

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
7 February 2008 (07.02.2008)

PCT

(10) International Publication Number
WO 2008/015685 A2

(51) International Patent Classification:
G06F 9/48 (2006.01) **G06F 9/52** (2006.01)

(21) International Application Number:
PCT/IL2007/000988

(22) International Filing Date: 7 August 2007 (07.08.2007)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/821,138 2 August 2006 (02.08.2006) US

(71) Applicant (for all designated States except US): **SanDisk IL Ltd.** [IL/IL]; 7 Atir Yeda St., 44425 Kfar Saba (IL).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **MOSEK, Amir** [IL/IL]; Shalom Alechem 39, 63344 Tel Aviv (IL).

(74) Agent: **FRIEDMAN, Mark**; 7 Jabotinsky St., 52520 Ramat Gan (IL).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,


AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, 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, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, 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, 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, MT, 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
- with information concerning request for restoration of the right of priority in respect of one or more priority claims

(54) Title: SYSTEMS AND METHODS FOR SYNCHRONOUS CODE RETRIEVAL FROM AN ASYNCHRONOUS SOURCE



(57) Abstract: The present invention discloses a computer-readable storage medium having computer-readable code embodied on the computer-readable storage medium, the computer-readable code including: program code for delivering, in response to a CPU request, from a host-system processor of a host system, for a command code, an SWI that is different than the command code. Preferably, the computer-readable code further includes: program code for selecting between providing the command code and providing the SWI based on an availability of the command code. Preferably, the computer-readable code further includes: program code for distinguishing between a command-code read-request and a data read-request according to a built-in command-codes table, wherein the command-code read-request and the data read-request are different. Preferably, the computer-readable code further includes: program code that uses SWI machine code to respond to the host-system processor, when the command code is a part of a missing-code segment, without causing latency.

WO 2008/015685 A2

SYSTEMS AND METHODS FOR SYNCHRONOUS CODE RETRIEVAL FROM AN ASYNCHRONOUS SOURCE

FIELD AND BACKGROUND OF THE INVENTION

5 The present invention relates to systems for code-delivery management capable of handling demand-paging code, independent of a host-system processor and operating system, and without latency.

Memory-management systems are well-known in the art of computer engineering as means for enhancing the performance of fast, random-access memory
10 (RAM) by dynamically swapping sub-areas of RAM with other larger (but slower) storage memory, so that the fast RAM can always be available for use in program execution. Such systems are controlled by a memory-management unit (MMU).

In the prior art, a typical memory-management system includes a host processor (having an internal or external MMU), non-volatile memory (NVM, which
15 is a large, but slow, memory), and RAM (which is typically fast, small volatile memory (VM) for code execution).

The MMU typically maps areas in the RAM to the NVM. The host-system processor loads code in the following ways:

- (a) upon application request, when one application intentionally attempts
20 to load and run another application that happens to be stored in the NVM (*e.g.* the dynamic-link library, or DLL, in MS WindowsTM);
- (b) upon page fault, when an unintentional “page fault” situation occurs, the MMU is responsible for loading the data from the NVM, and does so using the following sequence of operations:

(i) code caching: when possible, the MMU keeps the active code in a cache memory, eliminating the need for the host-system processor to read the code from an external resource (*e.g.* DRAM/RAM or any external storage device); this reduces the number of transactions (reducing power consumption), and improves processor performance and throughput;

(ii) mapping virtual addresses to physical addresses, using virtual-to-physical conversion tables enables software programs to implement a mechanism for demand paging, as explained below, and typically used by various operating systems (*e.g.* Windows Mobile, Symbian, and Linux);

(iii) a demand-paging mechanism is one of several implementations (enabled by the MMU) that makes the virtual-memory space equivalent to, or larger than, the physical-memory space;

(A) When a demand-paging mechanism is activated, and the host-system processor tries to read content (*e.g.* during code execution or data used by code) from the virtual-memory space, the MMU calculates the requested virtual page, and looks for the page in the internal cache;

(B) If the page is not in the cache, the MMU looks in the virtual-to-physical table for the physical address in the RAM where the code or data is located;

(C) If the code or data is not mapped to a RAM physical address (and therefore cannot be retrieved), then a page fault occurs; and

(D) In response to the page fault, the MMU invokes an exception-interrupt to the host-system processor, which is expected to provide the missing page from the NVM.

The main problem with the prior-art approach described above is the complexity of the process. Such a process becomes difficult to handle due to the plurality of management systems, drivers, and protocols that are involved. In the prior art, Microsoft Windows CE (4.2 and higher versions) supports demand paging. To support the demand-paging mechanism, a developer needs to provide specialized integration for the operating system. During development, the developer needs to install a file-system driver and a block device driver, and format the storage memory in a unique way (*i.e.* not similar to any integration required for storage memory that does not store code for demand paging).

A memory-management system in which the host-system processor requests data from one location, and always gets the data from that location (either directly or indirectly), is needed. However, the host-system processor demands to get the data synchronously, and within a very short response time (*e.g.* typically 4-5 clock cycles). Since some of the system components (*e.g.* the NVM) cannot work synchronously and/or are unable to provide data within the required response time, such a solution is difficult to provide. Such a scenario, together with the fact that the host-system processor cannot be externally interrupted during an operation, is why the cumbersome approach of the prior art appears to be the only possible solution.

It would be desirable to provide a memory-management system that is able to synchronously serve the demand of a host-system processor for code, providing the code wherever the code is stored (*i.e.* in the cache, in the RAM or in the NVM).

SUMMARY OF THE INVENTION

It is the purpose of the present invention to provide systems for code-delivery management capable of handling demand-paging code, independent of a host-system processor and operating system, and without latency.

5 For the purpose of clarity, several terms which follow are specifically defined for use herein. The terms “program counter” and “PC” are used herein to refer to an internal register in a host-system processor that includes the address of the current command to be executed by the host-system processor. The PC is incremented automatically by the host-system processor when the processor completes command
10 execution. The PC can be set to any address by the host-system processor using special “branch” or “jump” commands.

The terms “software interrupt” and “SWI” are used herein to refer to a processor command that automatically sets the PC to a “reset-vector” address. Upon jumping to the reset-vector address, the host-system processor immediately starts
15 executing code located at that address. The term “software-interrupt handler” is used herein to refer to code that is located in the reset-vector address, and is executed upon software-interrupt execution by the host-system processor.

The term “processor cache” is used herein to refer to a portion of memory that has the fastest performance, and typically a limited volume, in a computing system.
20 The term “page” is used herein to refer to the smallest unit of data used for memory management in a given system (typically 2-4 Kbytes). The term “segment” is used herein to refer to a set of sequential pages having the same type of usage (*e.g.* loading or storing method in VM and NVM).

The term “code” is used herein to refer to a set of processor commands that
25 can be executed by a processor. The term “resident code” is used herein to refer to

code that once loaded in VM, remains resident in the VM, and does need to be loaded from the NVM upon request. The execution of resident code does not involve latency.

The term "CPU atomic operation" is used herein to refer to a CPU operation that cannot be interrupted by any event (*e.g.* hardware or software interrupt), and thus, if started, will always be completed. The terms "code-delivery server" and "CDS" are used herein to refer to a hardware controller that is connected to a host-system processor by a main-memory bus. The CDS is connected to the main system memory by a system memory bus, or embedded inside main system memory. Every access request to the main system memory goes through the CDS before the request reaches the main system memory. In the present invention, the CDS is responsible for delivering one of the following in response to a CPU request for a given command code: (a) a command code from main system memory, or (b) a SWI that is different from the requested command code.

The term "code segment table" is used herein to refer to a table, maintained in the host-system processor, in which each code segment is represented by a start address or segment number and length of the command code.

In a preferred embodiment of the present invention, a memory-management system is taught that satisfies synchronous CPU requests for command code, providing the code from a variety of sources (some of which are asynchronous).

In another preferred embodiment of the present invention, a memory-management mechanism is taught that provides a host-system processor with a synchronous response for any code read-request, whether the requested code is available in VM or not. Every access request to VM is routed through the CDS. Upon receiving a code read-request from the host-system processor, the CDS forwards the requests to the main VM, receives the requested data from the main VM, and delivers

the data from the VM to the host-system processor. The host-system processor requires code (*i.e.* data read-requests) to be delivered within a few clock pulses (typically 3-5 pulses in DRAM). If the code is already located in VM, this timeframe is typically adequate. However, if the code is not resident in VM, and has to be loaded
5 from NVM, this timeframe is insufficient for the CDS to retrieve the code from NVM.

In a preferred embodiment of the present invention, in the case that requested code cannot be delivered on time by the CDS from VM (since the code might be stored in NVM or another slow storage device), the CDS avoids a situation of latency by providing the host-system processor with "replacement code" in the form of an
10 SWI. The host-system processor reads and executes the SWI instead of the actual code. Upon execution of the SWI, the host-system processor, in one CPU atomic operation, disables hardware interrupts, stores all the CPU registers and statuses, and jumps to a reset-vector address.

The reset vector includes code that causes the host-system processor to
15 execute program code in a finite loop ("wait-in-loop"). When the host-system processor completes the reset-vector program, the host-system processor restores the CPU registers and statuses, enables the hardware interrupts, and jumps back to the address that initiated the SWI (performed all within the time of one CPU atomic operation).

20 If the CDS has not managed to load the actual requested code from NVM to VM while the host-system processor was executing the reset-vector program, the CDS provides an SWI again, so that the host-system processor will jump again, as mentioned above, to the reset-vector wait-in-loop program. These processes are repeated until the CDS has completed the retrieval of the requested code from NVM.
25 When the CDS has completed the retrieval of the actual requested code, the CDS

waits for the next time the host-system processor requests the code again. Upon the next request, the code is delivered.

The results of such a procedure are that the host-system processor's need for fast response is fully satisfied, while the memory-management system has enough time to retrieve the data. While such a procedure may seem inefficient in terms of "visits per page", the procedure is very efficient in terms of programming, integration, and code maintenance.

In another preferred embodiment of the present invention, the CDS distinguishes between resident code and data. The CDS can distinguish between code and data according to parameters loaded during boot time from internal NVM (*e.g.* ROM or flash memory), loaded by registers that are programmed during production or development stages, or initialized by an application, executed by the host-system processor, at boot time (*i.e.* the application loads the parameters from NVM).

The CDS holds an internal virtual-to-physical table that maps between virtual-memory addresses of resident-code segments and physical-memory addresses of resident-code segments. Each entry in the table represents a segment number, and includes the physical address of the virtual-memory segment-number in the main system memory.

Upon an access request from the host-system processor to read or write data, the CDS carries out the following sequence of steps:

- (1) the CDS detects the access request, and determines whether the request is a read- or write-request;
- (2) if the request is not a read-request, the CDS directs the request to the main system memory in order to avoid any latency;

- (3) if the request is a read-request, the CDS checks if the requested data address belongs to a segment that may not be in VM, according to the *a priori* parameters mentioned above;
- (4) if the requested data address is in VM, then the sequence stops, and the request is forwarded to the main system memory for delivery;
- (5) if the requested data address is not in VM, then the CDS converts the virtual address to the physical address;
- (6) the CDS then sends a read request from the calculated physical address to the main system memory;
- (7) the CDS then checks the virtual-to-physical table to determine whether the requested data is already located in the main system memory ("data existing"), or should be loaded to the main system memory ("data missing");
- (8) if the requested data exists in VM, then the CDS stops the sequence of retrieving the data from NVM, and waits for the data to be retrieved from VM; and
- (9) if the requested data is missing from VM, the CDS proceeds with the following sub-sequence:
 - (a) the CDS responds to the host-system processor with data that includes an SWI whenever the requested code has not yet been retrieved and is ready for delivery;
 - (b) the PC of the host-system processor points to the SWI, reads the SWI and executes the SWI;
 - (c) since the command that the host-system processor retrieved from the main system memory is an SWI, the host-system

processor disables the interrupt, stores all the registers, and then jumps automatically to the software-interrupt handler in the reset vector (performed all within the time of one CPU atomic operation);

- 5 (d) the CDS identifies the address that caused the SWI, and calculates the virtual-memory segment-number to be loaded accordingly;
- (e) the CDS then loads the data of the calculated virtual-memory segment-number from NVM to the system main memory; and
- 10 (f) once the data is loaded into the system main memory, the CDS updates the virtual-to-physical table (*i.e.* updates the main-memory physical address of the “newly-loaded number, and marks the former page as “missing”).

Therefore, according to the present invention, there is provided for the first
15 time a computer-readable storage medium having computer-readable code embodied on the computer-readable storage medium, the computer-readable code including: (a) program code for delivering, in response to a CPU request, from a host-system processor of a host system, for a command code, an SWI that is different than the command code.

20 Preferably, the SWI is operative to cause the host-system processor to request the command code again.

Preferably, the command code resides in an NVM.

Preferably, the command code is retrieved from a remote source.

Preferably, the program code is embedded in a main VM component of the
25 host system.

Most preferably, the computer-readable code further includes: (b) program code for immediately directing every read-request, after converting a virtual address of the read-request to a physical address, to the main VM component irrespective of whether or not the command code is a part of an existing-code segment or a missing-code segment.

Preferably, the computer-readable code further includes: (b) program code for selecting between providing the command code and providing the SWI based on an availability of the command code.

Preferably, the computer-readable code further includes: (b) program code for recognizing resident code; and (c) program code for avoiding checking a VM component for the resident code upon receiving access requests for the resident code.

Preferably, the computer-readable code further includes: (b) program code for distinguishing between a command-code read-request and a data read-request according to a built-in command-codes table, wherein the command-code read-request and the data read-request are different.

Preferably, the computer-readable code further includes: (b) program code for identifying whether a requested code data-segment is available in a main VM component of the host system.

Preferably, the computer-readable code further includes: (b) program code that uses SWI machine code to respond to the host-system processor, when the command code is a part of a missing-code segment, without causing latency.

Preferably, the computer-readable code further includes: (b) program code for providing a synchronous response to the CPU request.

Preferably, the computer-readable code further includes: (b) program code for converting virtual addresses to physical addresses, and for storing the addresses in a conversion table; and (c) program code for command-code management.

Most preferably, the program code for command-code management is
5 operative to update the conversion table in accordance with an update command sent by the host-system processor.

These and further embodiments will be apparent from the detailed description and examples that follow.

10 BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is herein described, by way of example only, with reference to the accompanying drawing, wherein:

Figure 1 shows a simplified flowchart of the retrieval process, according to preferred embodiments of the present invention.

15

DESCRIPTION OF THE PREFERRED EMBODIMENTS

The present invention relates to systems for code-delivery management capable of handling demand-paging code, independent of a host-system processor and operating system, and without latency. The principles and operation for code-delivery
20 management capable of handling demand-paging code, according to the present invention, may be better understood with reference to the accompanying description and the drawing.

Referring now to the drawing, Figure 1 shows a simplified flowchart of the retrieval process, according to preferred embodiments of the present invention. Upon

receiving a request from the host-system processor to read or write data (Step 20), the CDS checks the request for the following (Step 22):

- (1) Is the request a read-request?
- (2) Does the data belong to a segment that is included in one of the “code-
5 segment” areas?
- (3) Is the segment not located in VM (*i.e.* “data missing”)?

If the answers to all three questions in Step 22 are positive, the CDS responds to the host-system processor with the SWI (Step 24). If the answer to one or more of the questions in Step 22 is negative, the CDS forwards the access request to the main
10 system memory (Step 26). It is noted, regarding the first two questions in Step 22, that the present invention mainly relates to read-requests and to code, as opposed to data.

An example of filler code that has a processor command that changes the PC value to a specific address value (*i.e.* the reset-vector address) is the code: swi 0x1 (a software interrupt with input parameter 1). The processor command in this example is
15 an SWI. The software-interrupt handler knows how to handle the SWI according to the input parameter. When the host-system processor executes the SWI, all other interrupts are disabled, all registers are stored automatically to a specific location, and the PC value is changed to the reset vector. Changing the PC by invoking the SWI, the PC stores all registers in a stack, and jumps automatically to the reset-vector
20 address.

A sample of reset-vector code (*i.e.* the software-interrupt handler), which is located at the reset-vector address, and executed upon receiving an SWI after all other interrupts are disabled, is shown below.

- (1) push r4; (push register)
- 25 (2) wait in loop of 1000;

13

	(3)	mov r4, #1000;	(init r4 (counter) to 1000)
	(4)	L1 subs r4, r4, #1;	(label 1: decrement r4)
	(5)	nop;	(wait)
	(6)	bne %B1;	(if r4 $\neq$ 0, jump to L1)
5	(5)	end of loop;	
	(6)	pop r4;	(pop register)
		(currently lr value is SWI+4 (performed by SWI))	
	(8)	subs lr, lr, #4;	(set lr value to the SWI address)
	(9)	ret;	(jump back to SWI address)
10			(enable interrupts)

When the host-system processor jumps to the reset-vector code, the interrupts are already disabled, and the registers of the host-system processor are already stored (including the “last-return” (LR) register which has the “jump-back” address for the host-system processor upon completion of executing the reset-vector code).

15 The reset-vector code instructs the host-system processor to loop for 1000 cycles, executing a nop command (*i.e.* “no operation” command). When the host-system processor completes executing the loop, the LR register is set to the address from where the SWI was called from. In the case that the host-system processor is an ARM processor, the ARM processor automatically keeps the “SWI address + 4” in

20 the LR register upon SWI execution. For calculating the SWI address, the host-system processor subtracts four from the value in the LR register. When the host-system processor executes the return command, interrupts are re-enabled, the registers are restored, the host-system processor returns to the SWI address (performed all within the time of one CPU atomic operation).

While the invention has been described with respect to a limited number of embodiments, it will be appreciated that many variations, modifications, and other applications of the invention may be made.

WHAT IS CLAIMED IS:

1. A computer-readable storage medium having computer-readable code embodied on the computer-readable storage medium, the computer-readable code comprising:
 - (a) program code for delivering, in response to a CPU request, from a host-system processor of a host system, for a command code, an SWI that is different than said command code.
2. The storage medium of claim 1, wherein said SWI is operative to cause said host-system processor to request said command code again.
3. The storage medium of claim 1, wherein said command code resides in an NVM.
4. The storage medium of claim 1, wherein said command code is retrieved from a remote source.
5. The storage medium of claim 1, wherein said program code is embedded in a main VM component of said host system.
6. The storage medium of claim 5, the computer-readable code further comprising:
 - (b) program code for immediately directing every read-request, after converting a virtual address of said read-request to a physical address,

to said main VM component irrespective of whether or not said command code is a part of an existing-code segment or a missing-code segment.

7. The storage medium of claim 1, the computer-readable code further comprising:

- (b) program code for selecting between providing said command code and providing said SWI based on an availability of said command code.

8. The storage medium of claim 1, the computer-readable code further comprising:

- (b) program code for recognizing resident code; and
- (c) program code for avoiding checking a VM component for said resident code upon receiving access requests for said resident code.

9. The storage medium of claim 1, the computer-readable code further comprising:

- (b) program code for distinguishing between a command-code read-request and a data read-request according to a built-in command-codes table, wherein said command-code read-request and said data read-request are different.

10. The storage medium of claim 1, the computer-readable code further comprising:

- (b) program code for identifying whether a requested code data-segment is available in a main VM component of said host system.

11. The storage medium of claim 1, the computer-readable code further comprising:

- (b) program code that uses SWI machine code to respond to said host-system processor, when said command code is a part of a missing-code segment, without causing latency.

12. The storage medium of claim 1, the computer-readable code further comprising:

- (b) program code for providing a synchronous response to said CPU request.

13. The storage medium of claim 1, the computer-readable code further comprising:

- (b) program code for converting virtual addresses to physical addresses, and for storing said addresses in a conversion table; and
- (c) program code for command-code management.

14. The storage medium of claim 13, wherein said program code for command-code management is operative to update said conversion table in accordance with an update command sent by said host-system processor.


15. A method for delivering code to a host system, the method comprising the steps of:

- (a) accepting a CPU request, from a host-system processor of the host system, for a code segment;
- (b) initiating a retrieval process to retrieve said code segment;
- (c) upon expiration of a predetermined time, checking whether said code segment is ready for delivery;
- (d) upon said predetermined time expiring before said code segment is ready for delivery, providing an SWI that is different than said code segment; and
- (e) upon said predetermined time expiring after said code segment is ready for delivery, providing said code segment.

16. The method of claim 15, wherein said SWI causes said host-system processor to jump to a reset-vector address.

17. The method of claim 16, wherein reset-vector code, located at said reset-vector address, includes a command to request said code segment again.

1/1

**Figure 1**