




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification⁴ : G06F 12/00, 12/06, 12/10 G06F 7/00, 7/06, 7/10 G06F 7/22</p>	<p>A1</p>	<p>(11) International Publication Number: WO 89/ 01663 (43) International Publication Date: 23 February 1989 (23.02.89)</p>
<p>(21) International Application Number: PCT/US88/02801 (22) International Filing Date: 16 August 1988 (16.08.88) (31) Priority Application Number: 086,417 (32) Priority Date: 17 August 1987 (17.08.87) (33) Priority Country: US (71) Applicant: MAXIMUM STORAGE, INC. [US/US]; 5025 Centennial Blvd., Colorado Springs, CO 80906 (US). (72) Inventors: WOOTEN, David, R. ; 318 West Cheyenne Nt. Blvd., Colorado Springs, CO 80906 (US). TAYLOR, Thomas, M., III ; 3365 Marble Terrace, Colorado Springs, CO 80906 (US).</p>		<p>(74) Agent: BEATON, Glenn, K.; 8400 E. Prentice Ave., Suite 900, Englewood, CO 80111 (US). (81) Designated States: AT (European patent), AU, BE (European patent), CH (European patent), DE (European patent), FR (European patent), GB (European patent), IT (European patent), JP, LU (European patent), NL (European patent), SE (European patent). Published <i>With international search report.</i></p>

(54) Title: SYSTEM FOR ACCESSING INFORMATION STORED AS A LINK-LIST WITH BACK-POINTERS ON AN OPTICAL DISK BY USING THE BACK-POINTERS TO GENERATE A DIRECTORY

(57) Abstract

A system for recording and retrieving data on a recording medium (10), particularly useful in write-once media such as optical disks on which data (42) pertaining to a particular subject is periodically written or changed and is interrupted by other intervening information (39A). A set of data (42) or information (39A) affecting a directory is written onto the disk in a transaction together with back-pointers to a previous transaction affecting that data or directory. The system reads transaction in reverse order until encountering a transaction affecting a desired data (42) or directory (39A), and then reads and stores a description of the data (42) or the directory information (39A) in the transaction. The back-pointer directs the system to another previous transaction affecting that data (42) or directory (39A), and so on until the system has established a complete description of a desired data or directory information by reading all necessary portions of relevant transactions. Information pertaining to particular data or directories is periodically consolidated and recorded onto the medium (10) to reduce the back-pointing process.



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	FR France	ML Mali
AU Australia	GA Gabon	MR Mauritania
BB Barbados	GB United Kingdom	MW Malawi
BE Belgium	HU Hungary	NL Netherlands
BG Bulgaria	IT Italy	NO Norway
BJ Benin	JP Japan	RO Romania
BR Brazil	KP Democratic People's Republic of Korea	SD Sudan
CF Central African Republic	KR Republic of Korea	SE Sweden
CG Congo	LI Liechtenstein	SN Senegal
CH Switzerland	LK Sri Lanka	SU Soviet Union
CM Cameroon	LU Luxembourg	TD Chad
DE Germany, Federal Republic of	MC Monaco	TG Togo
DK Denmark	MG Madagascar	US United States of America
FI Finland		

-1-

5 --SYSTEM FOR ACCESSING INFORMATION STORED AS A LINK-LIST
 WITH BACK-POINTERS ON AN OPTICAL DISK BY USING THE
 BACK-POINTERS TO GENERATE A DIRECTORY--.

BACKGROUND OF THE INVENTION

10 This invention relates to a method and
 apparatus for use with computer information storage
 devices. In particular the invention may be used for
 the organization and control of signals on write-once
 optical disks.

15 Optical information storage devices generally
 include a rotatable disk coated with a recording
 material. The disk is divided into concentric or
 spiral tracks for the digital recording of information.
 The information is recorded on these tracks during the
20 write phase of the machine by modulation of a laser
 light beam radiated on the recording material. The
 thermal energy of the modulated laser radiation
 produces predictable variations in the reflective
 characteristics of the recording material. These
25 variations are detected and decoded by the machine
 during the read phase. This writing and reading
 process is well known in the art, and allows the
 storage of vast quantities of information as compared
 to the quantity of information stored on a magnetic
30 disk.

 A method is necessary to quickly locate and
 read the stored information. One such method is to
 write a block of information at the beginning of the
 disk which addresses the stored information. The
35 address information may contain the disk location for
 the various information blocks on the disk, and may
 also indicate any relationships between those
 information blocks. This method is satisfactory when

40

45

5 it is known at the outset the nature and location of
each block of information that will be written onto the
disk. For example, this method is adequate for music
recordings on optical disks, since the information
format can be determined and the corresponding
10 addresses can be written at the outset.

The method described in the preceding
paragraph is not workable if information is written of
a nature and at a disk location not determinable of the
outset. This is the typical situation in most
15 applications. Information is written from time to
time, beginning on the next available portion of the
disk. The disk ultimately may have hundreds or
thousands of discrete blocks of information scattered
through it. The addresses of these scattered blocks of
20 information could not be written at the outset since
their nature and location was not known at that time.

This problem is peculiar to write-once media.
Magnetic and other erasable media may have address
blocks that can be updated by being erased and
25 rewritten each time an additional information block is
written. But because write-once media cannot be
erased, rewriting the address blocks each time an
information block is added would consume substantial
additional storage capacity. Even in the case of
30 erasable magnetic media, it is sometimes desirable to
continuously add and modify data rather than erase it.
For example, archiving operations using magnetic tapes
often require a complete historical set of data. A
method is useful for tying this historical data
35 together.

Accordingly, it is an object of this invention
to provide a storage device and method with improved
information organization and control features to
overcome or alleviate the disadvantages and problems
40 described above.

5

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a plan view of a typical optical disk capable of utilizing the invention.

10

FIG. 2 shows in schematic form an example of the tree directory system of the invention.

FIG. 3 shows in schematic form the relative arrangement of various blocks of information in a single transaction on one track of the disk.

15

FIG. 4 shows an array of bad sectors in diagrammatical form together with a bad sector map array corresponding thereto.

FIG. 5 shows a flow chart of the steps in reading information recorded in accordance with the present invention.

20

DESCRIPTION OF THE PREFERRED EMBODIMENTS

25

30

FIG. 1 illustrates a write-once optical disk organized in accordance with the invention. While FIG. 1 shows an optical disk, it will be apparent to those skilled in the art that the invention is adaptable to media in other configurations or modes such as magnetic tapes, disks and drums. Further, while the invention is described under the assumption that the disk rotates at a constant angular velocity, it will be apparent to those skilled in the art that the invention is equally applicable to a disk the tracks of which pass the optical head at a constant linear velocity.

35

This disk 10 has a plurality of tracks 12 for the recording of information through the modulated laser light beam (not shown) system described under Background of the Invention. The tracks shown in FIG. 1 are greatly exaggerated for clarity; an actual 5-1/4 inch disk would typically include approximately 15,000 closely adjacent tracks.

40

45

5 The disk 10 is also divided into a plurality
of sectors, shown as sectors 1 through 8 in FIG. 1.
Information blocks on the disk are addressed and
accessed by track and sector number. Thus, each
information block should begin on a sector boundary.
10 For this reason, it is desirable to have a large number
of small sectors when the disk is used for storing
small information blocks, in order to reduce the wasted
storage between the information blocks and the sector
boundaries. Any number of sectors is possible, and the
15 eight sectors in FIG. 1 are merely for illustration.

 The first sector of the first track on the
disk, sector 1 in the outermost track 12 of FIG. 1,
contains the disk volume name. That name ordinarily
occupies only a few bytes in the first sector, but the
20 user may choose a long name for the disk which may
occupy any number of bytes or sectors. The volume
information also contains a pointer directing the
machine to the partition information described in the
following paragraph.

25 The partition information begins on the next
available sector after the volume information. This
information divides the entire disk, other than the
sectors containing volume and partition information,
into a number of discrete partitions. Each partition
30 serves as a single working unit so that information
writing and reading operations are conducted on that
partition without regard to other partitions. The
partition information can be written at the outset.
Alternatively, it can be written later or written
35 partially at the outset and added to later as the need
for partitioning or further partitioning becomes
evident, assuming sufficient sector space is reserved
for the additional partition information. The function
of the partitions will be further discussed below in
40 describing reading operations.

5 The data to be stored on the disk together
with the directories for that information begin on the
first available sector in a partition. Each entry of
file data and directory information from time to time
is referred to as a transaction. The writing and
10: reading of various transactions is described below, and
may be understood with reference to the schematic
diagrams of FIG. 2 and FIG. 3.

 The information directory system is based on
the information directory tree structure shown in
15: FIG. 2. That figure shows a directory arrangement
containing a root "/" with directories A through D
under that root. Under directory A are directories G
through J. Of course, directories B, C and D could
also have directories under them. Similarly, each
20: directory under another directory, such as directories
G through J under directory A, could in turn have
directories under it and so on. Directory A is
identified as "/A" which means directory A immediately
under root. Directory G is identified as "/A/G" which
25: means directory G, under directory A, under root.
Other directories are identified in the same manner.

 Each directory may also have data files listed
under it, indicated as 20 through 33 in FIG. 2. Each
file has a name, and its location in the directory tree
30: is ascertainable by the prefix to the file name. For
example, the files identified as 20, 21 and 22 in
FIG. 2 could have names a, b and c, respectively. The
names together with the appropriate prefix would be
/A/a, /A/b and /A/c. Similarly, the two files under
35: directory /B, identified as 23 and 24 in FIG. 2 could
also be named a and b, but would be distinguished by
their names together with their prefix, which would be
/B/a and /B/b. As yet another example, the two files
under directory /A/J, identified as 32 and 33 in
40: FIG. 2, could be identified as /A/J/a and /A/J/b.

5 This information directory tree structure
permits the creation of any number of files and
directories under root. The list of the files under
root, together with the list of directories under root,
is the root directory. In the same manner, each
10 subdirectory may have an indefinite number of files and
lower directories under it. The list of those files
together with the list of lower directories immediately
under that directory constitute the table of contents
for the directory.

15 It is important to recognize that after the
disk has been used and added to repeatedly, the root
directory and other directories are not normally
written in a single place, since they are created and
altered at various times. Instead, the root directory
20 and other directories are collections of information
scattered throughout the disk, which cumulatively
constitute the table of contents for root and each
directory.

 The root directory is assigned an "item
25 number" of one and each directory and file is assigned
a separate item number in numerical sequence as it is
created. Referring to the example of FIG. 2, assuming
all the directories and files shown were created at the
outset, the root directory would be item number one,
30 directories /A through /D would be item numbers two
through five, directories /A/G through /A/J would be
item numbers six through nine, and the files shown as
20 through 33 would be item numbers 10 through 22,
respectively. The item number system permits a short-
35 hand reference throughout the disk to the directory or
file to which each item number is assigned.

 The first transaction after the partition
information is a root summary which identifies each
file under root and each directory under root in

40

45

5 existence at the time the root summary is written.
After the root summary, the disk contains transactions
of file data and directory information arranged as
shown in FIG. 3. The file data 42 are the blocks of
information actually being stored. The directory
10 information 39A describes the relationship of the
various directories to the information directory tree.

The transaction also contains an End-of-
Transaction ("EOT") record 35. The EOT record includes
a bad sector map 34 and an entry list 36, the locations
15 and number of entries of which are specified by a
section of fixed length variables 38 at the end of the
EOT record. The bad sector map identifies the location
of each sector within the transaction containing
physical defects that prevent accurate writing and
20 reading. Those bad sectors are not used. FIG. 4 shows
a diagram of sectors containing several bad sectors,
together with a bad sector map array 40 for the sectors
identified in the diagram. The first "4" in the bad
sector map array indicates that the fourth sector is a
25 bad sector. The "3" in the bad sector map array
indicates that the next bad sector is the third next
sector. Similarly, the "1" indicates that the next bad
sector is the very next sector, and the last "4"
indicates that the next bad sector is the fourth next
30 one. If there is no bad sector for a predetermined
number of sectors, then a "0" would appear in the bad
sector map array.

Referring again to FIG. 3, the entry list 36
identifies each item number affected by the
35 transaction. As the machine searches through EOT
records for transactions affecting a particular item
number, the machine can immediately go onto another EOT
record without reading the remainder of the transaction
if the sought-after item number is not in the entry
40 list.

5 The entry list information 36 for each item
number directs the machine either to the directory
information 39A in the case of an item number assigned
to a directory, or to the file data description 39 in
10 the case of an item number assigned to a file. The
directory information 39A indicates the relation of the
file data 42 to the various files and directories on
the disk or partition. For example, the directory
information may establish a new file for the
information under a particular directory.
15 Alternatively, the directory information may alter an
existing file or rearrange the system of directories.
The file data description directs the machine to the
file data itself for the sought-after item number and
unbundles that file data.

20 When the disk is being read, it is first
necessary to determine whether the sought-after file is
indeed on the disk. This requires first creating a
table of contents in machine memory for the files and
directories and files under root. The machine cannot
25 merely read the root summary written at the outset,
since the directories and files under root may have
been changed by transactions written subsequent to the
writing of the root summary. Similarly, the machine
cannot first read the root summary and then read each
30 additional transaction affecting the root directory,
since the location of those additional transactions was
not known when the root summary was written. The
machine could read the root summary and then read
consecutively every transaction on the disk in search
35 of changes to the root directory, but that would be
very time consuming.

 Instead, the machine proceeds as follows,
referring to the flow chart of FIG. 5. The machine
locates the last-written sector on the disk through a

40

45

5 procedure called a binary search. Under this
procedure, the machine divides the disk into halves and
reads the sector at the boundary between the halves.
If that sector has been written on, it divides the last
10 half into two quarters and reads the sector at the
boundary between those two quarters. On the other
hand, if the sector between the halves has not been
written on, it divides the first half into two quarters
and reads the sector at the boundary between those two
15 quarters. This process continues, as the disk is
further divided into eighths, sixteenths and so on,
until the machine locates a written sector immediately
adjacent to an unreserved, unwritten sector. This
written sector is the last-written sector.

In the event the disk has been partitioned,
20 the binary search operates on only the partition that
is indicated as relevant in the disk partition
information. For example, if the disk is divided into
nine partitions and the partition information indicates
that the sought after file is on partition six if it is
25 anywhere on the disk, then the machine divides
partition six into halves, quarters and so on to locate
the last written partition six sector.

When the machine locates the last-written
sector on the disk, or on a partition in the case of a
30 partitioned disk, it reads the EOT record to determine
whether the transaction ending on that sector affects
item number one, which the machine associates with the
root directory. Reading the EOT record is accomplished
by reading consecutively the fixed length variable
35 portion 38, the bad sector map 34 and the entry list 36
shown in FIG. 4. If the entry list 36 indicates that
the transaction does not affect item number one, then
the machine moves to the next-to-last transaction and
reads the EOT record for that transaction. Each

40

45

5 transaction EOT record is read in reverse chronological
order until an EOT record is read indicating it affects
the root directory, item number one. When an EOT
record is read affecting item number one, the directory
information in that transaction is also read and stored
10 in machine memory. That directory information also
back-points to the previous EOT record that affected
item number one by directing the machine to the address
for that previous EOT record.

The machine goes through the same procedure in
15 the transaction to which it was back-pointed by the
last transaction affecting item number one; that is,
the machine reads the EOT record and reads the
transaction directory information. The directory
information of that transaction, in turn, back-points
20 it to the previous transaction affecting item number
one, and so on until a root summary is reached.
Because the machine reads only those EOT records that
affect item number one through this back-pointing
process, rather than reading every EOT record, the
25 back-pointing process is a great time saver.

When a root summary is reached, a list of all
files and a list of all directories immediately under
root together with their assigned item numbers have
been entered into machine memory. This constitutes the
30 root directory. With reference to FIG. 2, the machine
now has in memory a list of directories /A through /D.
If the sought-after file does not have a prefix of any
of directories /A through /D, then it is not on the
disk.

35 If the sought-after file does have a prefix of
one of the directories under root, which would be any
one of /A through /D in the example of FIG. 2, the
machine can then establish the list of files and
directories under that directory. As in the case of

40

45

5 the root directory, this is done by moving again to the
 EOT record for the last-written transaction and
 searching in reverse chronological order for the last
 transaction affecting the item number assigned to that
 directory, then following back-pointers to each
 10 previous transaction affecting that item number. For
 example, with reference to FIG. 2, the list of files
 and directories under directory /A would be established
 by searching EOT records for transactions affecting
 item number two. The machine has a complete list of
 15 files and directories under that directory when the
 directory summary for directory /A is reached. If the
 sought-after file is several tiers down in the
 directory arrangement, then this process of searching
 through EOT records and back-pointers must be repeated
 20 several times before reaching that lower tier file.

Upon establishing the list of files and
 directories under any given directory, it may be
 apparent that the sought-after file is not on the disk.
 For example, if the sought-after file is /A/P/a, then
 25 the file is not in the arrangement depicted in FIG. 2.
 It will be apparent that file /A/P/a is not in the
 FIG. 2 arrangement upon establishing the list of files
 and directories under directory /A. That list would
 include only files /A/a, /A/b and /A/c and directories
 30 /A/G, /A/H, /A/I and /A/J. Thus, prefix /A/P as
 required for file /A/P/a will be missing. The machine
 will stop further searching at that point, and will
 indicate that the file is not on the disk.

If it is not apparent that the file is not on
 35 the disk, the machine will continue down through the
 directory and file tree arrangement. Upon reaching the
 file, the machine must establish a table of contents
 for the file. This is done in much the same manner as
 establishing the list of files and directories for each
 40

5 directory above the file in the information directory
tree. The machine moves to the last-written
transaction and searches the EOT records of
10 transactions in reverse chronological order for the
last transaction affecting the item number assigned to
that file, and then follows backpointers to each
previous EOT record affecting that item number. As
each EOT record for transactions affecting that item
15 number is read, the machine also reads the portion of
the file data description pertaining to that item
number. In this manner, the machine gradually
establishes a complete table of contents in machine
memory for the file by the time the point at which the
file was created is reached.

The discussion above supposes that the
20 repeated search through a disk or partition to create a
directory or file table of contents ends each time on
the first summary transaction in the case of a
directory, or the transaction that created the file in
the case of a file. That is not true unless the
25 directory or file was created at the outset with the
root summary. Actually, it is common for files and
directories to be created or destroyed later. If a
file or directory is created later, then the last
search through the disk or partition to complete the
30 directory or file table of contents will end on the
transaction that created the file or directory rather
than at the root summary since nothing preceding that
transaction will affect the table of contents.

If a large number of changes have been made to
35 a directory, then the process of searching through many
transactions with the back-pointing procedure to
establish a directory can be quite time consuming. To
reduce that time, the machine may periodically
consolidate a directory into an updated directory

40

45

5 summary and write the directory summary onto the next
available sectors. In that event, the machine begins a
search in the same manner as if there is no updated
directory summary, by starting at the EOT record for
10 the last written transaction and searching backward for
a transaction affecting the sought-after item number
and then following backpointers to previous
transactions affecting that item number. But when the
machine reaches an updated directory summary, no
15 further searching is required to establish the file
directory. The machine is thereby allowed to skip all
directory information prior to the updated directory
summary. Of course, a separate search will still be
required to establish the directory or file table of
20 contents for each directory and file under that
directory.

The same procedure is applicable to the root
directory and to file tables of contents. If many
changes have been made affecting the root directory or
the file, the machine may consolidate those changes
25 into an updated root summary or file summary. The
machine can then establish the root directory or file
table of contents by following backpointers to the
updated root summary or file summary and no further.

The file system contains several important
30 features to reduce the time-consuming search through
each EOT record before the last transaction is reached
affecting the sought-after file. One feature is a
consolidated item list periodically written onto the
disk. The consolidated item list is a list of the
35 last, for example, 100 items that were affected by
transactions since the previous consolidated item list
written onto the disk, and the location of each such
transaction. This enables the machine to avoid reading
every EOT record until it encounters an EOT record
40

40

45

5 affecting the sought-after item number. Instead it
reads every EOT record until it either encounters an
EOT record affecting the sought-after item number or
encounters a consolidated item list. The consolidated
item list will direct it immediately to the last EOT
10 record affecting the sought-after file.

Alternatively, if the sought-after file is not
on the item list, because it is not one of the last 100
items that were affected by transactions since the last
consolidated item list, then the consolidated item list
15 will direct the machine back to the preceding
consolidated item list. The machine will then search
that consolidated item list, which will direct the
machine to the last EOT record affecting the sought-
after item number or, if that item number is not on
20 that consolidated item list, then will direct the
machine to the preceding consolidated item list and so
on. By this process the machine, in the worst case,
will read EOT's until encountering a consolidated item
list and then jump backward through consolidated item
25 lists.

The machine may also periodically write onto
the disk a full item list naming the address of the
last transaction affecting each item number ever used
on the disk. The full item list ensures that the
30 machine will never need to go further than the point at
which the full item list is written in order to
ascertain the location of the last-written transaction
affecting any particular item number. The full item
list consumes significant storage space, but may be
35 worthwhile on a disk with many file and directory
changes. Because the full item list contains the
addresses of all the item numbers, some storage space
can be saved by omitting the item numbers from the full

40

45

5 item list and instead simply listing the addresses in
sequence corresponding to the numerical sequence of
item numbers.

10 Another time-saving feature is a cache
structure established in machine memory. The cache
structure stores the location of the last-written EOT
15 record for a limited number of item numbers. The item
numbers chosen for the cache structure are generally
those that were written most recently or read most
recently, on the theory that those are the item numbers
most likely to be sought-after next. If the item
20 number is found in the cache structure, then the
machine can go immediately to the last-written EOT
record for that item number.

25 The cache structure may also contain a list
identifying the item number of a limited set of file
names. As in the case of item numbers, if a file name
is found in the cache structure, then the machine can
go directly to the appropriate EOT record without the
necessity of establishing a table of contents for root
and each directory above that file to determine the
30 item number of the file.

35

40

45

CLAIMS

5

What is claimed is:

1. A data processing method, using a recording and reading machine, comprising:

10

recording onto a medium a transaction containing information;

recording onto the medium a subsequent transaction containing subsequent information, related to information in a transaction previously recorded onto the medium; and

15

recording onto the medium in said subsequent transaction a back-pointer to the transaction of said previously recorded information.

2. The method of claim 1, further comprising:

20

establishing an information directory tree of at least one directory and at least one file, every directory and file being under a parent directory until reaching a root directory;

25

recording onto the medium at the time each file and directory is created a summary of the current contents of such file and directory;

assigning each directory and file an item number;

30

recording in each transaction as an entry list, the item number of each directory and file affected by such transaction;

35

in the case of a transaction affecting a file, recording in said transaction a file data description describing said file effect; and

in the case of a transaction affecting a directory, recording in said transaction a set of directory information describing said directory effect.

40

5 3. The method of claim 2, wherein the machine establishes a list of files and directories under a parent directory by:

10 reading the entry list of each transaction recorded onto the medium, beginning with the last-written transaction and continuing in reverse order, until reading the item number assigned to the parent directory;

15 reading and storing in machine memory the directory information pertaining to that item number for the transaction containing said item number, and following a back-pointer in said transaction to the previous transaction affecting that item number; and

20 continuing to read and store in machine memory the directory information pertaining to that item number for each transaction to which a back-pointer leads, and following another back-pointer in said transaction to a previous transaction affecting the item number assigned to that directory until reaching a directory summary.

25 4. The method of claim 3, further comprising, in the case that said parent directory has other parent directories above it in the information directory tree, performing the steps of claim 3 for each directory above said parent directory until reaching said parent directory, and then performing said steps for said parent directory.

30 5. The method of claim 2, wherein the machine establishes a table of contents for a file by:

35 reading the entry list of each transaction recorded onto the medium, beginning with the last-written transaction and continuing in reverse order, until reading the item number assigned to said file;

40

45

5 reading and storing in machine memory the
file data description pertaining to that item number
for the transaction containing said item number, and
following a back-pointer in said transaction to the
previous transaction affecting that item number; and

10 continuing to read and store in machine
memory the file data description pertaining to that
item number for each transaction to which a back-
pointer leads, and following another backpointer to a
previous transaction affecting that item number until
15 reaching a file summary.

20 6. The method of any of claims 3, 4 or 5,
further comprising recording a consolidated item list
of the item numbers that were affected by transactions
recorded after a predetermined number of transactions
are recorded and wherein the machine reads the entry
list of each transaction recorded onto the medium in
reverse order, only until either (i) reading an entry
list containing said item number whereupon the machine
follows back-pointers to the previous transaction
25 affecting said item number, or (ii) reading a
consolidated item list whereupon the machine is
directed to the last-written transaction affecting said
item number if said item number is on the consolidated
item list or the machine is directed to and reads the
30 preceding consolidated item list in the same manner if
said item number is not on the consolidated item list.

35 7. The method of any of claims 3, 4 or 5
wherein each of said summaries is recorded onto the
medium after a predetermined amount of changes have
been made to such summary, so that the machine can
avoid reading directory information and file data
descriptions beyond such predetermined amount.

40 8. The method of any of claims 3, 4 or 5
further comprising:

5 storing in a cache structure in machine
memory the item numbers of a predetermined number of
the item numbers last affected by transactions recorded
onto the medium and a back-pointer showing the location
of the last transaction affecting each such item
number;

10 prior to reading any transactions in
search of said item number, reading the cache structure
in search of said item number; and

15 if said item number is in the cache
structure, then following said back-pointer directly to
the last transaction affecting said item number.

20 9. A data storage medium device comprising:
a transaction containing information;
a subsequent transaction containing
subsequent information, related to information in a
transaction previously contained on the medium; and
a back-pointer in said subsequent
transaction to said transaction previously contained on
the medium.

25 10. The device of claim 9, wherein said
information and previously contained information is
organized in a directory information tree of at least
one directory and at least one file, every directory
and file being under a parent directory until reaching
a root directory, further comprising:

30 a summary of the current contents of each
file and directory, recorded onto the medium at the
time each such file and directory is created;

35 separate item numbers assigned to each
file and directory;

an entry list recorded in each
transaction, containing the item numbers of each
directory and file affected by such transaction;

40

45

5 in the case of a transaction affecting a directory, a set of directory information in such transaction describing said effect; and

10 in the case of a transaction affecting a file, a file data description in such transaction describing said effect.

11. The device of claim 10, further comprising means to establish a list of files and directories under a parent directory.

15 12. The device of claim 11, wherein such means to establish a list of files and directories under a parent directory, includes means for:

20 reading the entry list of each transaction recorded onto the medium, beginning with the last-written transaction and continuing in reverse order, until reading the item number assigned to the parent directory;

25 reading and storing in machine memory the directory information pertaining to that item number for the transaction containing said item number, and following a back-pointer in said transaction to the previous transaction affecting that item number; and

30 continuing to read and store in machine memory the directory information pertaining to that item number for each transaction to which a back-pointer leads, and following another backpointer to a previous transaction affecting the item number assigned to that directory until reaching a directory summary.

35 13. The device of claim 12, further comprising, in the case that said parent directory has other parent directories above it in the directory information tree, means for repeating the steps of claim 12 for said other parent directories above said parent directory until reaching said parent directory, and then performing said steps for said parent
40 directory.

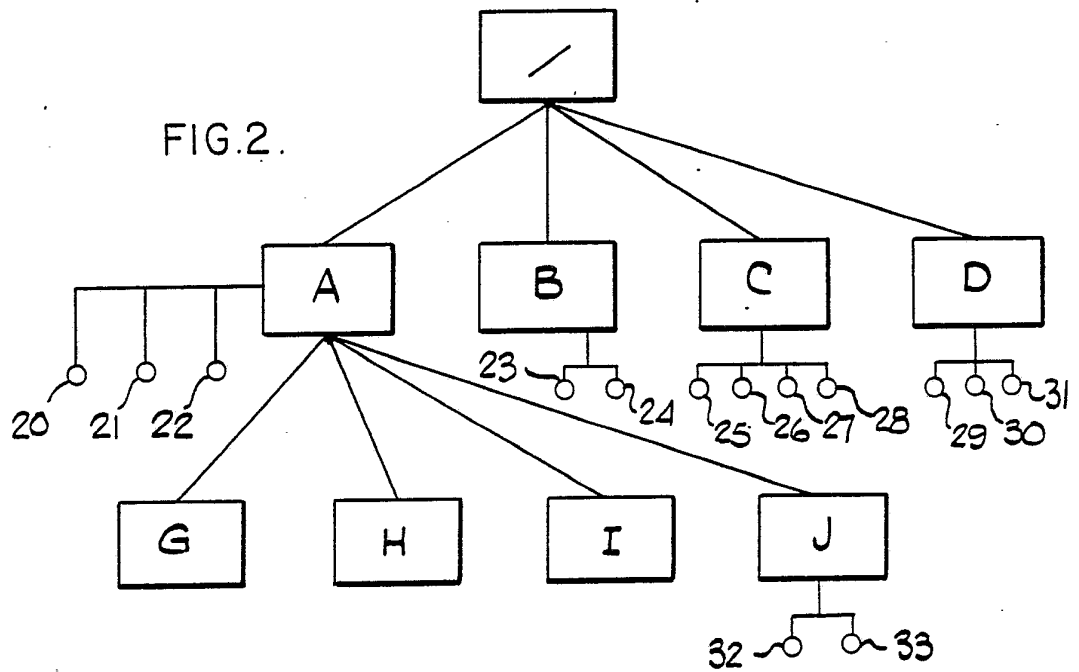

5 14. The device of claim 12 further comprising
a consolidated item list of the item numbers that were
affected by transactions recorded after a predetermined
number of transactions are recorded, so that the
machine reads the entry list of each transaction,
10 recorded onto the medium in reverse order, only until
either (i) reading an entry list containing said item
number whereupon the machine follows back-pointers to
the previous transaction affecting said item number, or
(ii) reading a consolidated item list whereupon the
15 machine is directed to the last-written transaction
affecting said item number if said item number is on
the consolidated item list or the machine is directed
to and reads the preceding consolidated item list in
the same manner if the item number is not on the
20 consolidated item list.

25 15. The device of claim 12, further
comprising an updated directory summary recorded onto
the medium after a predetermined amount of changes have
been made, so that the machine can avoid reading
information beyond such predetermined amount.

30 16. The device of claim 12, further
comprising a cache structure in machine memory
containing the item numbers of a predetermined number
of the item numbers last affected by transactions
recorded onto the medium and a back-pointer showing the
location of the last transaction affecting each such
item number so that prior to reading any transactions
in search of said item number, the machine can read the
cache structure in search of said item number and if
35 such item number is in the cache structure, then follow
said back-pointer directly to the last transaction
affecting said item number.

40

45






FIG. 3.


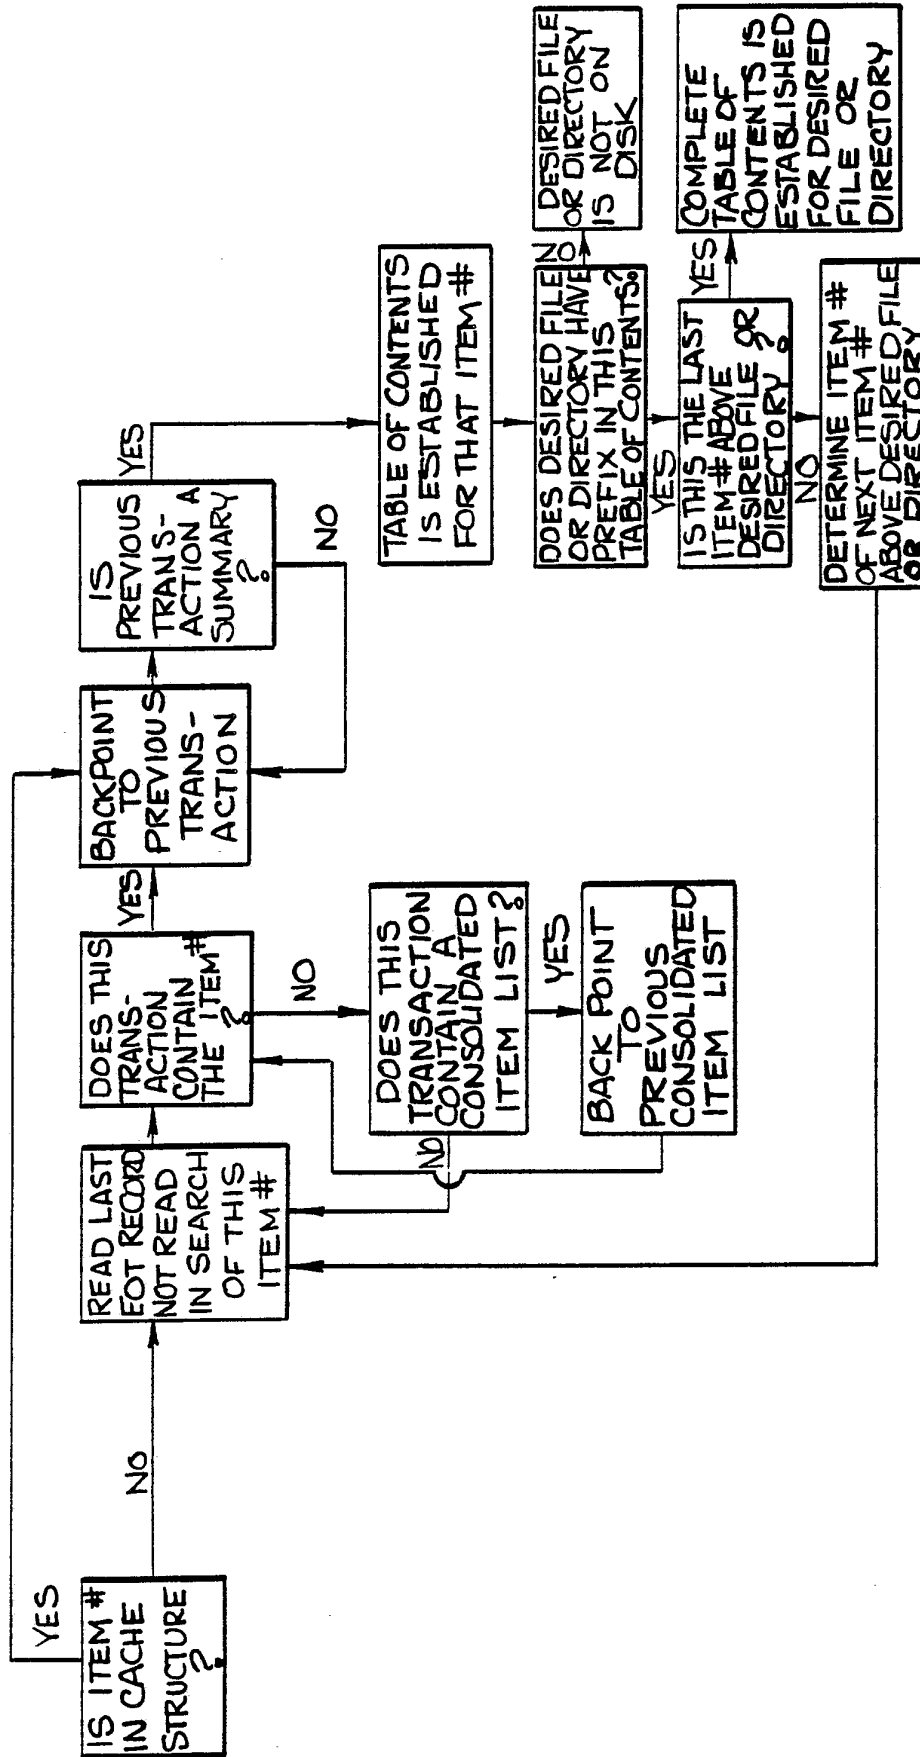


FIG. 4.

FIG. 5.



INTERNATIONAL SEARCH REPORT

International Application No. PCT/US88/02801

I. CLASSIFICATION OF SUBJECT MATTER (if several classification symbols apply, indicate all) ⁶				
According to International Patent Classification (IPC) or to both National Classification and IPC IPC(4): G06F 12/00, 12/06, 12/10, 7/00, 7/06, 7/10, 7/7/22 U.S. CI. 364/200				
II. FIELDS SEARCHED				
Minimum Documentation Searched ⁷				
Classification System	Classification Symbols			
U.S.	364/200 & 364/900, --MSFile. 364/300.			
Documentation Searched other than Minimum Documentation to the Extent that such Documents are Included in the Fields Searched ⁸				
III. DOCUMENTS CONSIDERED TO BE RELEVANT ⁹				
Category *	Citation of Document, ¹¹ with indication, where appropriate, of the relevant passages ¹²	Relevant to Claim No. ¹³		
P, X	US, A, 4,748,320 (YORIMOTO ET AL.) 31 May 1988 --See entire document--.	1-16		
X	US, A, 4,086,628 (WOODRUM) 25 April 1988 --See entire document--.	1-16		
<table style="width: 100%; border: none;"> <tr> <td style="width: 50%; vertical-align: top; border: none;"> <p>* Special categories of cited documents: ¹⁰</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> </td> <td style="width: 50%; vertical-align: top; border: none;"> <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> </td> </tr> </table>			<p>* Special categories of cited documents: ¹⁰</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>	<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>
<p>* Special categories of cited documents: ¹⁰</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>	<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>			
IV. CERTIFICATION				
Date of the Actual Completion of the International Search	Date of Mailing of this International Search Report			
29 SEPTEMBER 1988	22 NOV 1988			
International Searching Authority	Signature of Authorized Officer			
ISA/US	Robert B. Harrell			