



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)


(51) International Patent Classification ⁵ : G06F 12/08, 12/10, 13/00	A1	(11) International Publication Number: WO 91/08537 (43) International Publication Date: 13 June 1991 (13.06.91)
---	----	---

(21) International Application Number: PCT/US90/06498
 (22) International Filing Date: 2 November 1990 (02.11.90)
 (30) Priority data:
 443,933 30 November 1989 (30.11.89) US
 (71) Applicant: STORAGE TECHNOLOGY CORPORATION
 [US/US]; 2270 South 88th Street, Louisville, CO 80028
 (US).
 (72) Inventor: KITCHEN, John, Frederick ; No. 1 Northcote
 Road, Lindfield, NSW 2070 (AU).
 (74) Agent: GRAZIANO, James, M.; Dorr, Carson, Sloan &
 Peterson, 3010 East 6th Avenue, Denver, CO 80206 (US).

(81) Designated States: AT (European patent), AU, BE (European patent), CA, CH (European patent), DE (European patent), DK (European patent), ES (European patent), FR (European patent), GB (European patent), GR (European patent), IT (European patent), JP, LU (European patent), NL (European patent), SE (European patent).

Published
With international search report.

(54) Title: DATA RECORD COPY APPARATUS FOR A VIRTUAL MEMORY SYSTEM



(57) Abstract

The data record copy apparatus for a virtual memory system (120) instantaneously makes a copy of a selected data record (Record 4) by simply generating a new set of pointers (201-2) to reference the same physical memory location (214) as the original reference pointer (200-n). In this fashion, by simply generating a second set of pointers (201-2) referencing the same physical memory space (214), this apparatus effectively makes a copy of the data record (Record 4). A second physical copy of the data record (Record 4) can later be created as a background process or can be delayed until a modification is made to the data record (Record 4) by one of the two programs that access the data record (Record 4).

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	FI	Finland	ML	Mali
AU	Australia	FR	France	MN	Mongolia
BB	Barbados	GA	Gabon	MR	Mauritania
BE	Belgium	GB	United Kingdom	MW	Malawi
BF	Burkina Faso	GN	Guinea	NL	Netherlands
BG	Bulgaria	GR	Greece	NO	Norway
BJ	Benin	HU	Hungary	PL	Poland
BR	Brazil	IT	Italy	RO	Romania
CA	Canada	JP	Japan	SD	Sudan
CF	Central African Republic	KP	Democratic People's Republic of Korea	SE	Sweden
CG	Congo	KR	Republic of Korea	SN	Senegal
CH	Switzerland	LI	Liechtenstein	SU	Soviet Union
CI	Côte d'Ivoire	LK	Sri Lanka	TD	Chad
CM	Cameroon	LU	Luxembourg	TG	Togo
DE	Germany	MC	Monaco	US	United States of America
DK	Denmark	MG	Madagascar		
ES	Spain				

**DATA RECORD COPY APPARATUS FOR A
VIRTUAL MEMORY SYSTEM**

FIELD OF THE INVENTION

5 This apparatus relates to virtual memory systems and, in particular, to a data record copy apparatus that instantly makes a copy of an existing data record without having to read the data record from memory.

PROBLEM

10 It is a problem in the field of memory systems to expeditiously make copies of data records stored therein. The typical data record copy process entails the processor locating a data record then retrieving the file from memory. Once the data record has been retrieved from memory and the integrity of the data
15 record contents verified, the processor must then write this data record to a new memory location and again verify the integrity of the data record contents. Upon completion of this operation, the memory system now contains two identical copies of the
20 original data record: (1) in the original memory location, (2) in an alternate location. It is obvious that such an operation consumes a significant amount of processor time. In the case where the memory system is a rotating media, such as a magnetic disk
25 drive system, disk head positioning and data channel availability become constraints that can have a

-2-

negative impact on the time it takes to retrieve the data record from memory.

5 In disk drive memory systems for example, a processor is typically connected by a channel interface unit to a plurality of data channels. Each data channel is connected to one or more direct access storage device (DASD) units which function to store large quantities of data. Each DASD unit includes a control unit for interfacing with the data channel
10 that typically carries eight bit parallel data in a byte serial decoded input/output record format. The control unit performs a data management function to maintain records of the physical memory location of all the data stored in the DASD unit. The DASD unit
15 includes rotating data storage media consisting of a plurality of disks, each of which has associated therewith one or more moveable read/write heads. The data is stored on the rotating media in a track format that consists of a plurality of concentric rings of data. A control module is also included at the head
20 of a string of disk drive units and connected to the control unit to convert between the eight bit parallel data format of the data channel and the bit serial data format of the rotating media.

25 In this disk drive memory system, a data record copy process requires the processor to perform a data record read operation. In a data record read operation, the processor transmits a seek request, through the channel interface unit and an available
30 data channel, to the DASD control unit associated with the rotating media on which the requested data record is stored. The processor returns to processing other tasks upon completion of the transmission of the seek request. The DASD control unit responds to the seek
35 request by determining the physical location of the

-3-

requested data record on the plurality of disks. The control unit transmits information to the associated control module identifying the physical location and size of the data record in order to retrieve the requested data record from one of the disks of the rotating media. Since the data record is stored in a track format on a rotating media, the actuator associated with the identified disk of the rotating media on which the requested data record is stored must wait until the rotating media rotates a sufficient distance to present the beginning of the requested data record underneath the read/write head associated with the actuator.

A predetermined time before the beginning of the requested data record reaches the read/write head associated with the actuator, the control module requests the control unit to obtain a data communication path to the processor in order to transmit the retrieved data record from the rotating media directly to the processor over this data communication path. If a data communication path to the processor is not available, the control module must wait for one entire revolution of the rotating media before the media is again in the position where the requested data record is a sufficient distance away from the read/write head to establish a data communication path to the processor. This process is repeated until a data communication path to the processor is available and the data, as read by the read/write head, can be transmitted directly via the control module and control unit and an available data channel to the processor. It is obvious that there can be numerous delays in this data storage system while data communication paths are established.

Once the data record is retrieved, the processor

-4-

- must write the data record to a new physical memory location to make a copy of the data record. This is accomplished by the processor establishing a data communication path via an available data channel to
- 5 the control unit associated with the physical memory location that is selected to receive the retrieved data record. The processor then transmits the retrieved data record to this control unit for storage on the rotating media.
- 10 This data record copy process is time consuming and requires a significant expenditure of processor resources. In systems that require a large number of data record copy operations, a significant amount of processor time is consumed in this data record
- 15 retrieval process.

-5-

SOLUTION

The above described problems are solved and a technical advance achieved in the field by the data record copy apparatus for a virtual memory system of the present invention. The data record copy apparatus
5 instantaneously makes a copy of a selected data record by simply generating a new set of pointers to reference the same physical memory location as the original set of reference pointers. In this fashion,
10 by simply generating a second set of pointers referencing the same physical memory space, this apparatus effectively makes a copy of the data record.

Two different programs can thereby each access identical copies of the data record since the data
15 record is a single physical entity. This apparatus instantaneously makes available a copy of the original data record without the time penalty of having to download the data record to the processor and write a copy of the data record to a new physical memory
20 location. For the purpose of enabling a program to simply access the data record at a different memory address the use of this mechanism provides a significant time advantage. A physical copy of the original data record can later be written as a
25 background process to a second memory location so that two physical copies of the data record are present on the memory system, if so desired. Alternatively, when one of the two programs that can simultaneously access the single shared data record via their respective
30 memory addresses write data to or modify the shared data record in any way, this modified copy of the original data record can be written to a new memory location and the corresponding address pointers changed to reflect the new location of this rewritten
35 data record.

-6-

In this fashion, a data record can be instantaneously copied by simply creating a duplicate set of memory pointers and the actual physical copying of the data record can take place either as a background process or as necessary when the data record is rewritten by one of the programs that accesses the data record. This provides a more efficient use of processor time and defers making an actual physical copy of the data record until absolutely necessary.

-7-

BRIEF DESCRIPTION OF THE DRAWING

Figure 1 illustrates in block diagram form the architecture of the subject computer system;

Figure 2 illustrates in schematic form a typical
5 virtual memory arrangement; and

Figure 3 illustrates in schematic diagram form the virtual memory arrangement of the present invention.

-8-

DETAILED DESCRIPTION

The data record copy apparatus for a virtual memory system instantaneously makes a copy of a selected data record by simply generating a new set of pointers to reference the same physical memory location as the original set of reference pointers. In this fashion, by simply generating a second set of pointers referencing the same physical memory space, this apparatus effectively makes a copy of the data record.

Two different programs can thereby each access identical copies of the data record since the data record is a single physical entity. This apparatus instantaneously makes available a copy of the original data record without the time penalty of having to download the file to the processor and write a copy of the data record to a new physical memory location. For the purpose of enabling a program to simply access the data record at a different memory address the use of this mechanism provides a significant time advantage. A physical copy of the original data record can later be written as a background process to a second memory location so that two physical copies of the data record are present on the memory system, if so desired. Alternatively, when one of the two programs that can simultaneously access the single shared data record via their respective memory addresses write data to or modify the shared data record in any way, this modified copy of the original data record can be written to a new memory location and the corresponding address pointers changed to reflect the new location of this rewritten data record.

In this fashion, a data record can be instantaneously copied by simply creating a duplicate

-9-

set of memory pointers and the actual physical copying of the data record can take place either as a background process or as necessary when the data record is rewritten by one of the programs that
5 accesses the data record. This provides a more efficient use of processor time and defers making an actual physical copy of the data record until absolutely necessary.

System Architecture

10 Figure 1 illustrates the overall architecture of a data processing system that is equipped with an associated data storage system that includes a plurality of memory devices, one of which is illustrated in Figure 1 in the form of a conventional
15 3380-type of disk drive unit 120. Processor 100 is connected via bus 101 to a channel interface unit 102 that serves to interconnect processor 100 with a plurality of data channels 103-1 to 103-n, each of which is an eight bit parallel data bus. Connected to
20 one of these data channels 103-1 is the disk drive unit 120 mentioned above.

Disk drive unit 120 includes a control unit 104 that functions to provide overall management of the disk drive unit 120. This management function
25 includes storing information that identifies the exact physical location of all data stored on the rotating media 111-0 to 111-15 of disk drive unit 120. Processor 100 identifies a data record by volume identification and address. While this information
30 provides a general physical location of the data, the exact storage location on the rotating media 111-0 to 111-15 requires additional specificity. Control unit 104 provides this additional specific information by converting the volume and address information received

-10-

from the processor into head, actuator, track and sector identification information to precisely define the physical location of the requested data record on disk drive unit 120. The control module 106 responds to the control signals transmitted by control unit 104 identifying the head, actuator, track and sector for storing a data record by selecting one of the actuators 110-0 to 110-15 that are used by disk drive unit 120 to read/write data on the rotating media 111-0 to 111-15. The above-described conventional disk drive unit 120 of Figure 1 is also equipped with a plurality of read/write head buffers 108-0 to 108-15 on a one per read/write head basis. Read/write head buffer 108-0 serves to temporarily store the track image captured by the read/write head 110-0.

Memory Pointers

Control unit 104 stores a plurality of pointers, each of which identifies the exact physical location on the ones of rotating media 111-0 to 111-15 that contain an identified data record. For each data record stored in disk drive unit 120, there is a corresponding pointer stored in control unit 104 that identifies the physical memory location that contains this data record. This is illustrated schematically on Figure 2 wherein elements 200-0 to 200-n are pointers that are stored in control unit 104. Elements 211, 212 represents physical memory locations on rotating media 111-0 while 213, 214 and 215 represent physical memory locations on rotating media 111-1. Only two of the rotating media are shown for the purpose of illustration and memory space 211-215 on these rotating media are illustrated schematically. As can be seen from Figure 2, pointer 200-1 contains data that is indicative of the physical location of

-11-

data record 1. This data points to physical memory location 212 on rotating media 111-0. Thus, when the processor 100 requests access to data record 1, control unit 104 scans the table of pointers to locate the entry corresponding to data record 1. This entry is addressed by control unit 104 and the data contained therein read out to obtain the identification of the physical memory location that contains data record 1. In this case, the pointer entry indicates that data record 1 is located on rotating media 111-0 in the physical memory space designated on Figure 2 as 212.

Similarly, pointer 200-4 contains the pointer entries for data record 2. The data stored in pointer location 200-4 indicates that data record 2 is stored on rotating media 111-1 in physical memory location 215.

Instant Copy Operation

The data record copy apparatus of the present invention can perform a "instant copy" operation wherein a copy of a particular data record is produced in a significantly less time than in prior systems. This apparatus takes advantage of the fact that the disk drive memory system 120 accessed by processor 100 is a virtual memory system in that the data record Record 2 resides at a physical memory location 215 that differs from its memory address 200-4. In particular, the processor 100 addresses a data record Record 2 according to a particular address 200-4, which address is translated by the disk drive system 120 into an actual physical memory location that does not correspond to the address used by processor 100. Control unit 104 maintains the information that indicates the correspondence between each particular

-12-

pointer data record memory address 200-4 and the actual physical memory location 215 that contains the data record Record 2.

5 The data record copy performance improvement is accomplished by processor 100 generating a new pointer 201-2, the contents of which point to the memory location 215 of the original data record Record 2. Thus, upon the conclusion of this data record copy operation, there is a single physical memory location 10 215 that contains the data record Record 2 and two pointers 200-4, 201-2, both of which point to that same physical memory location 215 that contains the data record Record 2. Since this is a virtual memory system, the two pointers 200-4, 201-2 represent two 15 separate data records even though there is only a single physical memory location 215 that contains the data record Record 2. Thus, instead of making a copy of the physical data record 215 itself, a second pointer 201-2 is simply generated to point to the same 20 data record Record 2. This operation makes two copies of the data record Record 2 available to the software resident on processor 100 without suffering the time penalty of having to actually read the data record Record 2 from the physical memory location 215 and 25 then write a new copy of the data record Record 2 to a separate physical memory location (ex 213).

Processor 100 can, as a background process, make another physical copy of the data record Record 2 if so desired. Alternatively, a more efficient method of 30 managing the shared data record Record 2 is to wait until one of the programs on processor 100 modifies the data record Record 2. At this point, the modified data record (ex Record 5) can be written to disk drive memory 120 to a new physical memory location 213. The 35 corresponding pointer 201-2 is adjusted to indicate

-13-

this new physical memory location 213 for the data record Record 5. In this fashion, another physical copy 213 of the data record Record 5 is produced only when changes are made to the original data record Record 2 and there no longer is a shared data record that has identical contents. This eliminates duplicate data records as is commonly found in prior art systems.

While a specific embodiment of this invention has been disclosed, it is expected that those skilled in the art can and will design alternate embodiments of this invention that fall within the scope of the appended claims.

-14-

I CLAIM:

1. In a virtual memory system that contains one or more data records stored therein and a corresponding number of data record pointers, each of which identifies the physical memory location of the associated data record, a method of copying a data record in response to a data record copy request comprising the steps of:

identifying the physical memory location of the requested data record as specified by its associated data record pointer; and

generating a data record pointer that is duplicative of said associated data record pointer as the data record pointer for said copy of said requested data record.

2. The method of claim 1 further including the step of:

storing said duplicate data record pointer as the data record pointer for said copy of said requested data record.

3. The method of claim 2 further comprising the steps of:

writing, in response to a modification of said data record, the modified version of said data record in a new memory location that differs from its original location; and

generating a data record pointer associated with said modified version of said data record to identify said new memory location.

4. In a virtual memory system that contains one or more data records stored therein and a

-15-

corresponding number of data record pointers, each of which identifies the physical memory location of the associated data record, apparatus for copying a data record stored in said virtual memory system comprising:

means responsive to a request to copy a designated data record for retrieving the data record pointer associated with said designated data record;
means duplicating said retrieved data record pointer; and
means for storing said duplicate copy of said data record pointer as the data record pointer for said copy of said designated data record.

5. The apparatus of claim 4 further comprising:
means for writing, in response to a modification of said data record, the modified version of said data record in a new memory location that differs from its original memory location; and
means for generating a data record pointer associated with said modified version of said data record to identify said new memory location.

6. In a virtual memory system that contains one or more data records stored therein and a corresponding number of data record pointers, each of which identifies the physical memory location of the associated data record, a method of copying a data record in response to a data record copy request comprising the steps of:
identifying the physical memory location of the requested data record as specified by its associated data record pointer;
generating a data record pointer that is duplicative of said associated data record pointer as

-16-

the data record pointer for said copy of said requested data record;

15 storing said duplicate data record pointer as the data record pointer for said copy of said requested data record;

 writing, in response to a modification of said data record, the modified version of said data
20 record in a new memory location that differs from its original location; and

 generating a data record pointer associated with said modified version of said data record to identify said new memory location.

7. In a virtual memory system that contains one or more data records stored therein and a corresponding number of data record pointers, each of which identifies the physical memory location of the
5 associated data record, apparatus for copying a data record stored in said virtual memory system comprising:

 means responsive to a request to copy a designated data record for retrieving the data record
10 pointer associated with said designated data record;

 means for duplicating said retrieved data record pointer;

 means for storing said duplicate copy of said data record pointer as the data record pointer
15 for said copy of said designated data record;

 means for writing, in response to a modification of said data record, the modified version of said data record in a new memory location that differs from its original location; and

20 means for generating a data record pointer associated with said modified version of said data record to identify said new memory location.

1/2




FIG. 1.

2/2




FIG. 2.


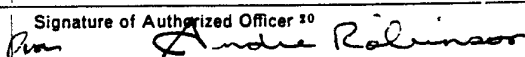


FIG. 3.

INTERNATIONAL SEARCH REPORT

International Application No PCT/US90/05498

I. CLASSIFICATION OF SUBJECT MATTER (if several classification symbols apply, indicate all) ³		
According to International Patent Classification (IPC) or to both National Classification and IPC		
IPC (5): G06F 12/08, 12/10, 13/00		
U.S. CL. 364/200,900		
II. FIELDS SEARCHED		
Minimum Documentation Searched ⁴		
Classification System	Classification Symbols	
U.S. CL.	364/200,900	
Documentation Searched other than Minimum Documentation to the Extent that such Documents are Included in the Fields Searched ⁵		
III. DOCUMENTS CONSIDERED TO BE RELEVANT ¹⁴		
Category *	Citation of Document, ¹⁶ with indication, where appropriate, of the relevant passages ¹⁷	Relevant to Claim No. ¹⁸
Y	US, A, 4,577,274 (HO et al) 18 MARCH 1986 See the entire document	1-7
Y	US, A, 4,742,450 (DUVALL et al) 03 MAY 1988 See the entire document	1-7
Y	US, A, 4,300,192 (COULEUR et al) 10 NOVEMBER 1981 See the entire document	1-7
Y	US, A, 3,317,898 (H. HELLERMAN) 02 MAY 1967 See the entire document	1-7
Y, P	US, A, 4,956,771 (NEUSTAEDTER) 11 SEPTEMBER 1990 See the entire document	1-7
<div style="display: flex; justify-content: space-between;"> <div style="width: 48%;"> <p>* Special categories of cited documents: ¹⁵</p> <p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier document but published on or after the international filing date</p> <p>"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p> </div> <div style="width: 48%;"> <p>"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention</p> <p>"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step</p> <p>"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.</p> <p>"&" document member of the same patent family</p> </div> </div>		
IV. CERTIFICATION		
Date of the Actual Completion of the International Search ²	Date of Mailing of this International Search Report ³	
03 JANUARY 1991	13 MAR 1991	
International Searching Authority ¹	Signature of Authorized Officer ²⁰	
ISA/US	 GOPAL C. RAY	