

(12) **United States Patent**
Devine et al.

(10) **Patent No.:** **US 6,397,242 B1**
 (45) **Date of Patent:** **May 28, 2002**

(54) **VIRTUALIZATION SYSTEM INCLUDING A VIRTUAL MACHINE MONITOR FOR A COMPUTER WITH A SEGMENTED ARCHITECTURE**

5,761,477 A * 6/1998 Wahbe et al. 709/1
 5,832,205 A 11/1998 Kelly et al.

OTHER PUBLICATIONS

Goldberg, "Survey of Virtual Machine Research," Computer, Jun. 1974, pp. 34-45.
 Ebciglu et al., "IBM Research Report—Daisy: Dynamic Compilation for 100% Architectural Compatibility", RC 20538, Aug. 5, 1996.
 Bugnion, "Disco: Running Commodity Operating Systems on Scalable Multiprocessors," ACM Trans. on Computer Systems, vol. 15, No. 4, Nov. 1997, pp. 412-447.
 Bressoud, "Hypervisor-based Fault-tolerance," SIGOPS '95, Dec. 1995, pp. 1-11.
 Rosenblum et al., "Using the SimOS Machine Simulator to Study Complex Computer Systems," ACM Trans. on Modeling and Computer Simulation, vol 7, No. 1, Jan. 1997, pp. 78-103.
 Creasy, "The Origin of the VM/370 Time-Sharing System," IBM J. Res. Develop., vol. 25, No. 5, Sep. 1981.
 Intel Architecture Software Developer's Manual, vol. 3, 1997.

(75) Inventors: **Scott W. Devine**, Palo Alto; **Edouard Bugnion**, Menlo Park; **Mendel Rosenblum**, Stanford, all of CA (US)

(73) Assignee: **VMWare, Inc.**, Palo Alto, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/179,137**

(22) Filed: **Oct. 26, 1998**

Related U.S. Application Data

(60) Provisional application No. 60/085,685, filed on May 15, 1998.

(51) **Int. Cl.**⁷ **G06F 9/00**

(52) **U.S. Cl.** **709/1; 709/214; 709/321; 703/27; 710/23; 711/148; 711/153**

(58) **Field of Search** **709/100, 1, 200, 709/224, 316, 320, 328, 330, 223; 717/131, 138, 140; 703/26, 27; 714/1, 2, 47**

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,747,040 A	5/1988	Blanset et al.	
4,787,031 A	11/1988	Karger et al.	
4,792,895 A	12/1988	Tallman	
4,926,322 A	5/1990	Stimac et al.	
4,974,159 A	11/1990	Hargrove et al.	
5,134,580 A	7/1992	Bertram et al.	
5,167,023 A	11/1992	de Nicolas et al.	
5,255,379 A	10/1993	Melo	
5,307,504 A	4/1994	Robinson et al.	
5,440,710 A	8/1995	Richter et al.	
5,488,716 A	1/1996	Schneider et al.	
5,522,075 A	5/1996	Robinson et al.	
5,560,013 A *	9/1996	Scalzi et al.	717/138
5,652,869 A	7/1997	Herdeg et al.	
5,652,872 A	7/1997	Richter et al.	
5,721,922 A	2/1998	Dingwall	

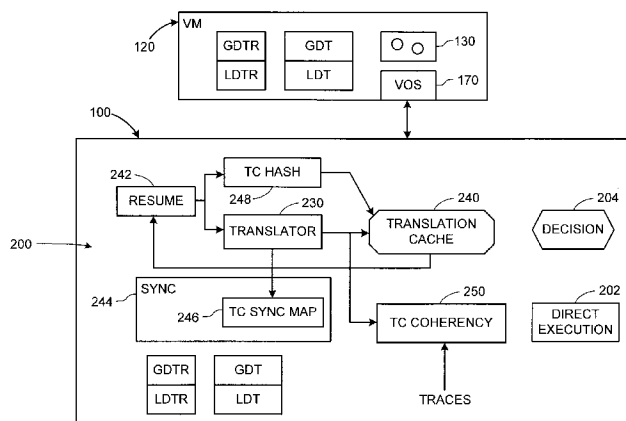
* cited by examiner

Primary Examiner—Majid Banankhah
 (74) *Attorney, Agent, or Firm*—Jeffrey Slusher

(57) **ABSTRACT**

In a computer that has hardware processor, and a memory, the invention provides a virtual machine monitor (VMM) and a virtual machine (VM) that has at least one virtual processor and is operatively connected to the VMM for running a sequence of VM instructions, which are either directly executable or non-directly executable. The VMM includes both a binary translation sub-system and a direct execution sub-system, as well as a sub-system that determines if VM instructions must be executed using binary translation, or if they can be executed using direct execution. Shadow descriptor tables in the VMM, corresponding to VM descriptor tables, segment tracking and memory tracing are used as factors in the decision of which execution mode to activate. The invention is particularly well-adapted for virtualizing computers in which the hardware processor has an Intel x86 architecture.

28 Claims, 6 Drawing Sheets



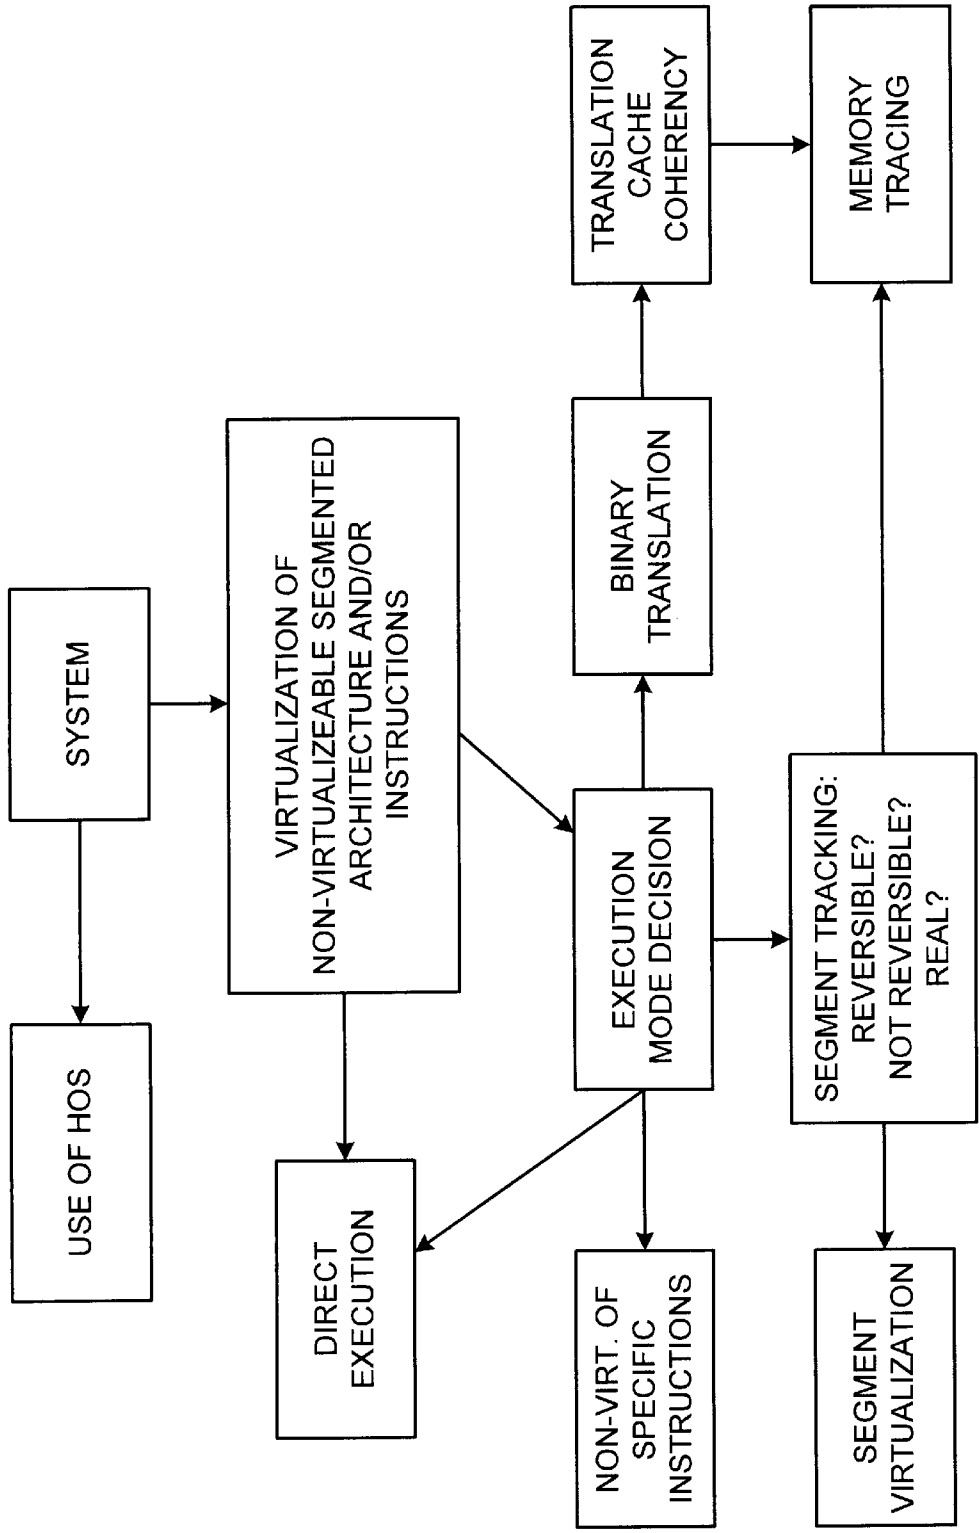


FIG. 1




FIG. 2




FIG. 3

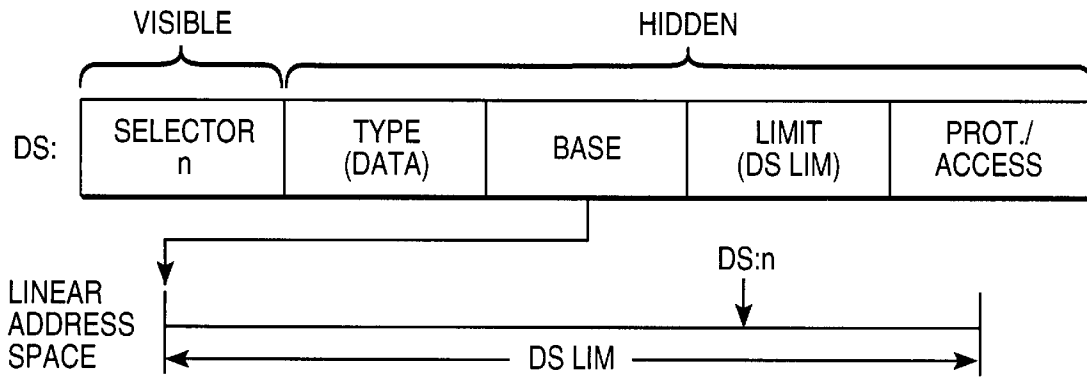


FIG. 4

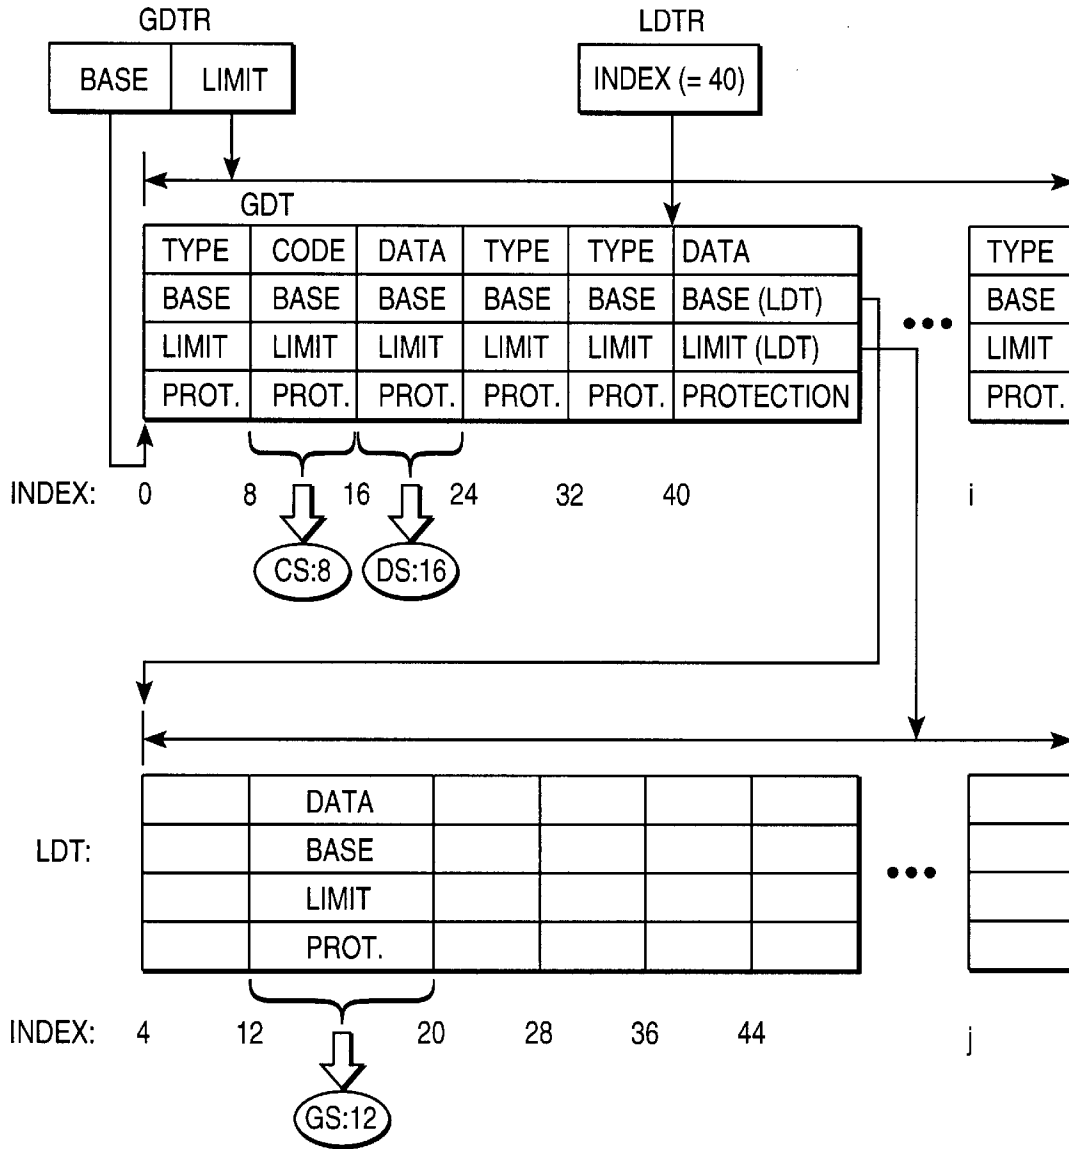


FIG. 5




FIG. 6

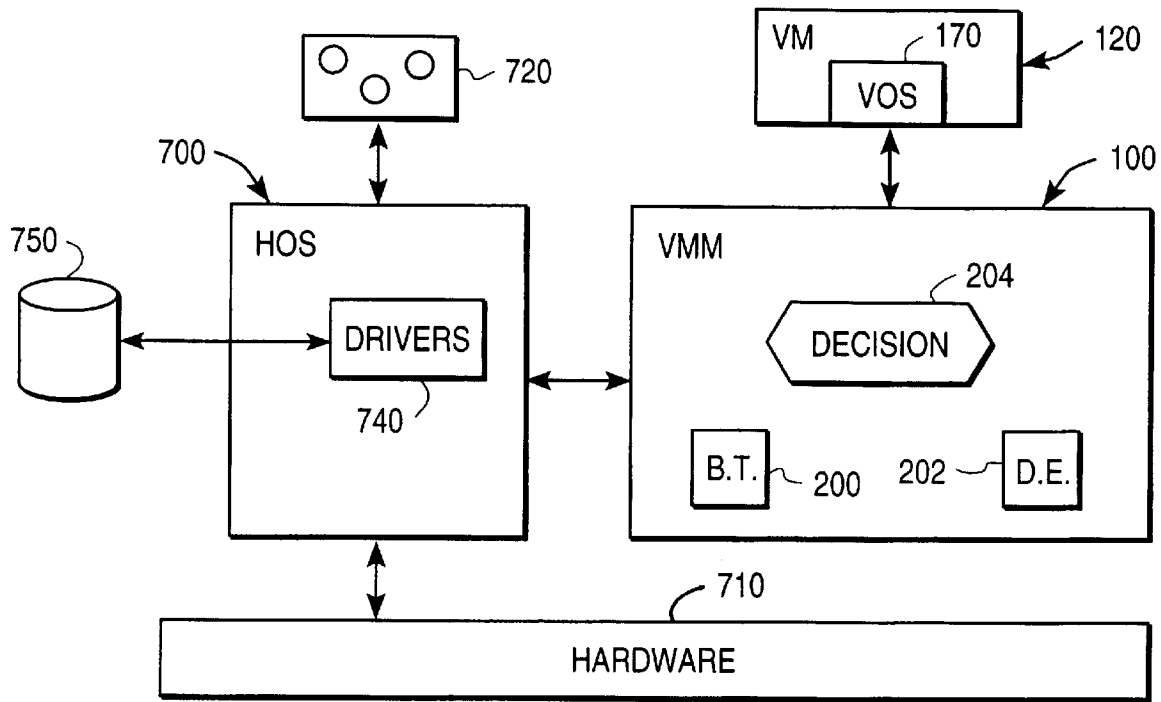


FIG. 7

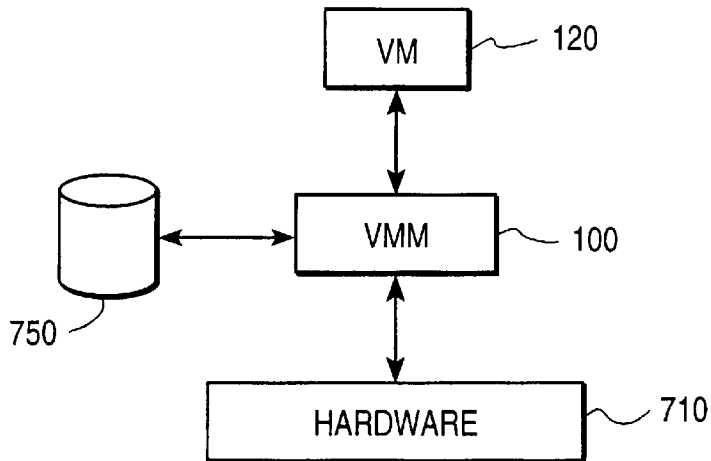


FIG. 8