and page faults that would not have occurred with direct usage of the guest page tables are referred to herein as 'induced' page faults.

[0054] Figure 3A illustrates a process of creating and maintaining metadata for a shadow PT hierarchy in a two-level hierarchy paging mode, according to one embodiment of the present invention.

[0055] Referring to Figure 3A, a number of physical page frames identified by distinct letters (letters A through W) is illustrated. Some guest page frames may contain a PD (e.g., frame A). Other guest page frames may contain a PT (e.g., frames A, B, C, and L). A hierarchy 302 is a guest PT hierarchy.

[0056] Figure 3A shows a shadow PT hierarchy 304 created based on a guest PT hierarchy 302. Each PD or PT in the guest PT hierarchy 302 includes a corresponding PD or PT in the shadow PT hierarchy 304. Note that in general a shadow page is not required for each page in the guest PT. Some embodiments may choose to restrict shadow pages according to usage statistics (e.g., only generate shadow pages for guest PT pages that have been used), or according to resource constraints (e.g., maintaining only a set of shadow pages based on available memory). Separate shadow tables are maintained for PD and PT tables derived from the same physical frame. For example, separate tables 330 and 332 are maintained for PD 306 and PT 308 that are derived from the same physical frame 314. The PD and PT entries in the shadow PT hierarchy 304 contain transformed mappings for the guest frames 314 through 324.

[0057] In the guest PT hierarchy 302, frames 316 and 318 are used as PTs 310 and 312, and frame 314 is used both as PD 306 and PT 308. This usage is illustrated as "PT" and "PD/PT" in the page frames 314 through 316 shown under the shadow PT hierarchy 304.

[0058] The shadow PT hierarchy 304 is associated with an active PTE list 342 and an active PDE list 344. In one embodiment, the active PTE list 342 identifies PT entries in

the shadow PT hierarchy 304 that map PT and PD page frames from the guest PT hierarchy 302. In particular, the active PTE list 342 identifies entries in the PT 332 that map page frames 314 through 318. In one embodiment, the active PDE list 344 identifies PD entries in the shadow PT hierarchy that point to PTs with entries identified in the active PTE list 342. In particular, the active PDE list 344 includes entries in the PD 330 that point to the PT 332. The active PTE list 342 and the active PDE list 344 are components of the metadata of the shadow PT hierarchy 304.

[0059] The shadow PT hierarchy 304 is associated with a PT bit vector (PTV) 362 and a PD bit vector (PDV) 364. In one embodiment, the PTV 362 tracks the guest page frames that are used as PTs. In particular, the PTV 362 includes page frames 314 through 318 which are used as PTs in the guest PT hierarchy 302. In one embodiment, the PDV 364 tracks the guest page frames that are used as PDs. In particular, the PDV 364 includes page frame 314 that is used as PD in the guest PT hierarchy 302. In one embodiment, the PTV 362 and PDV 364 represent all shadow PT hierarchies in the working set and track the capacity in which shadow pages are employed in the working set (e.g., if a shadow page has not been allocated for a guest PT, then the PTV will not reflect the guest PT page, even if it appears in the guest paging structures).

[0060] In one embodiment, if the guest OS adds a new PT to the guest PT hierarchy 302, the VMM may detect this addition (e.g., on the next or subsequent VM exit related to TLB manipulation) and add a corresponding PT to the shadow PT hierarchy 304. For example, if a new PT 352 derived from a frame 319 is added to the guest PT hierarchy 302, with a mapping for a new frame 354, the VMM may add a corresponding PT 360 with transformed mappings to the shadow PT hierarchy 304 and update the metadata to reflect this change. In particular, the VMM adds an entry mapping frame 319 in the PT 332 to the active PTE list 342, and an entry pointing to the PT 360 in the PD 330 to the active PDE list 344. Also, the VMM adds frame 319 to PTV 362, which tracks guest frames (i.e., here frame 319) used as PTs.

[0061] Figure 3B illustrates a process of creating and maintaining metadata for a shadow PT hierarchy in the EM64T paging mode, according to one embodiment of the present invention.

[0062] Referring to Figure 3B, the base of the paging structure is a PML4 page (e.g., frame A). Each entry in the PML4 page may reference a PDP page (e.g., frames B and C). Each entry in the PDP page may reference a page directory (PD) page (e.g., frame

D or E), each entry of which in turn may reference a page in a page table (PT) page (e.g., frame F, G, H or I).

[0063] Each PML4, PDP, PD or PT page may be 4KB in size. In order to support physical address spaces larger than 32 bits, the entry size may be increased relative to the 32-bit paging mode. Specifically, there may be 512 entries per page, requiring that 9 bits of the virtual address be used at each level to select the appropriate entry. This selector size may lead to a large page size of 2MB instead of 4MB as described previously.

[0064] In one embodiment, the creation of metadata in the EM64T paging mode includes the generation of several vectors, an active entry list, and several active directory lists. The vectors include a PML4V vector identifying frames used as PML4 pages, a PDPV vector identifying frames used as PDP pages, a PDV vector identifying frames used as PD pages, and a PTV vector identifying frames used as PT pages. The active entry list is an active PTE list including all mappings which map a PML4, PDP, PD or PT page. The active directory lists include lists identifying higher level mapping structures referencing a lower level structure through which the guest page corresponding to a shadow structure can be accessed. In particular, the active directory lists consist of an active PDE list including those PDEs that reference a page containing active PTE list entries, an active PDPE list including active PDPE entries which reference a PD containing an active PDE list entry, and an active PML4E list including entries which map a PDP containing elements in the active PDPE list.

[0065] In one embodiment, the synchronization of the shadow page tables begins with checking each entry in the active PML4E list associated with the used shadow PT hierarchy. If the entry has been accessed, each element in the active PDPE list corresponding to the accessed PML4 entry is checked, and then the processing continues as previously described.

[0066] In an alternative embodiment, active lists are not maintained and/or processed for one or more of the upper levels of the hierarchy. For example, in a system in which only a single entry is populated in the uppermost paging structure, the use of an active list for each level of the hierarchy will cause this single entry to be always accessed, thereby allowing no reduction in the amount of processing required for lower levels in the hierarchy. To accommodate this usage model, the synchronization may instead begin by processing an active list lower in the hierarchy. For example, in one embodiment, active PDPE list elements may first be processed followed by active PDE list elements or active

PTE list elements associated with a used shadow PT hierarchy. In one embodiment, the initial layer processed on synchronization may be predetermined. In another embodiment, the initial layer to be processed may be determined by dynamic profiling of the guest's page table usage.

5 **[0067]** Various other paging modes may be used with embodiments of the present invention. For example, IA-32 supports an additional paging mode in which a 32-bit virtual address is mapped to a larger physical address. In this additional mode of operation, the page table base register is configured to point to a PDP page which contains four elements. Entry sizes and behaviors in this additional mode of operations are similar
10 to those described above for the 64-bit virtual address mode. As this additional mode does not make use of PML4 pages, the PML4V and active PML4E list are not required.

[0068] **Figure 4** is a flow diagram of one embodiment of a process 400 for synchronizing a guest translation data structure and a shadow translation data structure. The process may be performed by processing logic that may comprise hardware (e.g.,
15 dedicated logic, programmable logic, microcode, etc.), software (such as that run on a general purpose computer system or a dedicated machine), or a combination of both. In one embodiment, the process is performed by an address translation module 126 of **Figure 1**.

[0069] Referring to **Figure 4**, process 400 begins with processing logic receiving
20 control transitioned from a VM due to an event pertaining to manipulation of the TLB (processing block 402). Examples of such events may include a request to change the current address space (e.g., CR3 load), a request to adjust inconsistent translations for a specified virtual address in the TLB (e.g., INVLPG), a page fault, etc.

[0070] At processing box 404, processing logic determines whether the event
25 pertaining to the manipulation of the TLB should be handled by the VM. If so (e.g., the event is a page fault caused by a problematic mapping in a guest translation data structure), control is returned to the VM for handling the event (processing block 406). If not, processing logic determines whether the event is associated with a specified problematic address (processing box 408).

30 **[0071]** If the event does not need to be handled by the VM, it may be associated with a specified problematic address. Examples of such an event may include an event caused by the INVLPG instruction that takes a problematic address as an operand, an event caused by an induced page fault (e.g., a page fault resulting from an inconsistency

between the two translation data structures with respect to a specific mapping, a page fault caused by a need to virtualize A/D bits in the guest translation data structure, etc.), etc. If the event is associated with a specified problematic address, processing logic makes corrections in the shadow translation data structure for the specified address (e.g., removes a stale mapping for the specified address or adds a new mapping for the specified address) to conform to the guest translation data structure (processing block 410). One embodiment of a process for synchronizing entries of two translation data structures for a specified address is discussed in more detail below in conjunction with **Figure 7**.

[0072] If the event is not associated with any specific address (e.g., the event is caused by a request of the VM to change the address space, which flushes all TLB entries in IA32), processing logic determines which entries of the guest translation data structure have been modified (processing block 412). The determination is made using metadata extracted from the shadow translation data structure and attributes associated with the entries of the shadow translation data structure (processing block 412). The metadata includes vectors and active lists for various levels of the shadow translation data structure. A vector for a specific level of the shadow translation data structure identifies frames used as pages at this level of the guest translation data structure. The active lists include an active entry list and one or more active directory lists. The active entry list includes mappings that map pages used by the guest in forming the guest translation data structure. The active directory lists identify higher level mapping structures referencing a lower level structure through which a guest page corresponding to a shadowed paging structure can be accessed. As discussed above, in the two-level hierarchy paging mode, the metadata includes, in one embodiment, vectors PTV and PDV, an active entry list (a PTE list), and an active directory list (a PDE list). In the EM64T paging mode, the metadata includes, in one embodiment, vectors PTV, PDV, PDPV and PML4V, an active entry list (a PTE list), and active directory lists (an active PDE list, an active PDPE list and an active PML4E list).

[0073] One embodiment of a process for identifying recently modified entries of the guest translation data structure using metadata is discussed in more detail below in conjunction with **Figure 6**.

[0074] At processing block 414, processing logic synchronizes corresponding entries in the shadow translation data structure with the modified entries of the guest translation data structure. Accordingly, processing logic only needs to synchronize the

entries that were modified, rather than re-populating the entire content of the shadow translation data structure.

[0075] In one embodiment extra storage is used to maintain some guest PD and/or PT contents as they were last synchronized. This permits the VMM to determine where
5 modifications have been made without calculating or looking up additional relocation or permission information.

[0076] Note that certain modifications to the guest page tables do not require modifications to the shadow page tables. For example, if a guest PT contains a not present mapping which is subsequently modified, no change is required to the corresponding
10 shadow PT.

[0077] **Figures 5-14** illustrate various processes performed to support address translation in a virtual machine environment using the two-level hierarchy paging mode, according to different embodiments of the present invention. These processes may be performed by processing logic that may comprise hardware (e.g., dedicated logic,
15 programmable logic, microcode, etc.), software (such as that run on a general purpose computer system or a dedicated machine), or a combination of both. In one embodiment, each of these processes is performed by an address translation module 126 of **Figure 1**.

[0078] **Figure 5** is a flow diagram of one embodiment of a process 500 for maintaining metadata for a shadow translation data structure such as a shadow PT
20 hierarchy.

[0079] Referring to **Figure 5**, process 500 begins with processing logic creating a shadow page for each PD or PT page from the guest PT hierarchy (processing block 502).

[0080] At processing block 504, processing logic tracks page frames used as PDs or PTs in the guest PT hierarchy. In one embodiment, processing logic sets an entry in the
25 PDV if a corresponding PFN is a PD in the guest PT hierarchy. Similarly, processing logic sets an entry in the PTV if a corresponding PFN is a PT in the guest PT hierarchy.

[0081] At processing block 506, processing logic tracks mappings to any Dynamic Random Access Memory (DRAM) backed page (to identify pages that can potentially be PDs or PTs). In one embodiment, processing logic tracks mappings to DRAM based
30 pages using an inverted page table (IPT) and an inverted page directory (IPD). The IPT is indexed by a PFN of a data page frame, with each entry containing a list of addresses of PTEs that map the data page frame. The IPD is indexed by a PFN of the page table, with each entry containing a list of addresses of PDEs that reference the PFN as a page table.

[0082] In one embodiment, at processing block 508, processing logic identifies 4MB pages in the guest PT hierarchy and creates a page table in the shadow PT hierarchy for each 4MB page to avoid large page mappings and thereby reduce future synchronization time. Otherwise, an update of a 4MB page would cause the
5 synchronization of every PD and PT page within the 4MB. In one embodiment, an inverted expansion table (IET) is used to track which PDEs in the guest PT hierarchy point to a 4MB page. The IET is indexed by a PFN and attribute bits, with every entry listing PDEs that point to the exploded 4MB page.

[0083] In an embodiment of the invention the IPD may be indexed by the address
10 of the shadow PFN to minimize required address translation steps.

[0084] In IA32, memory type information (e.g., cacheability information) can be stored in PAT bits within the PDE/PTE that maps a page. This type information is not captured in a PDE that is a page-table pointer. Hence, if two 4MB pages were to map the same region with different PAT attributes, then separate page tables would be required to
15 convey the correct PAT attributes. Using separate expansion tables for each set of attributes resolves this issue.

[0085] At processing block 510, processing logic identifies PTEs in the shadow PT hierarchy that map pages used as PD or PT in the guest PT hierarchy and creates an active PTE list.

[0086] At processing block 512, processing logic identifies PDEs in the shadow
20 PT hierarchy that point to PTs with PTEs identified in the active PTE list and creates an active PDE list.

[0087] Subsequently, at processing block 514, if the guest OS modifies the structure of the guest PT hierarchy (e.g., adds or removes a PD or PT), processing logic
25 changes the above active PTE and PDE lists accordingly.

[0088] **Figure 6** is a flow diagram of one embodiment of a process 600 for facilitating a change of an address space. Note that in IA32 the same CR3 value may also be reloaded to force a flush of stale TLB mappings. Similar processing steps are taken for a change of CR3 or for a CR3 reload.

[0089] Referring to **Figure 6**, process 600 begins with processing logic
30 determining that a VM exit occurred due to a request of the VM to enable a different guest PT hierarchy (e.g., by issuing a CR3 load request) (processing block 602).

[0090] In response, processing logic scans all active PDEs corresponding to the currently in-use shadow PT hierarchy identified in the active PDE list of the metadata to find which of these PDEs have been accessed (have an access attribute set to an access value) (processing block 604), and then initializes the access attributes of the accessed PDEs (processing block 606). In IA32, non-leaf paging tables do not support a dirty bit. If the accessed bit is clear, then no page within the 4MB region has been read or written, so any guest page table or page directory cannot have been modified. However, the accessed bit does not distinguish between reads and writes, so 4MB regions which have been accessed should be further processed even though it is possible that nothing has been modified. In architectures supporting a dirty bit for non-leaf page tables, the dirty bit is checked instead, and only regions which had been written to require further processing.

[0091] Next, for each accessed PDE, processing logic scans all shadow PTEs corresponding to the accessed active PDE in the active PTE list of the metadata to find which of these PTEs include mappings for an updated page (have an update attribute set to an update value) (processing block 608).

[0092] Further, for each updated page, processing logic compares PD/PT entries in the guest PT hierarchy with corresponding entries in the shadow PT hierarchy (processing block 610) and changes the corresponding entries of the shadow PT hierarchy to conform to the modified entries of the guest PT hierarchy (e.g., by removing from the shadow PT hierarchy a PTE/PDE absent in the guest PT hierarchy, by adding to the shadow PT hierarchy a new PTE/PDE recently added to the guest PT hierarchy, etc.) (processing block 612). Note that adding PDEs may require the allocation and initialization of additional shadow PTs. This in turn may require updates to the various metadata structures maintained by the address translation module 126.

[0093] At processing block 614, processing logic initializes update attributes that were set to an update value. Updated mappings identify the pages that were modified by the guest OS.

[0094] At processing block 616, processing logic synchronizes the shadow mappings based on modified guest pages and updates the metadata if needed due to the above modifications.

[0095] At processing logic 618, processing logic determines whether a working set maintained by the VMM includes a shadow PD corresponding to the new guest PD requested by the VM. If so, processing logic requests the processor to load the base

address of this shadow PT hierarchy (processing block 620). If not, processing logic allocates a new shadow PT hierarchy corresponding to the requested guest PT hierarchy (processing block 622), adds the PD of the new shadow PT hierarchy to the PDV (processing block 624), adds each valid PDE to the PD of the new shadow PT hierarchy (processing block 626), configures the active PDE and PTE lists to monitor the PTEs that map this PD for PD coverage (processing block 628), and then requests the processor to load the base address of this shadow PT hierarchy (processing block 620). One embodiment of a process for monitoring a PTE is discussed in more detail below in conjunction with **Figure 13**.

[0096] **Figure 7** is a flow diagram of one embodiment of a process 700 for synchronizing entries of two translation data structures for a specified address. Process 700 may be performed, for example, as a result of the INVLPG instruction issued by the VM or as a result of an induced page fault.

[0097] Referring to **Figure 7**, process 700 begins with processing logic determining whether the mapping in the shadow PT hierarchy for the specified address is stale (i.e., there is valid mapping that does not correspond to the current contents of the guest page table) (processing box 702). If not, processing logic proceeds to processing box 712. If so, processing logic determines whether the stale entry mapped a PD or PT page (processing box 704). If the stale entry did not map a PD or PT page, processing logic removes the stale entry (processing block 710) and proceeds to processing box 712.

[0098] If the stale entry did map a PD or PT page, processing logic further determines whether the mapped page has been updated (processing box 706). If not, processing logic proceeds to processing block 710. If so, processing logic updates, synchronizes, or removes the modified PD or PT shadow(s) (processing block 708) and proceeds to processing block 710. In one embodiment, the page is marked for future synchronization.

[0099] At processing box 712, processing logic determines whether the guest PT hierarchy contains a new mapping for the specified address. If not, process 700 ends. If so, processing logic adds the new mapping as a corresponding PTE or PDE and, if necessary, creates a shadow page and updates the metadata according to the addition (processing block 714).

[00100] **Figure 8** is a flow diagram of one embodiment of a process 800 for removing a shadow PT hierarchy from a working set of shadow PT hierarchies maintained by the VMM.

[00101] A shadow PT hierarchy may be removed from the working set upon
5 detecting a deactivation of a corresponding process by the VM. The deactivation may be detected using heuristic defined for a relevant OS or employing a set of checks based on clues provided by the behavior of the guest VM with respect to the current address space. If the VM supports an interface through which the OS or a driver notifies the VMM of deactivations, then a heuristic may be avoided. A shadow PT hierarchy may also be
10 removed due to resource constraints, e.g., because the amount of memory used for shadow structures exceeds a target threshold.

[00102] Referring to **Figure 8**, processing logic begins with removing each valid PDE from the PD in the shadow PT hierarchy (processing block 802).

[00103] At processing block 804, processing logic clears a corresponding
15 entry in the PDV.

[00104] At processing block 806, processing logic deallocates the PD page and removes the translation from a PD translation table (PDTT). The PDTT is used to track the address and type (e.g., PD or PT) of a page. The PDTT is indexed by a guest PFN, with each entry containing a physical PFN and metadata.

[00105] At processing block 808, processing logic removes monitoring from
20 the PTEs that map the PD. One embodiment of a process for removing monitoring from a PTE is discussed in more detail below in conjunction with **Figure 14**.

[00106] **Figure 9** is a flow diagram of one embodiment of a process 900 for adding an entry to a PD of a shadow PT hierarchy. For illustration, we will consider a
25 present entry which maps a page table.

[00107] Referring to **Figure 9**, processing logic begins with adding an entry for the PDE to the IPD (processing block 902).

[00108] At processing box 904, processing logic determines if the PT mapped by this PDE is set in the PTV. If so, the appropriate shadow PT is looked up in
30 the PTTT (processing block 916), the new shadow PDE is created (processing block 914) and process 900 ends. If not, processing logic sets a corresponding vector in the PTV (processing block 906), allocates a shadow page and initializes the translation (processing block 908), populates the new shadow page table (processing block 910), updates active

PTE/PDE lists and metadata to reflect that the guest page used as a page table by the current guest PDE is to be monitored (processing block 912), and adds the new PDE, adding it to the active PDE list if the shadow page table contains any active PTE list elements (processing block 914). One embodiment of a process for monitoring a PTE is discussed in more detail below in conjunction with **Figure 13**.

[00109] **Figure 10** is a flow diagram of one embodiment of a process 1000 for removing an entry from a PD of a shadow PT hierarchy.

[00110] Referring to **Figure 10**, processing logic begins with removing an entry for this PDE from the IPD PDE list (processing block 1002). If the PDE is in the active PDE list, then the active PDE list must be updated.

[00111] At processing box 1004, processing logic determines whether the PDE was the last entry to map the corresponding PT. If not, process 1000 ends. If so, processing logic clears the entry for the PT in the PTV (processing block 1006), removes each valid PTE (processing block 1008), updates the active PTE/PDE lists that map this PT for PT coverage (processing block 1010), and removes the shadow page translation and free the memory used to store the PT shadow page (processing block 1010).

[00112] **Figure 11** is a flow diagram of one embodiment of a process 1100 for adding an entry to a PT of a shadow PT hierarchy.

[00113] Referring to **Figure 11**, processing logic begins with adding an entry for this PTE to the IPT (processing block 1102).

[00114] At processing box 1106, processing logic creates the shadow mapping and proceeds to processing box 1108.

[00115] At processing box 1108, processing logic determines whether a corresponding entry in the PDV or PTV is set. If not, process 1100 ends. If so, processing logic adds this entry to the active PTE list and updates associated metadata indicating if it maps a PD and/or PT page (processing block 1110). If the active PTE entry just created is the first for this page table, then the IPD must be consulted and each PDE which maps this page table page added to the active PDE list.

[00116] **Figure 12** is a flow diagram of one embodiment of a process 1200 for removing an entry from a PT of a shadow PT hierarchy.

[00117] Referring to **Figure 12**, processing logic begins with determining whether this PTE maps a page set in the PDV or PTV (processing block 1202). If not, processing logic proceeds to processing block 1206. If so, processing logic removes the

PTE from the active PTE list. If this was the last active PTE list element in the PT, then PDEs referencing this PT are removed from the active PDE list. (processing block 1204), and proceeds to processing block 1206.

[00118] At processing block 1206, processing logic removes the
5 corresponding entry from the IPT.

[00119] **Figure 13** is a flow diagram of one embodiment of a process 1300 for monitoring a PTE of a shadow PT hierarchy. The steps shown in **Figure 13** represent the processing that may be required when the monitor recognizes that a page which has been mapped as a data page is being used as a page directory or page table page. This
10 process will therefore be triggered by a status change for the page mapped by the PTE.

[00120] Referring to **Figure 13**, processing logic begins with determining whether the PTE is identified in the active PTE list (processing box 1302). If so, processing logic adds the previously missing coverage (processing block 1304). This flow is triggered by a status change of the mapped page. Since this PTE was already in the
15 active PTE list, it must be the case that this PTE was previously in use as a PT or PD, and is now in use in the other capacity as well. Such information may be explicitly stored with the entry or in associated metadata. If the PTE is not in the active PTE list, processing logic adds the PTE to the active PTE list and updates metadata accordingly (processing block 1306).

[00121] Next, at processing box 1308, processing logic determines whether the PTE is the first active PTE list entry for this PT. If not, process 1300 ends. If so, processing logic adds, to the active PDE list, entries that map this PT (as found through the IPD) (processing block 1310).
20

[00122] **Figure 14** is a flow diagram of one embodiment of a process 1400 for decreasing the monitoring coverage provided by a PTE of a shadow PT hierarchy. This process might be invoked when a process is removed from the working set, or the last PDE to reference a page table is removed, resulting in a change of status of a previously monitored page directory or page table page.
25

[00123] Referring to **Figure 14**, processing logic begins with determining whether this PTE had monitored a page that was both a page table and page directory page (processing box 1402). If so, processing logic reduces the coverage level, indicating that the PTE now monitors a page as either a PT or as a PD, but not both (processing block 1404). If not, processing logic removes the PTE from the active PTE list (processing
30

block 1406). Note that if the PTE was an element for a page tracked for its use in a single capacity, then it must now be the case that the page no longer requires monitoring.

[00124] Next, if the last active PTE list element in the PT was removed (processing box 1408), processing logic removes the corresponding entries which mapped
5 this page table from the active PDE list (as found through the IPD) (processing block 1410).

[00125] As discussed above, the physical or virtual platform may comprise multiple processors. Each processor may in turn comprise one or more threads or logical processors. The processes discussed above can be used in a single-threaded system
10 supporting a single-threaded VM or in a physical system with multiple logical processors that supports one or more VMs each containing a single virtual logical processor. Note that each VM has its own set of metadata, shadow page tables, etc. and that synchronization steps are confined to a given VM.

[00126] Thus, a method and apparatus for supporting address translation in a
15 virtual machine environment have been described. It is to be understood that the above description is intended to be illustrative, and not restrictive. Many other embodiments will be apparent to those of skill in the art upon reading and understanding the above description. The scope of the invention should, therefore, be determined with reference to the appended claims, along with the full scope of equivalents to which such claims are
20 entitled.

CLAIMS

What is claimed is:

1. A method comprising:

5 receiving control transitioned from a virtual machine (VM) due to a privileged event pertaining to a translation-lookaside buffer (TLB);
determining which entries in a guest translation data structure were modified by the VM, based on metadata extracted from a shadow translation data structure maintained by a virtual machine monitor (VMM) and attributes associated with entries in the shadow translation data structure; and
10 synchronizing entries in the shadow translation data structure that correspond to the modified entries in the guest translation data structure with the modified entries in the guest translation data structure.

2. The method of claim 1 wherein:

15 the guest translation data structure is used by the VM for address translation operations; and
content of the shadow translation data structure is used by a processor to cache address translations in the TLB.

20 3. The method of claim 1 further comprising:

maintaining a working set of shadow translation data structures, each shadow translation data structure in the working set being associated with one of a plurality of active processes of the VM; and
reusing content of a shadow translation data structure from the working set that is
25 associated with one of the plurality of active processes when a guest translation data structure associated with the one of the plurality of active processes is enabled.

4. The method of claim 3 wherein the privileged event is any one of a request of the VM to enable a different guest translation data structure, a page fault caused by one or
30 more inconsistencies between entries of the guest translation data structure and entries of the shadow translation data structure, and a request of the VM to invalidate one or more address translations in a translation-lookaside buffer (TLB).

5. The method of claim 3 further comprising:

determining that the privileged event is any one of a page fault, an INVLPG instruction, and a request to enable a new guest translation data structure;

creating a new shadow translation data structure based on the new guest translation data structure; and

deriving metadata from the new shadow translation data structure.

6. The method of claim 3 further comprising:

determining that one of the plurality of the active processes of the VM is

deactivated; and

removing a shadow translation data structure associated with the active process being deactivated from the working set.

7. The method of claim 3 further comprising:

determining one of the shadow translation data structures is no longer used; and

removing the one of the shadow translation data structures from the working set.

8. The method of claim 1 wherein the metadata includes a set of vectors, an active entry list, and one or more active directory lists.

9. The method of claim 8 wherein:

each vector in the set identifies frames used as pages at a corresponding level of the guest translation data structure;

the active entry list identifies mappings that map pages used in forming the guest translation data structure for which a shadow translation data structure exists; and

the one or more active directory lists identify higher level mapping structures referencing a lower level structure through which the shadow translation data structure can be accessed.

10. The method of claim 9 wherein:

the active entry list is an active page table (PT) entry list, the active PT entry (PTE) list identifying PTEs in the shadow translation data structure that map PT pages and PD pages from the guest translation data structure, and

the one or more active directory lists include an active page directory (PD) entry list, the active PD entry (PDE) list identifying PDEs in the shadow translation data structure that point to PTs having the PTEs from the active PTE list.

5 11. The method of claim 9 wherein:

the active entry list is an active page table (PT) entry list, the active PT entry (PTE) list identifying mappings that map any of page map level 4 (PML4) pages, page directory pointer (PDP) pages, page directory (PD) pages, and PT pages; and

10 the one or more active directory lists include an active PD entry (PDE) list containing PDEs that reference a page with active PTE list entries, an active PDP entry (PDPE) list containing active PDPE entries which reference a PD with an active PDE list entry, and an active PML4E entry (PML4E) list containing entries which map a PDP with elements from the active PDPE list.

15 12. The method of claim 10 wherein the attributes associated with entries in the shadow translation data structure include access attributes associated with PDEs in the shadow translation data structure and update attributes associated with PTEs in the shadow translation data structure.

20 13. The method of claim 12 wherein determining which entries in the guest translation data structure were modified by the VM comprises:

identifying one or more PDEs from the active PDE list corresponding to the active PT hierarchy being synchronized that have access attributes set to an access value; and

25 for each of the identified PDEs, finding corresponding PTEs from the active PTE list that have update attributes set to an update value.

14. The method of claim 13 further comprising:

initializing the access attributes; and

initializing the update attributes.

30

15. A method comprising:

creating a shadow page table (PT) hierarchy based on a guest PT hierarchy used by a guest operating system for address translation operations; and

deriving metadata from the shadow PT hierarchy to determine subsequently which entries of the guest PT hierarchy that are represented in the shadow PT hierarchy were modified, the metadata comprising an active entry list identifying mappings that map pages used by the guest operating system in forming the guest PT hierarchy, and one or
5 more active directory lists identifying higher level mapping structures referencing a lower level structure through which the shadow PT hierarchy is to be accessed.

16. The method of claim 15 wherein:

the active entry list is an active PT entry (PTE) list identifying PTEs in the shadow
10 PT hierarchy, which map PT pages and page directory (PD) pages from the guest PT hierarchy, and

the one or more active directory lists include an active PD entry (PDE) list identifying PDEs in the shadow PT hierarchy, which point to PTs with the identified PTEs.

15

17. The method of claim 15 wherein:

the active entry list is an active PT entry (PTE) list identifying PTEs in the shadow PT hierarchy, which map any of page map level 4 (PML4) pages, page directory pointer (PDP) pages, page directory (PD) pages, and PT pages; and

20

the one or more active directory lists include an active PD entry (PDE) list containing PDEs that reference a page with active PTE list entries, an active PDP entry (PDPE) list containing active PDPE entries which reference a PD with an active PDE list entry, and an active PML4E entry (PML4E) list containing entries which map a PDP with elements from the active PDPE list.

25

18. The method of claim 16 further comprising:

tracking pages used by the guest operating system in forming the guest PT hierarchy; and

tracking mappings to any Dynamic Random Access Memory (DRAM) backed
30 page frame.

19. The method of claim 15 further comprising:

identifying one or more large-size pages in the guest PT hierarchy; and

creating a PT table in the shadow PT hierarchy for each of the one or more large-size pages.

20. The method of claim 15 further comprising:

5 detecting a change in a structure of the guest PT hierarchy; and
modifying the metadata to reflect the change.

21. The method of claim 15 wherein:

10 content of the shadow PT hierarchy is used by a processor to cache address
translations in the TLB.

22. An apparatus comprising:

a guest translation data structure used by a virtual machine (VM) for address
translation operations;

15 a shadow translation data structure maintained by a virtual machine monitor
(VMM), the shadow translation data structure deriving a format and structure from the
guest translation data structure; and

20 an address translation module to determine, based on metadata extracted from the
shadow translation data structure, which entries in the guest translation data structure were
modified by the VM, and to synchronize entries in the shadow translation data structure
that correspond to the modified entries in the guest translation data structure with the
modified entries in the guest translation data structure.

23. The apparatus of claim 22 wherein the address translation module is further to

25 maintain a working set of shadow translation data structures, each shadow translation data
structure in the working set being associated with one of a plurality of active processes of
the VM, and to reuse content of a shadow translation data structure from the working set
that is associated with one of the plurality of active processes when a guest translation data
structure associated with the one of the plurality of active processes is enabled.

30

24. The apparatus of claim 23 wherein the privileged event is any one of a request of
the VM to enable a different guest translation data structure, a page fault caused by one or
more inconsistencies between entries of the guest translation data structure and entries of

the shadow translation data structure, and a request of the VM to invalidate one or more address translations in a translation-lookaside buffer (TLB).

25. A system comprising:

5 a dynamic random access memory (DRAM) to store a guest translation data structure used by a virtual machine (VM) for address translation operations, and a shadow translation data structure deriving a format and structure from the guest translation data structure; and

10 a processor, coupled to the DRAM, to determine, based on metadata extracted from the shadow translation data structure, which entries in the guest translation data structure were modified by the VM, and to synchronize entries in the shadow translation data structure that correspond to the modified entries in the guest translation data structure with the modified entries in the guest translation data structure.

15 26. The system of claim 25 wherein the processor is further to maintain a working set of shadow translation data structures, each shadow translation data structure in the working set being associated with one of a plurality of active processes of the VM, and to reuse content of a shadow translation data structure from the working set that is associated with one of the plurality of active processes when a guest translation data structure
20 associated with the one of the plurality of active processes is enabled.

27. The system of claim 26 wherein the privileged event is any one of a request of the VM to enable a different guest translation data structure, a page fault caused by one or more inconsistencies between entries of the guest translation data structure and entries of
25 the shadow translation data structure, and a request of the VM to invalidate one or more address translations in a translation-lookaside buffer (TLB).

28. A machine-readable medium containing instructions which, when executed by a processing system, cause the processing system to perform a method, the method
30 comprising:

receiving control transitioned from a virtual machine (VM) due to a privileged event pertaining to a translation-lookaside buffer (TLB);

determining which entries in a guest translation data structure were modified by the VM, based on metadata extracted from a shadow translation data structure maintained by a virtual machine monitor (VMM) and attributes associated with entries in the shadow translation data structure; and

5 synchronizing entries in the shadow translation data structure that correspond to the modified entries in the guest translation data structure with the modified entries in the guest translation data structure.

29. The machine-readable medium of claim 28 wherein the method further comprises:

10 maintaining a working set of shadow translation data structures, each shadow translation data structure in the working set being associated with one of a plurality of active processes of the VM, and

 reusing content of a shadow translation data structure from the working set that is associated with one of the plurality of active processes when a guest translation data
15 structure associated with the one of the plurality of active processes is enabled.

30. The machine-readable medium of claim 29 wherein the privileged event is any one of a request of the VM to enable a different guest translation data structure, a page fault caused by one or more inconsistencies between entries of the guest translation data
20 structure and entries of the shadow translation data structure, and a request of the VM to invalidate one or more address translations in a translation-lookaside buffer (TLB).


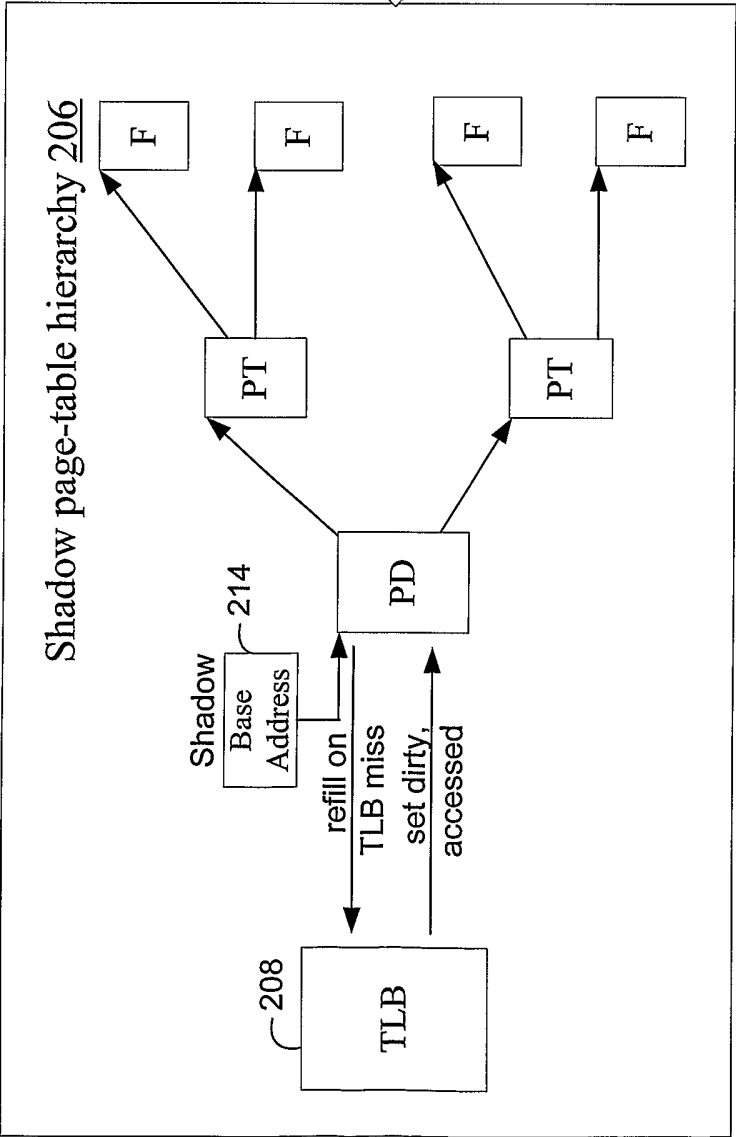


FIG. 1

200

Guest page-table hierarchy 202



"Virtual TLB" 204

FIG. 2





FIG. 3A



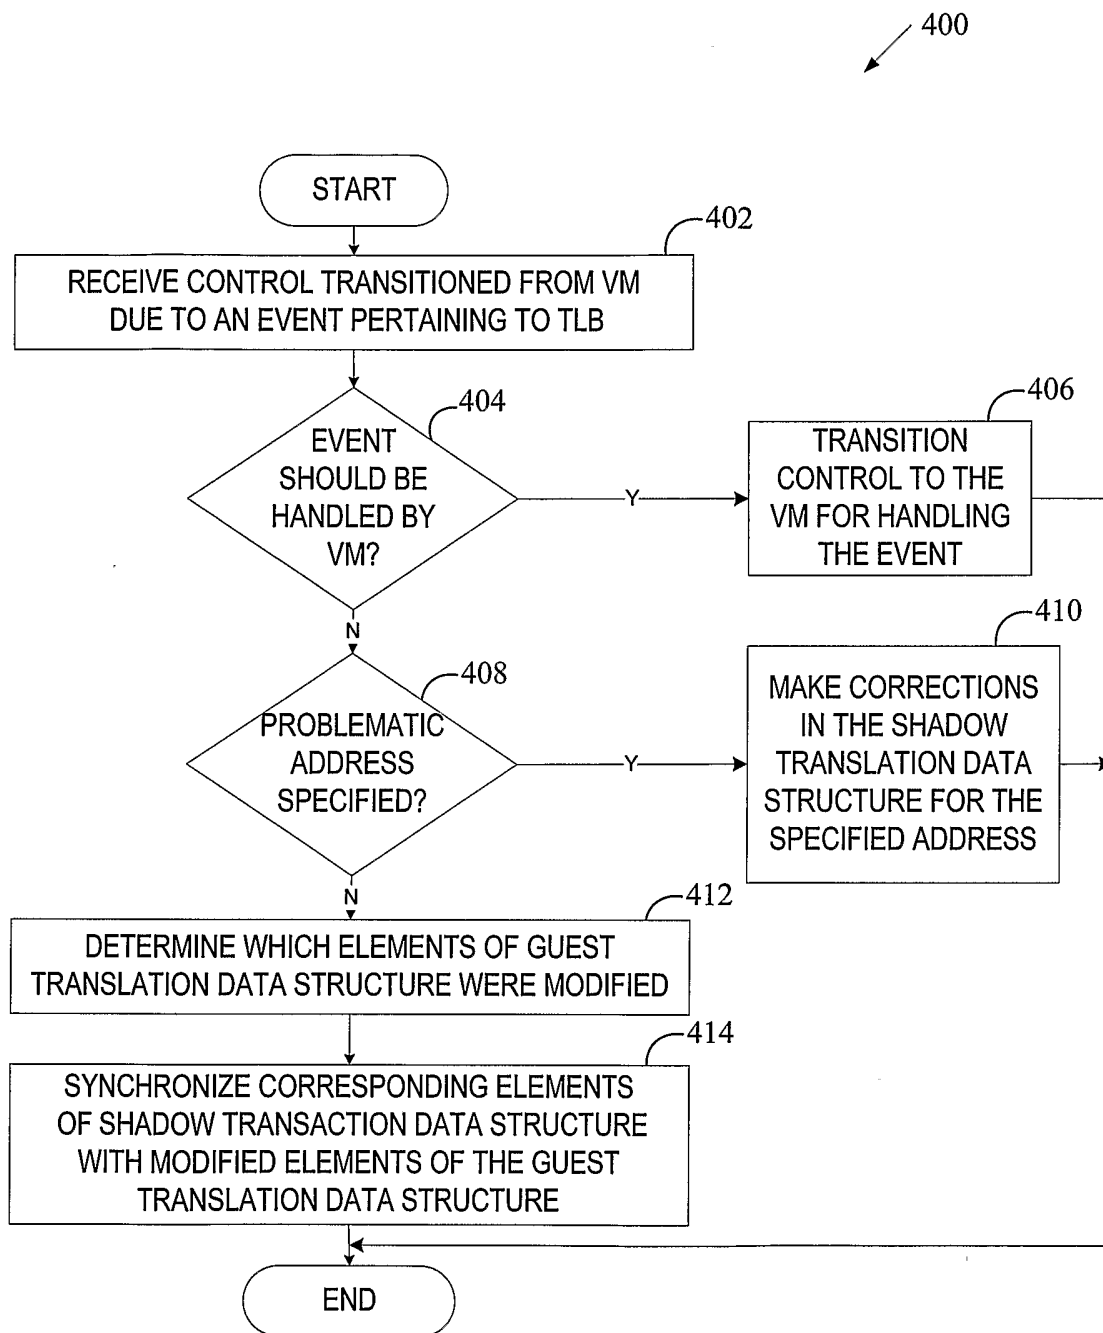


FIG. 4

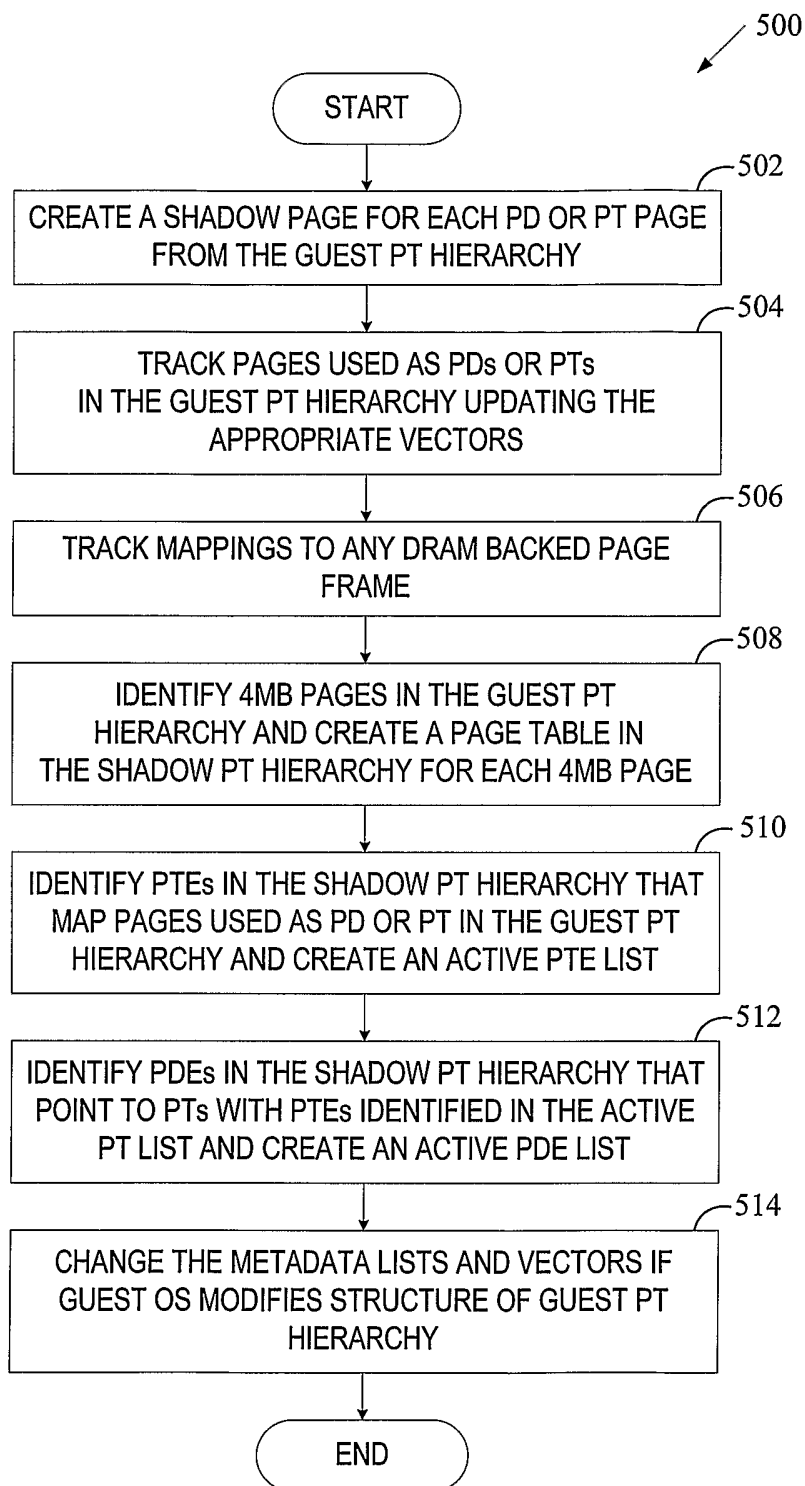


FIG. 5

7/15

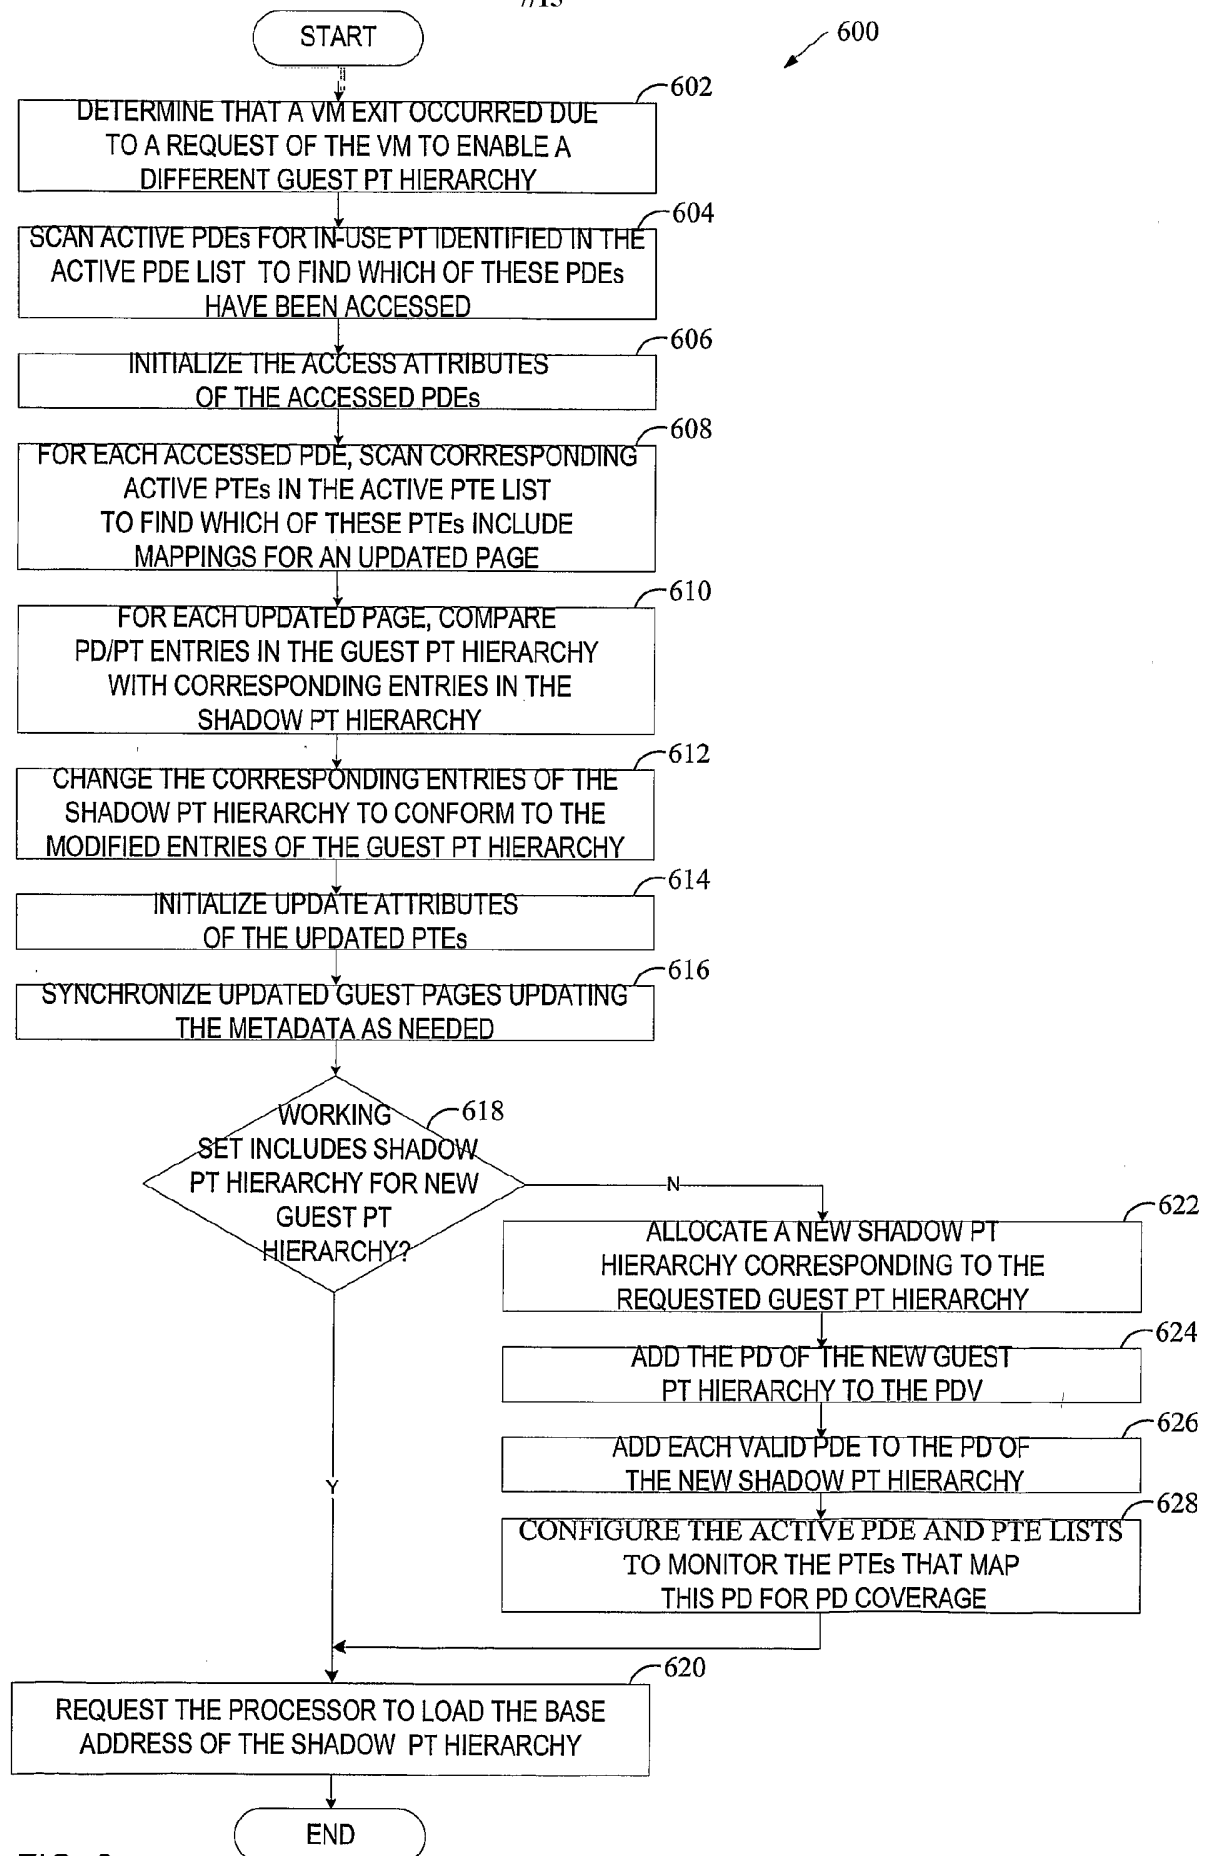


FIG. 6

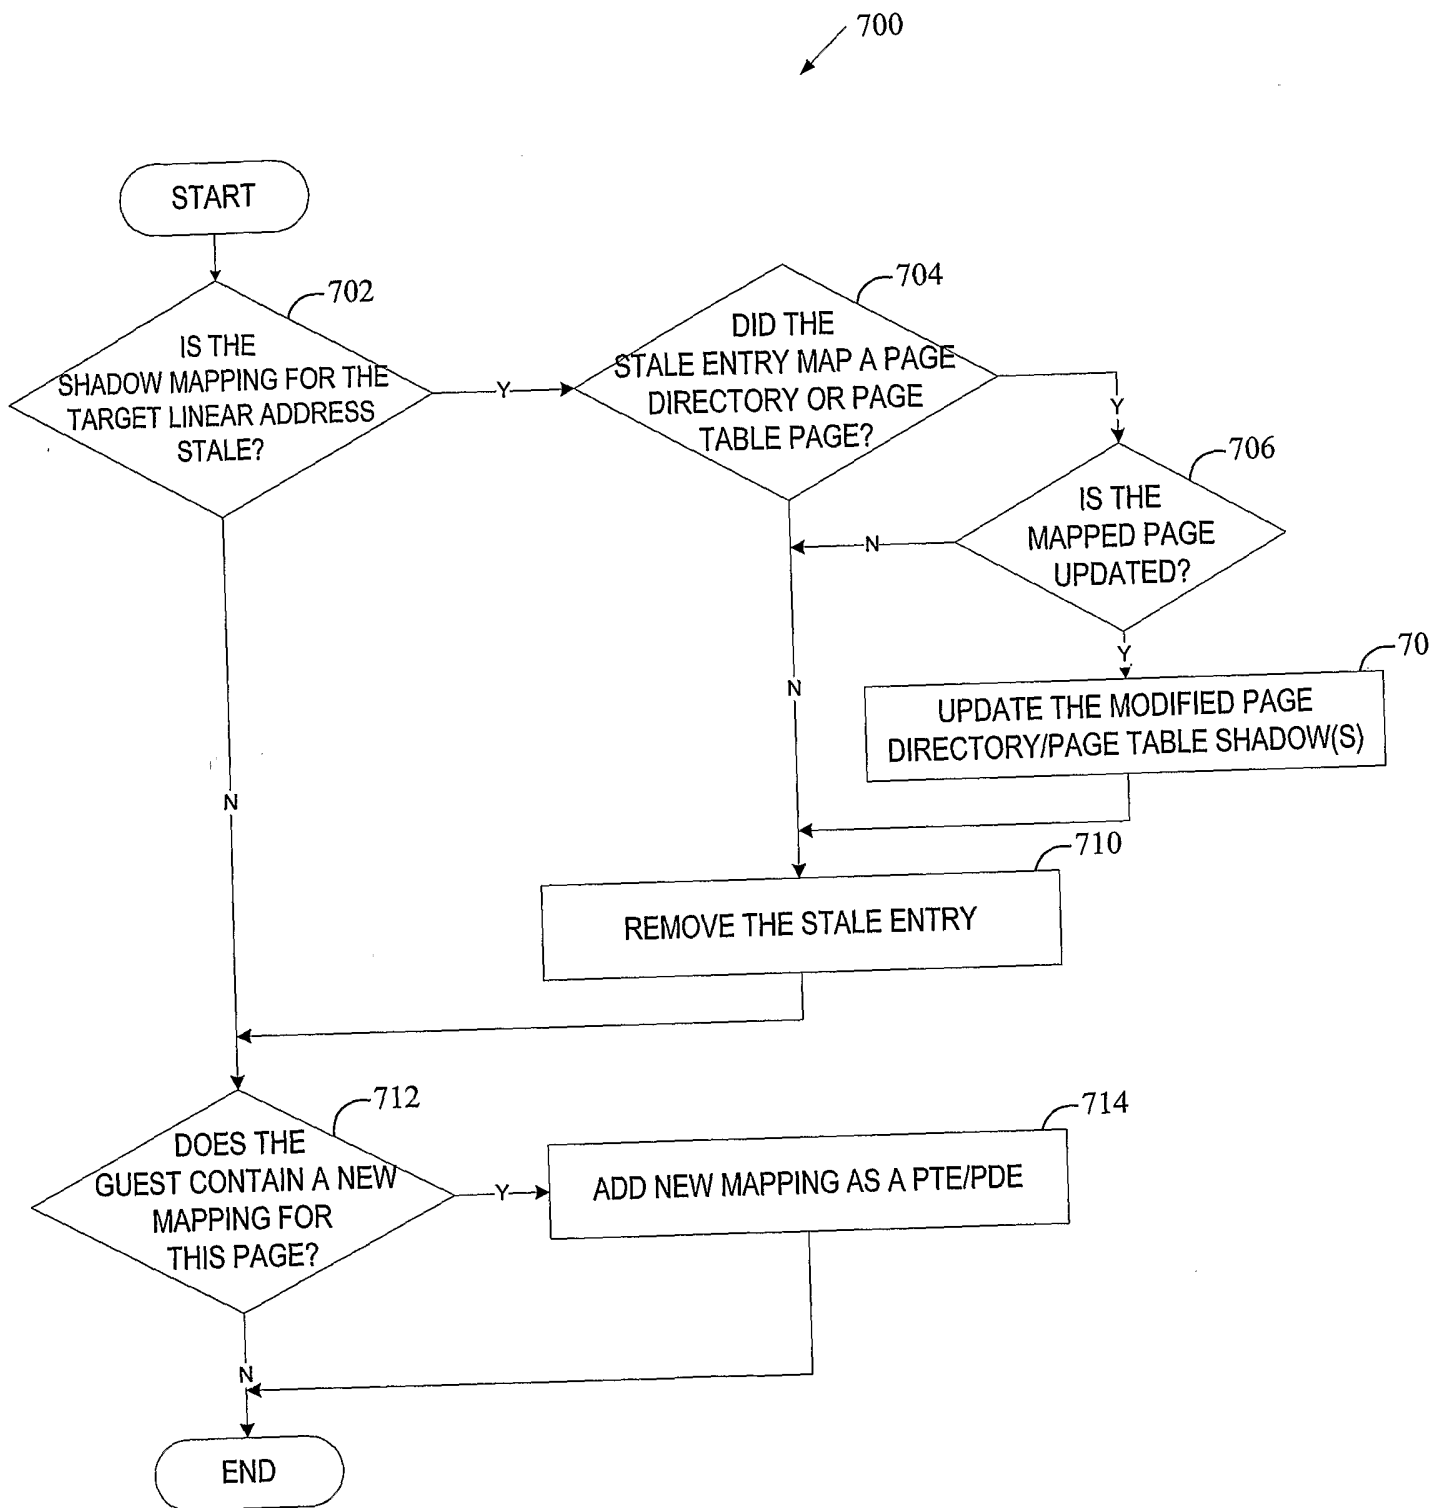


FIG. 7




FIG. 8

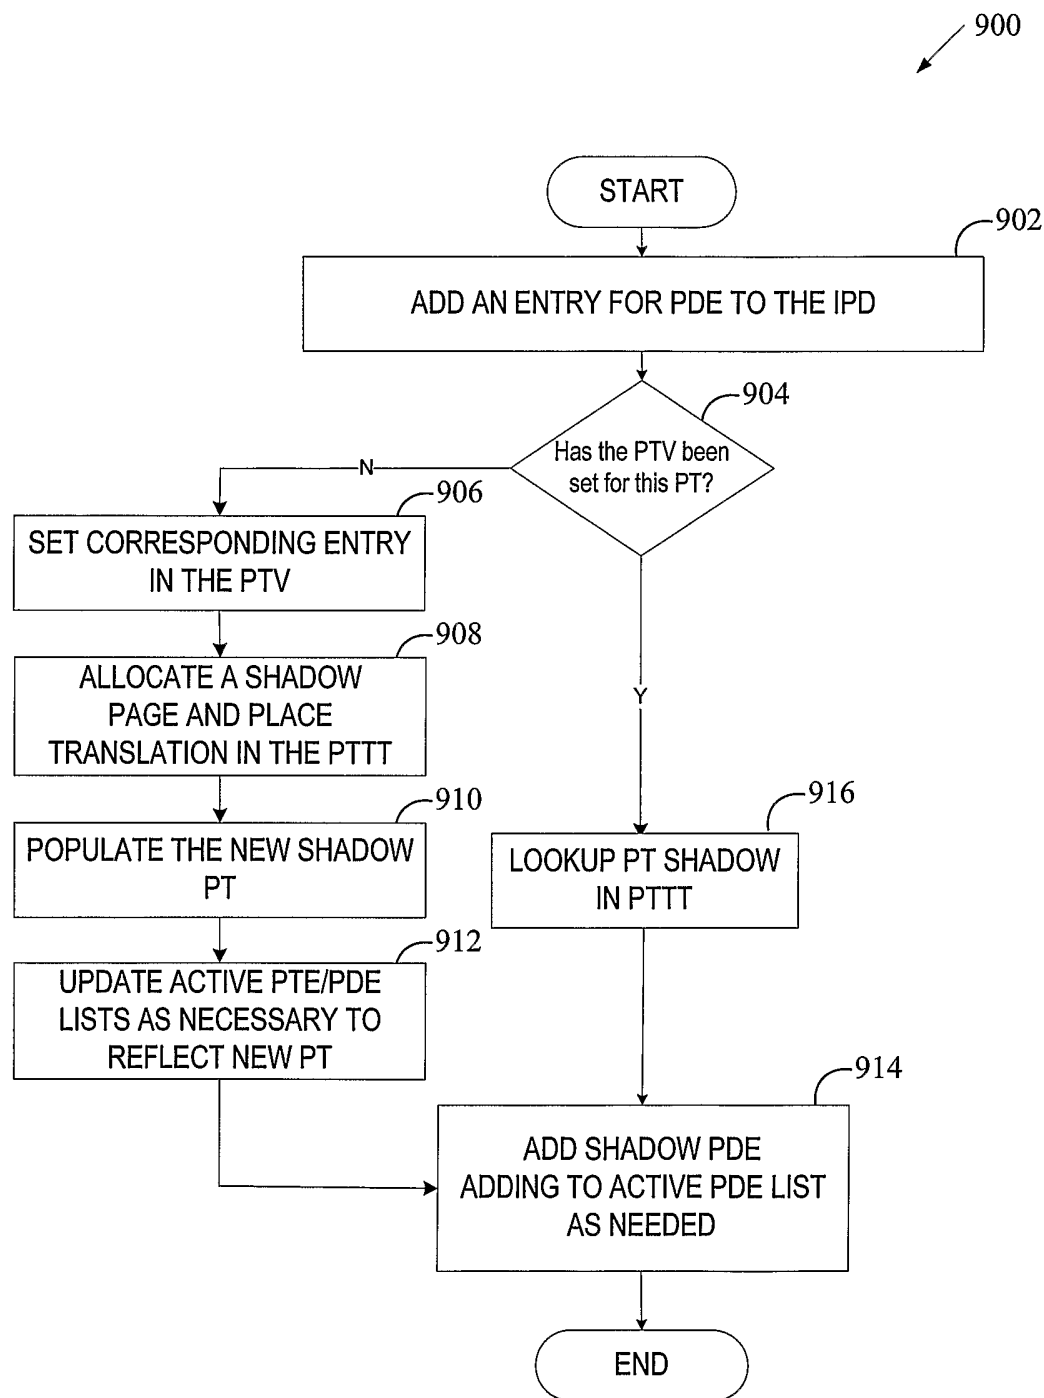


FIG. 9

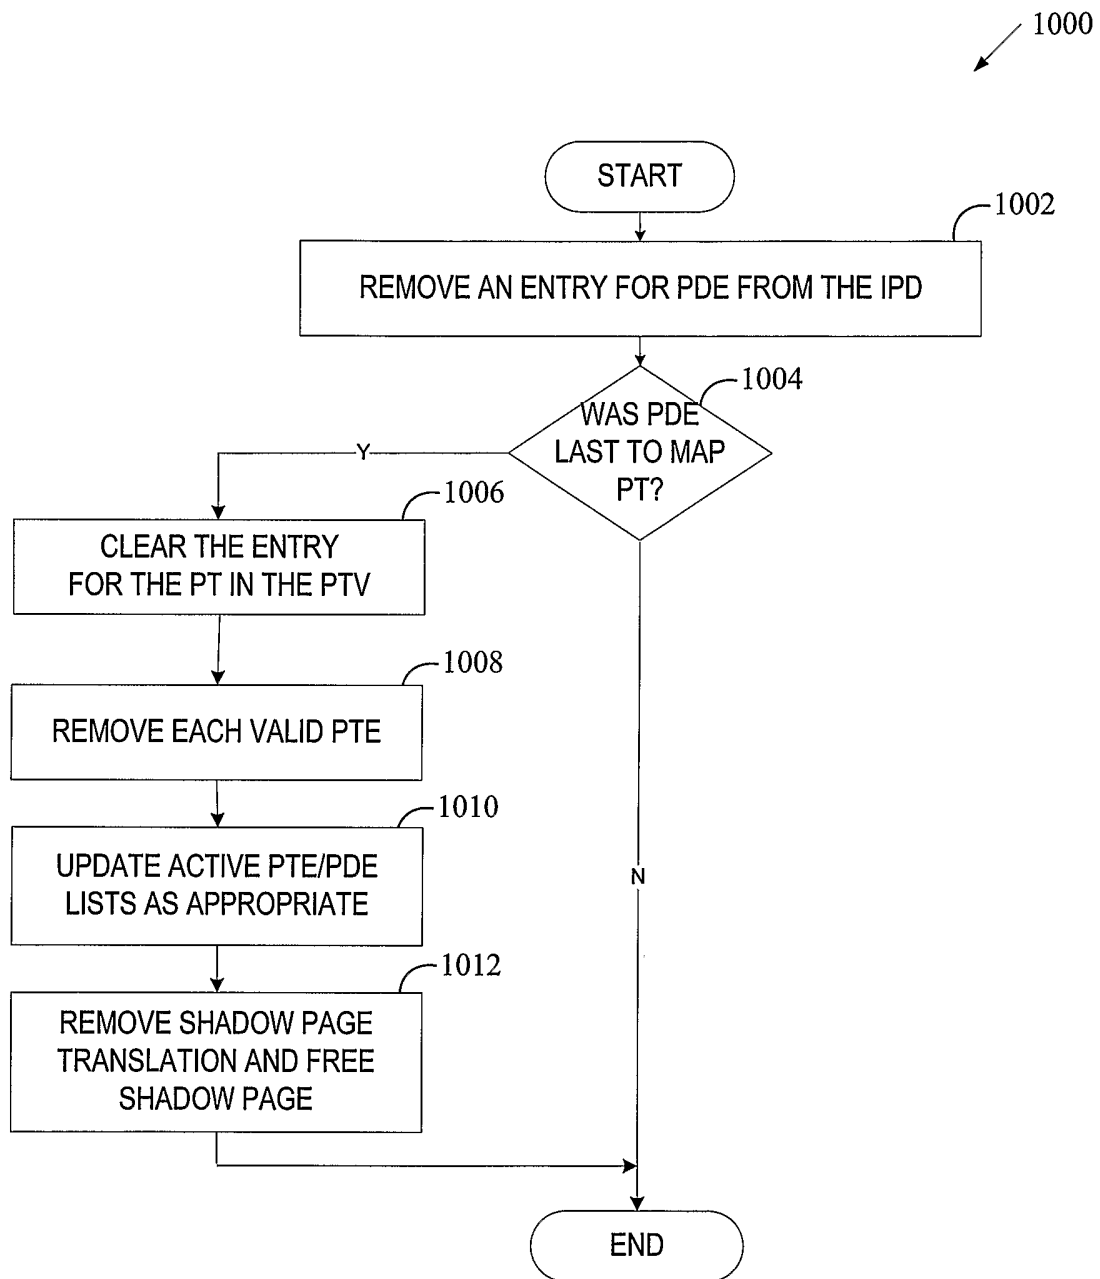


FIG. 10

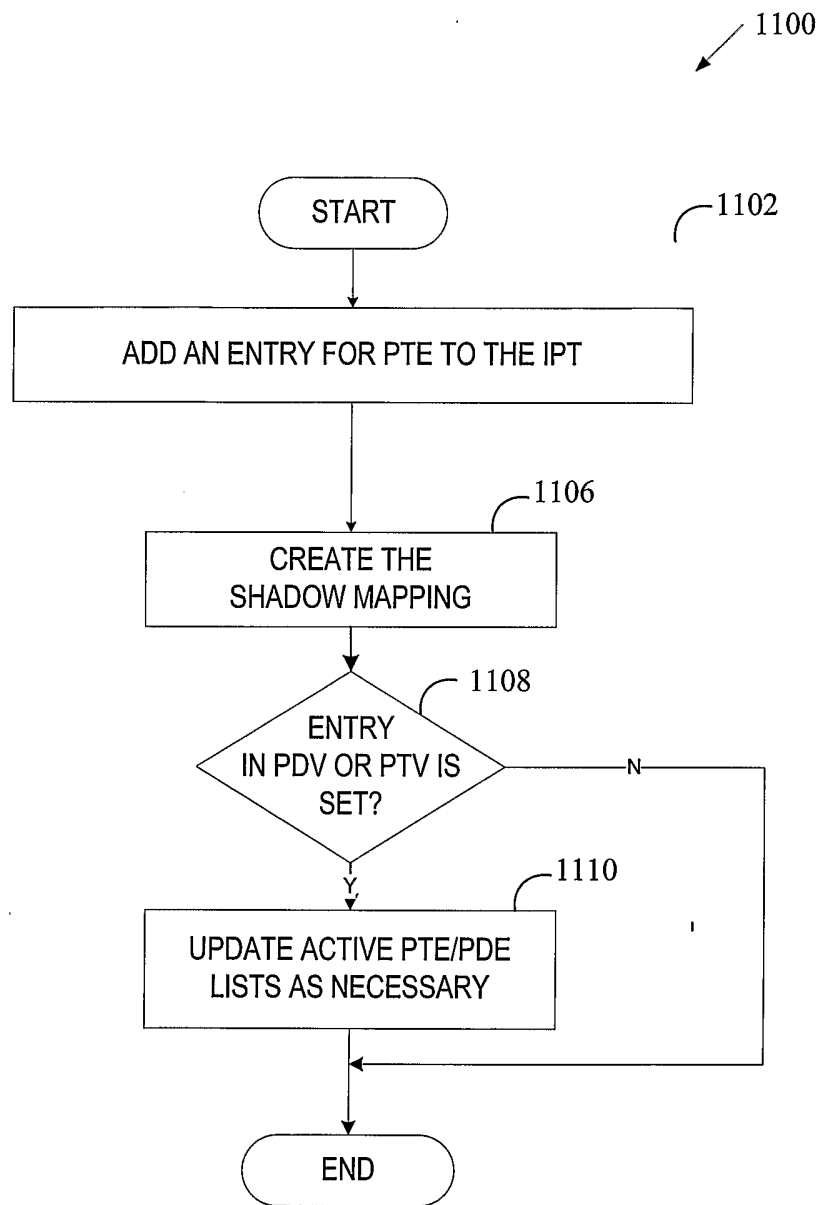


FIG. 11

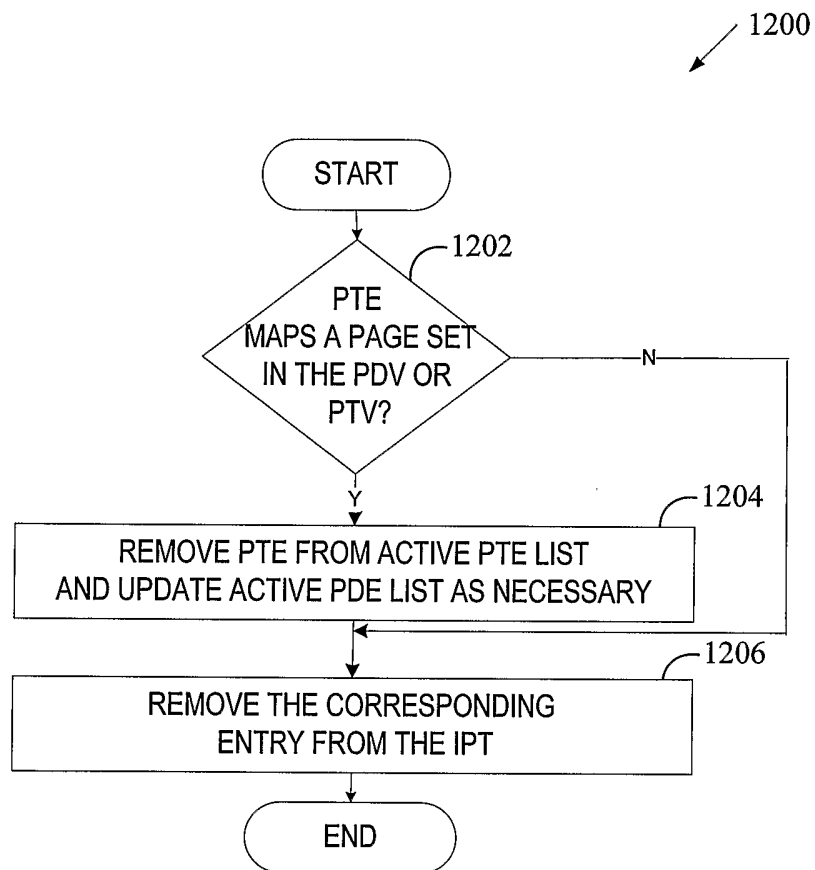


FIG. 12

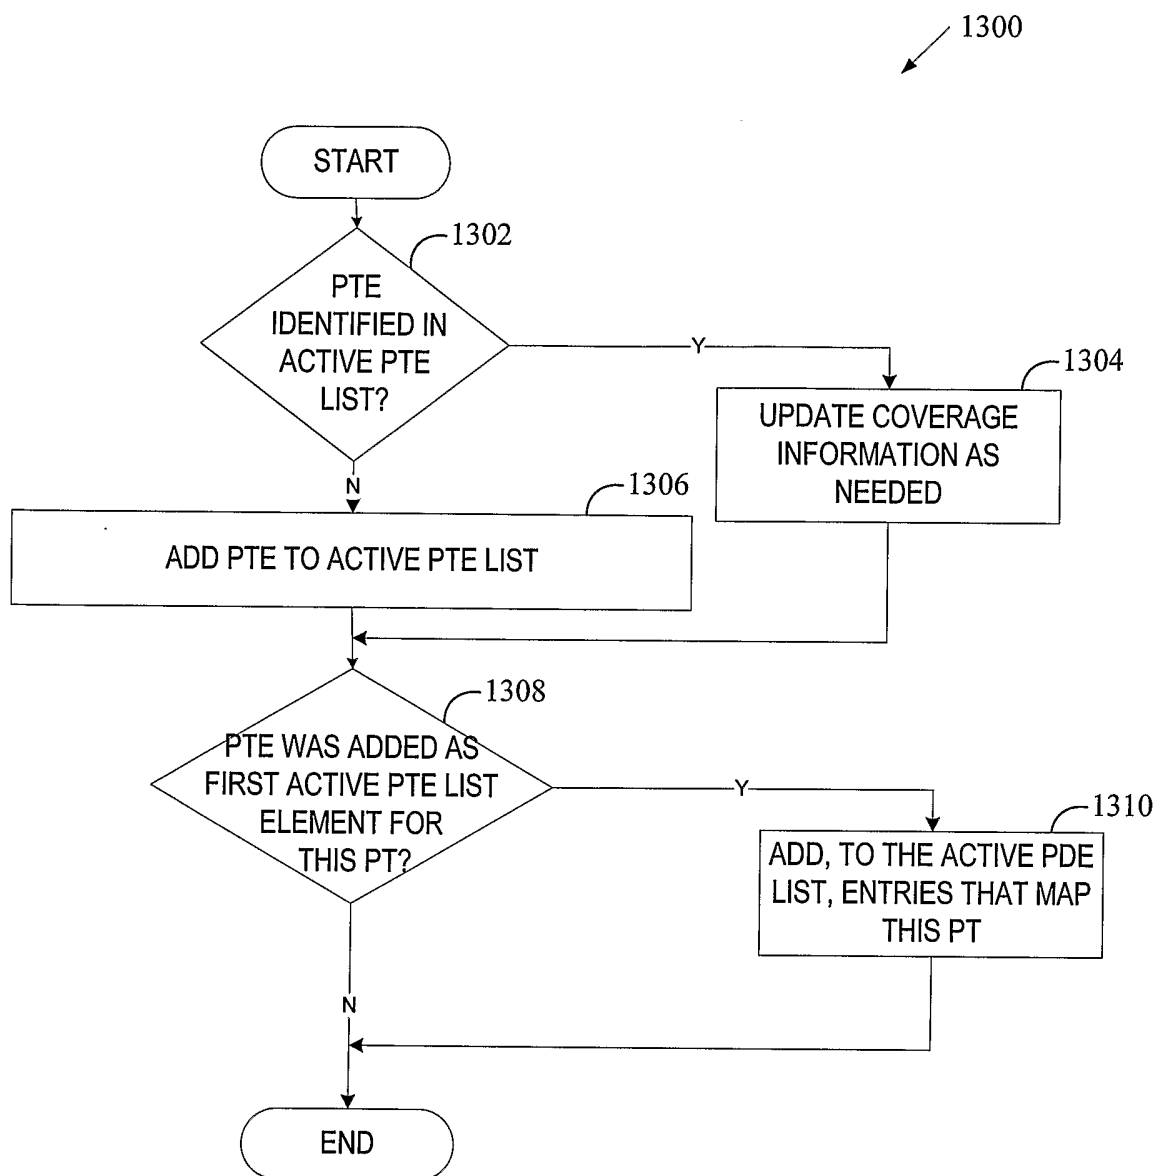


FIG. 13

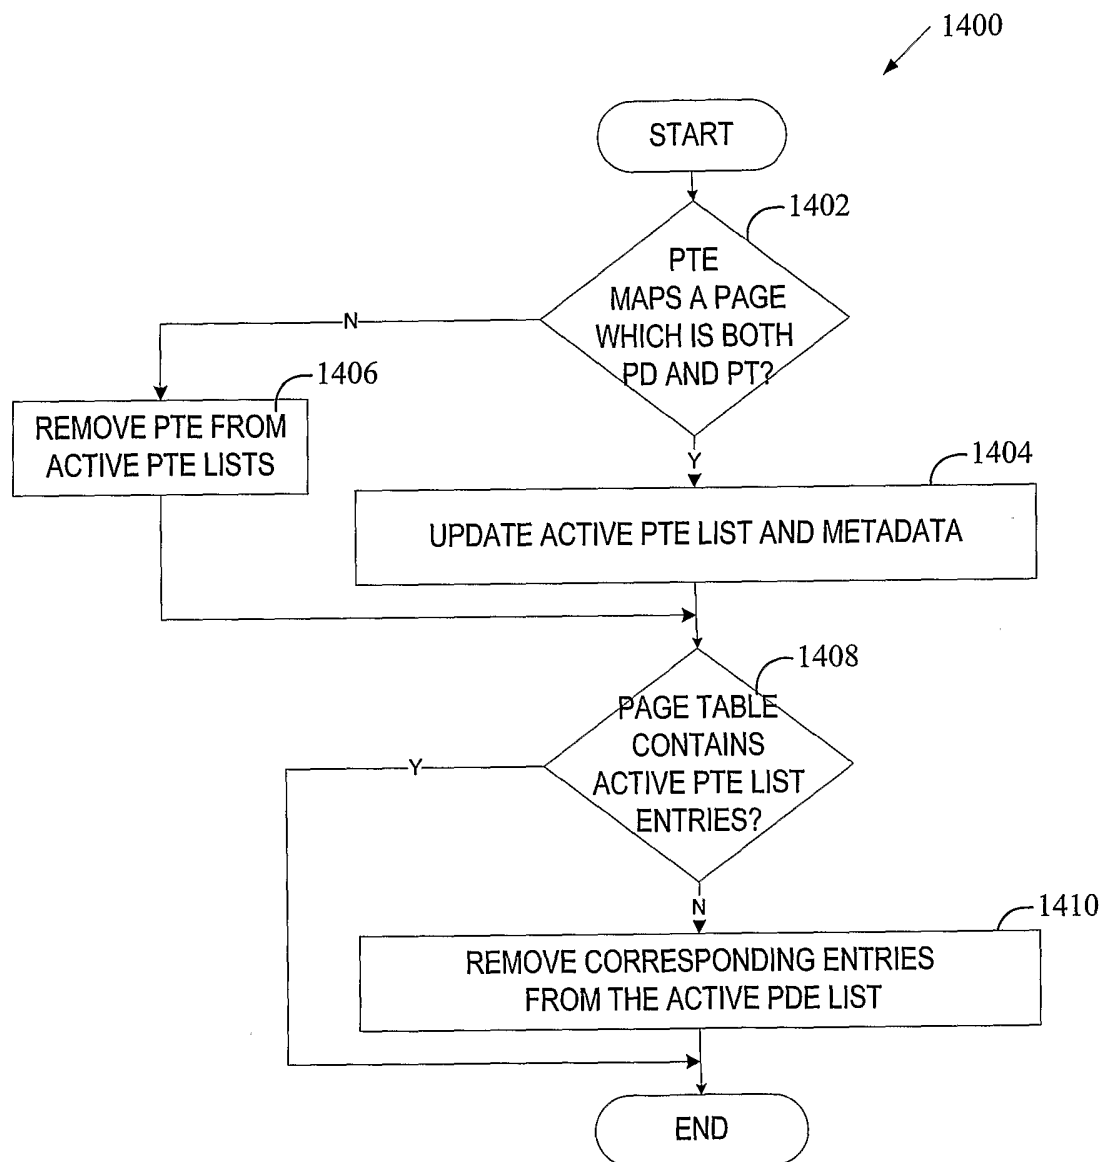


FIG. 14