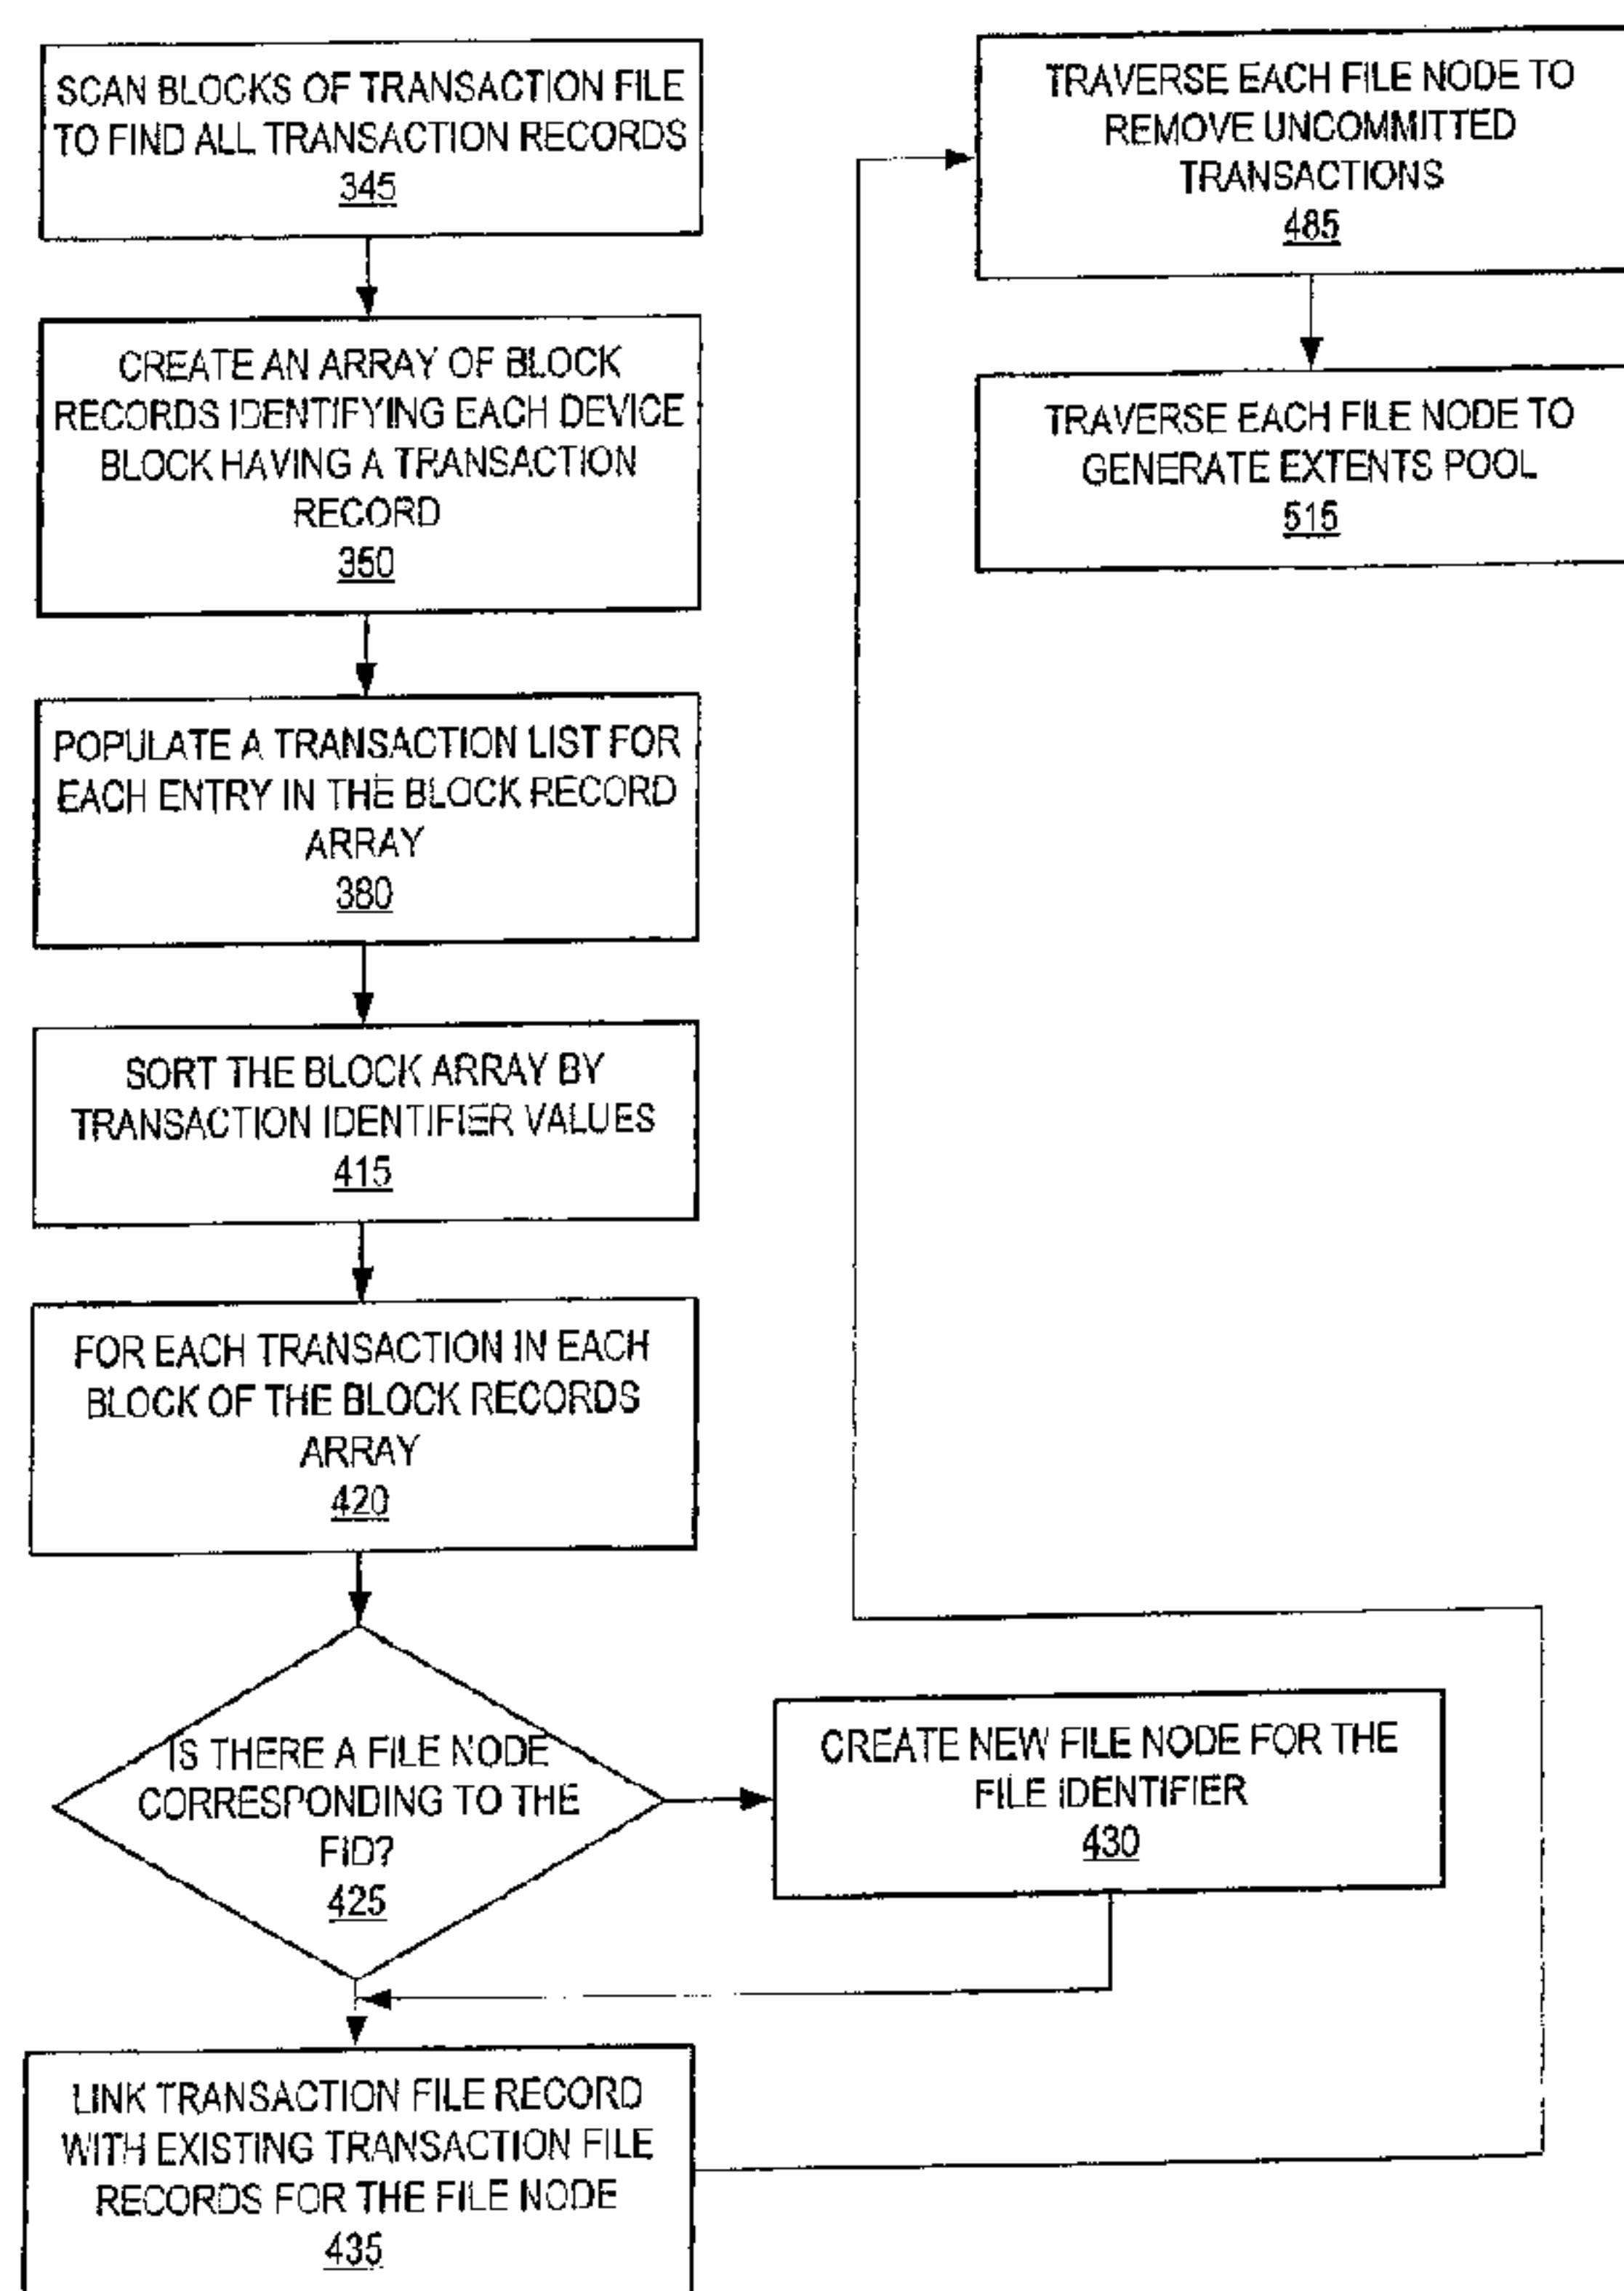




(22) Date de dépôt/Filing Date: 2006/06/07
 (41) Mise à la disp. pub./Open to Public Insp.: 2007/01/01
 (45) Date de délivrance/Issue Date: 2015/10/06
 (62) Demande originale/Original Application: 2 550 974
 (30) Priorité/Priority: 2005/07/01 (US11/173798)

(51) Cl.Int./Int.Cl. *G06F 17/30* (2006.01),
G06F 11/00 (2006.01)
 (72) Inventeur/Inventor:
DODGE, DAN, CA
 (73) Propriétaire/Owner:
2236008 ONTARIO INC., CA
 (74) Agent: RIDOUT & MAYBEE LLP

(54) Titre : VERIFICATION OPTIMISEE AU DEMARRAGE DE L'INTEGRITE DU SYSTEME D'ARCHIVAGE
 (54) Title: OPTIMIZED STARTUP VERIFICATION OF FILE SYSTEM INTEGRITY



(57) **Abrégé/Abstract:**

A computer system having a transaction based file system is disclosed. The computer system includes the system software that manages the file data and the file system structure of files stored on a persistent data storage device and maintains a transaction file that includes a plurality of transaction records. Each of the transaction records has a header section and a data section. The header section of each transaction record includes one or more fields that are designated to store information corresponding to a file transaction that is represented by the transaction record. The file system software executes a startup process in which a reconstructed file system is generated in random access memory. The startup process skips verification of the data section of a transaction record when the transaction record meets one or more predetermined criterion.

ABSTRACT

A computer system having a transaction based file system is disclosed. The computer system includes file system software that manages the file data and the file system structure of files stored on a persistent data storage device and
5 maintains a transaction file that includes a plurality of transaction records. Each of the transaction records has a header section and a data section. The header section of each transaction record includes one or more fields that are designated to store information corresponding to a file transaction that is represented by the transaction record. The file system software executes a startup process in which a
10 reconstructed file system is generated in random access memory. The startup process skips verification of the data section of a transaction record when the transaction record meets one or more predetermined criterion.

OPTIMIZED STARTUP VERIFICATION OF FILE SYSTEM INTEGRITY

BACKGROUND OF THE INVENTION

TECHNICAL FIELD

5 **[0001]** This invention is generally directed to a file system for use in a computer, embedded controller, or the like. More particularly, this invention is directed to a transaction based file system in which the startup verification of the file system integrity is optimized.

RELATED ART

10 **[0002]** Computers, embedded controllers, and other microprocessor based systems are typically constructed from a variety of different hardware components. The hardware components may include a processor, I/O devices, human interface devices, and the like. Additionally, such systems use memory storage units to
15 maintain the data used in the system. The memory storage units may take on a variety of different forms including, but not limited to, hard disk drives, floppy disk drives, random access memory, flash memory, and the like.

[0003] High-level application programs that are executed in such systems must often interact seamlessly with these hardware components, including the memory storage units. To this end, many systems run an operating system that
20 acts as an interface between the application programs and the system hardware. File system software may be included as part of the operating system, or it may be provided as an ancillary software component that interacts with the operating system. In either instance, the file system software organizes the data within the memory storage units for ready access by the processor and the high-level
25 application programs that the processor executes.

[0004] There are a number of different file system classifications since there are many ways to implement a file system. For example, a transaction based file system is one in which the file system is always maintained in a consistent state

since all updates to the file system structure and the data are logged as transactions to a transaction file. More particularly, all updates to the file system are made as transactions within the transaction file, and the contents of the file system are dynamically re-constituted by successively applying all of the
5 transactions that have been committed.

[0005] A transaction in the transaction file is either committed or it has not been completed. If the operation of the file system is interrupted, such as due to a power outage, for example, the state of the file system can be restored by consulting the contents of the transaction file. Any committed transactions are
10 used by the file system, and any transactions that are not complete are rolled back, restoring the file system to the state it was in prior to the attempted update.

[0006] Restoration of the file system to a consistent state requires that the file system software execute a predetermined startup process. During a typical startup process, the integrity of each transaction stored in the transaction file is
15 verified before it becomes part of the file system. Additional file system operations also may be executed during the startup process. The traditional manner in which transaction verification and other file system operations are performed after a file system interruption, however, is often sub-standard in that the operations are time, process and resource intensive.

20 **SUMMARY**

[0007] A computer system that may be used in implementing a transaction based file system is disclosed. The computer system includes a processor, random access memory that is accessible by the processor, and a persistent data storage device that is likewise accessible by the processor. The computer system also may
25 includes file system software. The file system software may be executed by the processor and operates to manage the file data and the file system structure of the files stored on the persistent data storage device. Additionally, the file system software may maintain a transaction file that includes a plurality of transaction records. Each of the transaction records has a header section and a data section.

The header section of each transaction record may include one or more fields that are designated to store information corresponding to a file transaction that is represented by the transaction record. If the operation of the file system software is interrupted, the file system software resumes its execution using a startup
5 process in which a reconstructed file system is generated in the random access memory. During system restart, the startup process may skip verification of the data section of a transaction record when the transaction record meets one or more predetermined criterion. For example, the startup process may make a distinction between transaction records merely affecting file data versus transaction records
10 that affect the metadata of the file system. As transactions are found during the startup process, the file system software may identify whether a transaction impacts file data or metadata. Since only the metadata is required to ensure that the file system is in a consistent state after startup, the transaction records relating to metadata may be selected as the only subset of transaction records that are
15 subject to complete verification. Verification of other transaction records may, for example, be limited to a check of the information contained in the header section of each remaining transaction record.

[0008] The transaction file may be stored, for example, in flash memory. In such instances, the startup process may be further enhanced. For example, the
20 startup process may limit its header information verification to the first transaction record of a sequence of transaction records in the same block of the flash memory device. Neither the header nor data sections of the trailing transaction records of the sequence are verified during startup and the startup process moves on to processing the transaction records of the next device block, if any. Still further, the
25 startup process may check the header section information to determine whether the memory locations in a device block have been erased or retired. If the memory locations in the device block have been erased or retired, startup processing continues with the next device block.

[0009] Other systems, methods, features and advantages of the invention will
30 be, or will become, apparent to one with skill in the art upon examination of the following figures and detailed description. It is intended that all such additional

systems, methods, features and advantages be included within this description, be within the scope of the invention, and be protected by the following claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] The invention can be better understood with reference to the following
5 drawings and description. The components in the figures are not necessarily to scale, emphasis instead being placed upon illustrating the principles of the invention. Moreover, in the figures, like referenced numerals designate corresponding parts throughout the different views.

[0011] Figure 1 is a block diagram of a computer system that may implement
10 a transaction based file system in which startup verification of the file system integrity is optimized.

[0012] Figure 2 is a tree diagram showing one example of an arrangement of files and directories that may be implemented in the transaction based file system.

[0013] Figure 3 is a block diagram illustrating one manner in which records of
15 a metafile may be arranged to implement the file system structure shown in Figure 2.

[0014] Figure 4 illustrates one manner of logically arranging a transaction record in a transaction file of the transaction based file system.

[0015] Figure 5 shows the physical arrangement of memory in one type of
20 flash media device.

[0016] Figures 6 and 7 illustrate various manners in which transaction records may be arranged in flash media devices for use in the transaction based file system.

[0017] Figure 8 illustrates a number of interrelated processing steps that may
25 be used to generate an extents pool that, in turn, is employed in a reconstructed file system that is created by the computer system during startup.

[0018] Figures 9 through 11 are directed to exemplary formats for various record types used in the processing steps shown in Figure 8.

[0019] Figure 12 is directed to an exemplary format for a directory node record of the regenerated file hierarchy used in the reconstructed file system.

5 **[0020]** Figure 13 is directed to an exemplary format for a file node record of the regenerated file hierarchy used in the reconstructed file system.

[0021] Figure 14 illustrates a number of interrelated processing steps that may be used to construct the regenerated file hierarchy used in the reconstructed file system.

10 **[0022]** Figure 15 is a logical representation of a reconstructed file system that has been generated in the manner set forth in connection with Figures 8 through 14 as applied to the exemplary file and directory arrangement shown in Figure 2.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0023] Figure 1 illustrates the components that may be employed in an exemplary transaction based computer system 10. As shown, the exemplary system 10 includes a processor 15, read only memory 20, and a persistent storage unit 30. Computer system 10 also may include random access memory 35, an I/O interface 40, and a user interface 45. The specific components that are used in computer system 10 may be tailored to the particular function(s) that are to be executed by the computer system 10. Accordingly, the presence or absence of a component, other than processor 15, may be specific to the design criterion imposed on the computer system 10. For example, user interface 45 may be omitted when the computer system 10 takes the form of an embedded controller or the like.

25 **[0024]** Read only memory 20 may include operating system code 43 that controls the interaction between high-level application programs executed by the processor 15 and the various hardware components, including memory devices 20

and 35, the persistent storage unit 30, and the interface devices 40 and 45. The operating system code 43 may include file system software for organizing files stored on the persistent storage unit 30. Alternatively, the file system software may be provided as a separate software component that merely interacts with the
5 operating system code 43. In the latter case, the code corresponding to the file system software may be stored in read only memory 20, persistent storage unit 30 or the like. When computer system 10 is networked with other computers and/or storage devices through I/O interface 40, the file system software may be stored remotely and downloaded to computer system 10 as needed. Figure 1, however,
10 illustrates storage of the file system software 47 in read only memory 20.

[0025] The persistent storage unit 30 may take on any number of different forms. For example, the persistent storage unit 30 may take the form of a hard disc drive, floppy disk drive, and the like. It also may be in the form of a non-rotating media device, such as non-volatile memory implemented in an integrated
15 circuit format (*e.g.*, flash memory, and the like.). Still further, persistent storage unit 30 need not be limited to a single memory structure. Rather, the persistent storage unit 30 may include a number of separate storage devices of the same type (*e.g.*, all flash memory) and/or separate storage devices of different types (*e.g.*, one or more flash memory units and one or more hard disk drives).

20 **[0026]** The files stored in the persistent storage unit 30 include data that is interpreted in accordance with a predetermined format used by an application program or by the operating system code 43. For example, the data stored within a file may constitute the software code of an executable program, the ASCII text of a database record, data corresponding to transactions executed (or not executed)
25 by computer system 10, and the like.

[0027] In this exemplary system 10, the file system software 47 organizes the files stored on the persistent storage unit 30 using an inverted hierarchical structure. Figure 2 is a diagram showing one manner in which the inverted hierarchical structure, shown generally at 50, may be implemented. In the
30 traditional hierarchical structures used by many file systems, the top level of the file

structure begins with the root directory and each directory points downward to the files and subdirectories contained within the directory. In the exemplary inverted hierarchical structure 50, however, the child files and child directories contained within a parent directory point upward to the parent directory. Depending on
5 where the file system begins its organization, the root directory may constitute the lowest level of the file system structure.

[0028] The exemplary inverted hierarchical structure 50 includes five files 55, 60, 65, 70 and 75 at the highest level of the file system structure. Files 55, 60 and 65 are contained within directory 80 while files 70 and 75 are contained within
10 directory 85. Accordingly, the file system software 47 organizes the file system so that the file system records representing child files 55, 60 and 65 point to the record for their parent directory 80. Similarly, file system records representing child files 70 and 75 point to the record for their parent directory 85.

[0029] At the next level of the exemplary inverted hierarchical structure 50,
15 files 90 and 95 as well as directory 80 are contained within directory 100, while directory 85 may be contained within directory 105. Accordingly, the file system software 47 organizes the file system so that file system records representing child directory 80 and child files 90 and 95 point to the record for their parent directory 100. Similarly, the file system record representing child directory 85 points to the
20 record for its parent directory 105.

[0030] The root directory 110 may form the trunk of the inverted hierarchical structure 50. In this example, directories 100 and 105 and file 115 are contained within the root directory 110. Accordingly, the file system software 47 organizes the file system so that file system records representing child directories 100 and 105
25 and child file 115 point to the record for their parent directory 105.

[0031] One manner in which the file system software 47 may organize the records of the file system to implement an inverted hierarchical structure is shown in Figure 3. In this implementation of the file system, the file system software 47 may generate one or more metafiles that include records corresponding to each file

and directory used in the file system. Figure 3 shows a single metafile 120 and an exemplary manner in which the records within the metafile 120 may be arranged and formatted. In this example, metafile 120 may be arranged as a table that includes a plurality of equal length record entries 125. Each record entry 125
5 corresponds to a single file or directory may be used in the file system. A unique file identifier, such as the one shown at 130, may be used by the file system software 47 to address a corresponding record 125 of the metafile 120. If each record entry 125 has the same record length, the format for the file identifier 130 may be chosen so that it may be used, either directly or indirectly, as an index to
10 the desired record in metafile 120. For example, file identifier 130 may constitute an offset value that may be used along with the memory address location of the first record of metafile 120 to calculate the memory address location of the first byte of the metafile record having the desired directory/file information.

[0032] In the example of Figure 3, the file identifier 130 is pointing to record
15 135 (Entry 7) in metafile 120. Record 135 is shown in Figure 3 in an expanded form adjacent to the metafile 120. The expanded form of record 135 also illustrates a basic record format that may be used for each record entry 125. In this example, record 135 includes a number of different fields containing information relating to the file or directory represented by the record. This information, among
20 other things, corresponds to the logical location of the file or directory within the structure of the file system.

[0033] The inverted hierarchical structure of the file system may be implemented by employing a metafile record format in which each metafile record includes a pointer to the metafile record representing its parent directory. Figure 3
25 shows a metafile record format in which each metafile record includes a parent identifier field 140 that stores the file identifier of its parent directory. In this example, the parent record identifier 140 of metafile record 135 corresponds to the file identifier used to address record 145 (Entry 9). Record 145, in turn, includes information pertaining to the directory containing the file or directory represented
30 by record 135.

[0034] Each metafile record also may include other information pertaining to the directory or file that the record represents. In the exemplary record format of record 135, a number of different information fields are employed. The information fields include a mode field 150, user identification field 155, group identification field 160, access time field 165, modified time field 170, created time field 175, file size field 180 and short name field 185. The mode field 150 may be used to determine whether the file or directory represented by the record is a system file/directory, a hidden file/directory, a read only file/directory, and the like. The user identification field 155 and group identification field 160 contain information relating to user and group ownership of the represented file or directory. The access time field 165, modified time field 170, and created time field 175 contain information relating to the time at which the represented file or directory was last accessed, the time at which the represented file or directory was last modified and the time at which the represented file or directory was created, respectively. The size field 185 contains information on the size of the file represented by the record and is zero for directory records. Finally, the short name field 185 contains ASCII characters representing the short text name of the corresponding file or directory. The length of the short name field 185 may be chosen, for example, to conform to the POSIX standard. Additionally, each record may include hash values and/or name sums that correspond to the short name. Such hash values and/or name sums may be used by the file system software 47 to quickly search for a particular directory and/or file record.

[0035] Each record in metafile 120 also may include a field for an extended record identifier 190. The extended record identifier 190 may be used as a file identifier that points to an extended record in the metafile 120. The extended record may contain further information for the file or directory represented by the record and may be particularly useful in instances in which all of the information pertaining to a particular file or directory does not fit within the memory space allocated for a single metafile record.

[0036] Figure 3 illustrates one manner in which an extended record identifier 190 may be used. In this example, the extended record identifier 190 of record

135 corresponds to the file identifier (fid) used to access record 195 (Entry 11) in metafile 120. An exploded view of record 195 is shown adjacent the exploded view of record 135 in Figure 3. This exploded view illustrates one record format that may be used for the extended record. As shown, each extended record may include
5 its own parent identifier field 200. The parent identifier field 200 of an extended record, however, corresponds to the file identifier of the record which points to the extended record. In the example shown in Figure 3, the contents of the parent identifier field 200 may be used to point back to record 135 (Entry 7).

[0037] In those instances in which the memory space allocated for two record
10 entries is insufficient to hold all of the information pertaining to a file or directory, the extended record 195 may point to yet a further extended record using its own extended record identifier, such as the one included in field 205 of record 195. Although the format for the further extended record pointed to by extended file identifier 125 is not shown, the further extended record may likewise include a
15 parent record identifier that points back to record 195.

[0038] The type of information included in an extended record may vary between file systems. In Figure 3, the extended record 195 includes a long name field 210 that contains ASCII characters corresponding to the text of the long name of the file or directory represented by the record 135. Further fields may be
20 reserved in an expansion area 215 of each extended record, such as record 195, to store additional information relating to the corresponding file or directory.

[0039] In the previous example, the extended records used by the file system are stored in metafile 120. However, the extended records and any further extended records may alternatively be stored in a separate metafile, multiple
25 metafiles, and the like. The separate metafile(s) need not share the same storage medium with metafile 120 nor with each other. Rather, the metafiles may be stored in different storage media accessible to processor 15. Even the basic metafile records (directory and file records that do not have corresponding extended records) may be distributed among multiple files and/or multiple storage
30 media. As such, although the metafile records of the exemplary system are stored

in a single metafile, the metafile may alternatively be in the form of many individual files on the same or different storage media.

[0040] By organizing the files and directories of computer system 10 in an inverted hierarchical structure, it becomes possible to realize one or more file system advantages. For example, the file system is capable of being implemented in any manner in which typical file and directory transactions (*i.e.*, moving a file/directory, deleting a file/directory, creating a file/directory, copying a file/directory) are accomplished atomically as a change, addition or deletion of a single metafile record. In this implementation, for example, the file/directory represented by record 135 may be moved to another directory in the hierarchy merely by changing the parent identifier 140 so that it points to the metafile record for the new parent directory. This may be accomplished with a single write operation to record 135 in the metafile 120.

[0041] The inverted hierarchical structure may be employed to optimize a transactional or log-based system. An exemplary transactional or log-based system may be constructed from the components shown in Figure 1. In this example, a transaction file 220 may be maintained in the persistent storage unit 30 and may be used to keep records of the transactions associated with each file and directory of the file system. Updates to the file system are committed atomically based on the transaction records contained in transaction file 220. In one of its simplest forms, every transaction record may be stored as a single logical page that may be mapped to a physical block or sector of the persistent storage unit 30.

[0042] One manner in which a transaction record 225 may be formatted for use in computer system 10 is shown in Figure 4. Generally stated, each transaction record 225 of the transaction file 220 includes a header field 230 and a corresponding data field 230. The header field 230 may include a number of different sub-fields. The sub-fields shown in Figure 4 include a transaction sequence field 240, a file identification field 245, a transaction status field 250, a cluster high field 255, a cluster low field 260 and number of clusters field 265. Additionally, further sub-fields may be included in header 230 to verify the integrity

of the transaction and for error correction. These further sub-fields include a cluster sum field 247, a transaction sum field, an error correction code field 257 to check and correct header 230, an error correction code field 259 to check and correct data 235, and a further status field 262 indicative of the condition of the memory locations in which the transaction record may be stored.

[0043] Each of the sub-fields of header field 230 has a meaning to the file system software 47. In this example, the transaction sequence field 240 may be a monotonically increasing transaction identifier that may be assigned by the file system software 47. When a new transaction record may be added to the transaction file 220, the value stored in the transaction sequence field 240 of the new record may be increased by a predetermined amount over the value of the transaction sequence field of the chronologically preceding transaction record. Consequently, transaction records having larger transaction identifier values are considered to have been added to the transaction file 220 later in time than transaction records having lower transaction identifier values. This chronological sequencing of the transactions, as represented by the value of the transaction sequence field 240 (and, in certain circumstances, the position of the transaction record within a block of the transaction file 220), allows the file system software 47 to apply (*i.e.*, commit) the transactions in the proper order to maintain the integrity of the file system contents. Other ways of keeping track of the chronological sequencing of the transactions also may be used.

[0044] File system software 47 uses the transaction status field 250 to determine whether the transaction of a transaction record 225 has been committed. Once a transaction has been committed, further alteration of the committed transaction record 225 may be inhibited by the file system software 47. This ensures consistency of the file system and also allows the file system to store the transaction file 220 in, for example, write-once media, flash media, or the like.

[0045] The file identification field 245 of header 230 identifies the file that may be affected by the transaction record 225. The format for the file identification field 245 may be selected so that it is the same as the file identifiers used in the

metafile records. The cluster high field 255 and cluster low field 260 may be used by the file system software 47 to determine the starting address (or offset) at which the data 235 may be to be written into the identified file while the number of clusters field 265 may be used to determine how many clusters of the identified file are to be overwritten by the data 235.

[0046] As noted above, persistent storage unit 30 may include one or more flash memory devices. Flash memory devices store information in logic gates, called "memory cells," each of which typically stores one bit of information. More recent advances in flash memory technology have also enabled such devices to store more than 1 bit per cell, sometimes referred to as multi-level cell devices. Additionally, flash memory is non-volatile, which means that the contents of memory cells are not lost when power is withdrawn from the device.

[0047] Although flash device technology is continuously evolving, dominant technologies include NAND flash memory and NOR flash memory. NOR flash devices and NAND flash devices generally differ in the type of logic gate used for each storage cell. An exemplary logical architecture 270 of one type of NAND flash memory device 270 is shown in Figure 5. As illustrated, the available memory on the device 270 may be organized into contiguous physical blocks 280 each having an equal number of memory cells (*i.e.*, 16K bytes). NAND flash memory device 270 further divides each of the contiguous blocks 280 into a specific number of physical sectors or pages 290. Each physical page 290, in turn, may be further divided into a data area 295 and spare area 300. The data area 295 is normally reserved for storage of data, while the spare area 300 is typically reserved for maintenance of meta-information about the data stored in data area 295. The meta-information may include, for example, error-correcting codes used for verification and correction of sector contents, cyclic redundancy check data, and the like..

[0048] NOR flash devices have an architecture similar to that shown in Figure 5, except that the spare areas of each page are located on opposite sides of the data area. NOR flash devices also offer random access read and programming operations, allowing individual memory locations to be read on or read. However,

once a memory location in a block has been written, NOR flash devices do not allow the block to be rewritten a smaller granularity than a block. Likewise, NOR flash devices do not allow erase operations at a smaller granularity than a block. insert quick mark saved document

5 **[0049]** The data area 295 and spare area 300 are typically set to specific sizes in both NOR and NAND flash devices. For example, each page 290 of the exemplary NAND flash device 270 of Figure 5 includes a data area 295 of 512 bytes and a spare area 300 of 16 bytes for a total page size of 528 bytes. The NAND flash device 270 also employs 32 pages 290 per block 280. Other page sizes may
10 be used in computer system 10 and are commercially available. For example, many NAND devices include blocks having 64 pages where each page stores 2112 bytes so that the total data area per page is 2048 bytes and the spare area per page is 64 bytes.

[0050] Flash memory devices, such as NAND flash device 270, typically
15 perform erase operations on an entire block 280 of memory at a time. An erase operation sets all bits within the block 280 to a consistent state, normally to a binary "1" value. Programming operations on an erased block 280 of flash device 270 can only change the contents of an entire page 290 (although NOR flash devices may be programmed in a slightly different manner). Once a page 290 of a
20 NAND flash device is programmed, its state cannot be changed further until the entire block 280 may be erased again. Reading of the contents of flash device 275 also occurs at the page level.

[0051] Figure 6 illustrates one manner in which transaction records may be organized in a flash memory device, such as NAND flash device 270. In this
25 example, each transaction record 310 may be comprised of two or more contiguous logical pages 315. Each logical page 315, in turn, may be comprised of two or more contiguous physical pages 290 of a block 280 of device 270. Meta-data information for the transaction record 310 may be stored in spare area 300, and may include some of the fields described in connection with header 230 of Figure 4. Depending
30 on the size of the spare area 300 of each page 290, the meta-data information may

be divided among multiple spare areas 300 of the transaction record 310. A division of the meta-data information between the spare areas 300 of two consecutive physical pages 290 is shown in Figure 6. The transaction records shown in Figure 6 also may be organized so that each transaction 310 corresponds to a single logical page 315 that, in turn, may be comprised of, for example, two contiguous physical pages 290.

[0052] An alternative arrangement in which there may be a one-to-one correspondence between each logical page 315 and a physical page 290 of flash device 270 is shown in Figure 7. A difference between this arrangement and the one shown in Figure 6 is that all of the meta-data information 320 may be stored in a single spare area 300 of the first physical page 290 of the transaction 310. Arrangements of this type may be particularly suitable when large capacity flash devices are employed. However, the meta-data information 320 also may be divided between the spare areas 300 of the two contiguous physical pages 290 of the transaction record.

[0053] The sequence identifiers for the transaction records 310 stored in the same device block 290 may have the same values. In such instances, the sequence identifier provides chronological information that may be used to compare the time relationship between the transaction records of different device blocks. Chronological information on the transaction records 310 stored in the same block can be derived from the offset location of the transaction record 310 within the block 290, with later occurring transaction records 310 occurring at larger offsets.

[0054] After the computer system 10 has been started or powered on, the integrity of the file system may be verified by generating a reconstructed version of the file system in random access memory 35. The reconstructed file system, shown generally at 330 of Figure 1, may be generated using the valid, committed transactions stored in the transaction file 220 and from the file/directory information stored in metafile 120. In Figure 1, the reconstructed file system 330 includes a regenerated file hierarchy 335 and an extents table 340.

[0055] One manner of generating the extents table 340 is shown in Figures 8 through 11. Figure 8 illustrates a number of interrelated processing steps that may be used to generate the extents pool 340 while Figures 9 through 11 illustrate the logical organization of various tables and arrays generated and used in these operations.

[0056] Generation of the extents table 340 may commence at step 345 of Figure 8 by scanning the blocks of the transaction file 220 to find all of the transaction records. The blocks may be scanned in sequence from the lowest ordered block to the highest ordered block in which a committed transaction record is found. As transactions are found within the blocks, an array of block records identifying each device block having a transaction record may be generated at step 350.

[0057] As the file system software 47 scans the blocks of the transaction file 220 four transactions, the file system software may encounter a block that has been erased as a result of transactions that have been retired, or because the blocks have not yet been assigned for use in the file system. The transaction header may be structured so that there are no valid transactions that will have all of the bits of the header set to the erased value, typically a binary "1". As the file system software 47 scans the blocks of the transaction file 220, any transaction in which the header indicates an erased block may be skipped. This header invariant may be enforced by using a single bit as a flag to indicate the transaction is in use by the file system when it is the inverse of the erase value. Upon finding such an erase signature value in a transaction header, scanning of the remaining pages in the block may be skipped thereby saving the time that would otherwise be used to access the erased pages. The overall system startup time may be correspondingly decreased.

[0058] The organization of an exemplary block array 355 is shown in Figure 9. Each block array record 360 includes a sequence field 365, a begin transaction field 370 and a number of transactions field 375. The sequence field 365 may be used to store the transaction identifier value for the transaction records stored in the

block. The begin transaction field 370 may be used to store an index to the first transaction in the block and the number of transactions field 375 may be used to store the number of transactions found in the block.

[0059] At step 380 of Figure 8, the file system software 47 populates a transaction list table for each record entry in the block array 355. Figure 9 illustrates one manner in which the transaction list table 385 may be organized. In this example, each record 360 of the block array 355 points to at least one transaction list record 390 of the transaction list table 385. More particularly, a transaction list record 390 may be generated for each transaction found in the block represented by a given block array record 360. The value stored in the number of transactions field 375 of the given block array record 360 corresponds to the number of transactions in the given block and designates how many records 390 for the given block will be added to transaction list table 385.

[0060] Each transaction list record 390 of the transaction list table 385 may have the same record length and include the same record fields. The exemplary fields used in records 390 of Figure 9 include a file cluster offset field 395, a device cluster index field 400, a number of clusters field 405 and a file identifier/idx field 410. The file cluster offset field 395 may be used to identify the physical location of the transaction within the block. The device cluster index field 400 may be used to identify where the data for the transaction begins. The number of clusters field 405 may be used to identify how many clusters of data are present within the transaction. Finally, the file identifier/idx field 410, as will be set forth below, is multipurpose. Initially, however, the value stored in the file identifier/idx field 410 may be used to identify the file to which the transaction applies. The file identifier value stored in field 410 may directly correspond to the file identifier used to reference the record in metafile 120. Upon the completion of step 380, the records 360 of block array 355 will be arranged, for example, in increasing block order, while the records 390 for each block array record 360 will be arranged in increasing page order.

[0061] At step 415, the records 360 of block array 355 are sorted based on the values stored in the sequence fields 365. This operation may be performed to place the records 390 of the transaction list table 385 in chronological order (*i.e.*, the order in which the corresponding transactions are to be applied to the files of the file system).

[0062] A temporary file 440 storing file node information corresponding to the transaction records of the file system may then be generated in RAM 35 using the sorted records of block array 355 and transaction list table 385. To this end, a basic record corresponding to the root directory of the file system may be first added to temporary file 440. The information used to generate the root directory node in temporary file 440 may be obtained from the record corresponding to the root directory file stored in metafile 120.

[0063] A logical representation of one manner of arranging the file node records in temporary file 440 is shown generally at 445 of Figure 10. In this example, each file node record 450 includes a file node field 455 and a start field 460. The contents of the file node field 455 may be used to identify the file node to which various transaction records 390 of the transaction list table 385 may be linked. For the sake of simplicity, the contents of the file node field 455 may have the same format as the file identifiers used to access the corresponding record entries 125 of metafile 120. The contents of the start field 460 may be used to identify the location of the first transaction record 390 in transaction list table 385 that corresponds to the file identified in the file node field 455. As such, each file node record 450 identifies a file within the file system as well as the location of the first transaction relating to the identified file.

[0064] At step 420, each of the sorted records 360 and 390 of the block array 355 and transaction list table 385 are traversed to determine whether or not the temporary file 440 includes a file node record 450 corresponding to the file identifier stored in file identifier/idx field 410. If a file node record 450 with the same file identifier as the transaction record 390 is not found in the temporary file 440 at step 425, a new file node record 450 may be created at step 430. Once a

file node record 450 corresponding to the transaction list record 390 exists in temporary file 440, the transaction list record 390 may be linked into a list of transactions for the file node record 450. In this example, the transaction list record 390 may be linked into the list of transactions for the file node record 450 at step 435 of Figure 8. The manner in which a transaction list record 390 may be linked into the list of transactions for the file node may depend on whether the transaction list record 390 may be the first transaction list record of the file node or a subsequent transaction list record for the file node. If it is the first transaction list record of the file node, the start field 460 of the file node record 450 may be updated to identify the starting location of this first transaction list record 390. As such, the contents of the start field 460 of the file node record 450 may be used to point to a location in the transaction list table 385 that, in turn, contains extent information for the first transaction applied to the file. The function of the file identifier/idx field 410 changes when the transaction list record 390 may be to be appended to existing transaction list records for the file node (*i.e.*, when it is not the first transaction list record for the file node). More particularly, the value and the function of the field 410 may be changed so that it points to the last transaction record 390 associated with the file node. This is illustrated in Figure 10, where the start field 460 of file node record 450 points to the beginning of transaction list record 390. The file identifier/idx field 410 of record 390, in turn, points to the beginning of transaction list record 465, which contains the information on the location of the second transaction for the file represented by the file node record 450. Similarly, the start field 460 of file node record 470 points to the beginning of transaction list record 475. The file identifier/idx field 410 of transaction list record 475 points to the beginning of transaction list record 480, which contains the information on the location of the second transaction for the file represented by the file node record 470.

[0065] Once all of the transaction list records of the transaction list table 385 have been linked in the proper manner with the corresponding file node records, the transaction list records for each file node are traversed at step 485 to remove any transaction list records that reference uncommitted and/or bad file

transactions. Removal of such transaction list records may be accomplished in a variety of different manners. For example, the file system software 47 may check the status field of the last occurring transaction to determine whether or not it was committed. If the transaction has been committed, the corresponding record in the transaction list table 385 may be left undisturbed. If the transaction has not been committed, however, the corresponding record in the transaction list table 385 may be removed or otherwise ignored.

[0066] To expedite this type of transaction commitment checking, the file system software 47 only needs to ensure that the last occurring transaction has been committed. Commitment checking of all other records may be skipped since only the last occurring transaction is impacted by a power failure, improper system shutdown, or the like. By skipping commitment checking of all other records, the time required for system startup may be substantially reduced.

[0067] Although it is shown as part of a linear sequence, step 485 may be executed as each transaction list record may be processed for incorporation in the corresponding file node. For example, file system software 47 may check the status information included in the header of each transaction record to determine whether the transaction has been committed. This check may occur as each transaction record may be used to populate the corresponding transaction list record. Once the file system software 47 finds a transaction that has not been committed, no further processing of the transaction list table 385 in steps 420 through 485 of Figure 8 is necessary.

[0068] At step 490, entries are generated in extents pool 340 for each of the file nodes. One manner in which this may be accomplished is shown in Figure 11. In this example, the content of the start field 460 of each file node may be changed so that it now operates as an extents index field 487. The extents index field 487 points to the first location in the extents pool 340 containing information on the location of the transaction data for the first transaction for the file. Each extents record 490 may include a number of clusters field 495, a start cluster field 500, and a next extent field 505. The start cluster field 500 identifies the starting location in

device 270 where the first file transaction for the file corresponding to the file node may be stored. The number of clusters field 495 identifies how many contiguous clusters of device 270 are used to store the file transaction. The next extents field 505 identifies the extents index of the next extents record for the file represented
5 by the file node. In this example, extents index 487 points to extents record 510 while the next extents field 505 of extents record 510 points to extents record 515.

[0069] The data used to populate the records of the extents pool 340 may be derived, at least in part, from the data stored in the transaction list table 385. In the example shown here, the extents pool 340 may be a more compact form of the
10 transaction list table 385. To this end, file system software 47 may combine transaction list records having contiguous data into a single extents record entry if the transaction list records are part of the same file node. Similarly, there is no further need to maintain the block array 355 in RAM 35. Therefore, block array 355 may be discarded from RAM 35.

15 **[0070]** The integrity of the transactions in the transaction file 220 may be checked during the execution of the various steps used to generate extents pool 340. For example, integrity checking of the transaction records may be executed during either steps 350 or 380 of Figure 8. Common data checks include CRC and ECC techniques.

20 **[0071]** To decrease the startup time of the computer system 10, error checking techniques may be limited to the information included in the header for certain transactions. As transactions are found during the startup process shown in Figure 8, the file system software 47 may identify whether the transaction impacts file data or metadata, such as directory structure information in metafile 120. This
25 distinction may be based on the file identifier associated with the transaction. Normally, metadata will be represented by file identifiers that are well-known and hard coded into the file system software 47 (*e.g.*, they will identify the metafile 120 as the file that is the subject of the transaction). Since only the metadata is required to ensure that the files system is in a consistent state after startup, data
30 checking techniques on the data portion of the transaction are only performed when

the transaction relates to such metadata. If the transaction does not relate to a change of the metadata, data checking techniques may be initially limited solely to the checking of the header information. In the transaction record format shown in Figure 6, the principal header information that must be verified on system startup
5 may be stored in the first spare area 300 of each transaction record 310. This allows the file system software 47 to skip verification of the header information included in the second spare area of each transaction record 310 thereby further optimizing the startup sequence. As will be explained in further detail below, error checking of the data portion of each transaction may be deferred until the time that
10 the corresponding file may be first accessed by the file system software 47 after completion of the startup sequence.

[0072] Any startup verification of the transaction records may be further optimized by limiting error checking solely to the first transaction header of a series of sequential transactions. During startup scanning of the transaction file 220, when
15 a transaction header is found that indicates that a number of sequential transaction records for the same file follow, verification of the headers of the trailing transactions in the sequence may be skipped once the header for the first transaction record of the sequence has been verified. Scanning and verification of header information may then resume with the next block following the last of the
20 trailing transactions.

[0073] The next broad step in generating the reconstructed file system 330 in RAM 35 may be the construction of the regenerated file hierarchy 335. In this example, the regenerated file hierarchy 335 may be comprised of both file and directory node records. An exemplary format for a directory node record is shown
25 generally at 520 of Figure 12 while a corresponding exemplary format for a file node record is shown generally at 525 of Figure 13.

[0074] Directory node record 520 includes a number of different fields that are used by the file system software 47. More particularly, directory node record 520 may include a sibling field 530, a file identifier field 535, a parent identifier field
30 540, a child field 545 and a directory named field 550. Similarly, file node record of

Figure 13 includes a number of different fields that are used by the file system software 47. The file node record fields may include a sibling field 555, a file identifier field 560, an extents index field 565 and a name sum field 570.

[0075] Since the data contained in the records of metafile 120 may be used in the construction of the regenerated file hierarchy 335, the manner in which the metafile records are arranged in the metafile 120 will have an impact on the system startup performance. To this end, the records of metafile 120 are arranged in a single metafile as contiguous records having the same length and are all stored in the same storage media. This arrangement enhances the speed with which the file system software 47 may access the metafile data and reduces the amount of processing that is required for such access.

[0076] One sequence of steps that may be used to populate the fields for each file node record 525 and directory node record 520 of the regenerated file hierarchy 335 is shown in Figure 14. The illustrated sequence may be executed for each record in metafile 120 and may start at step 575. At step 575, a file identifier may be generated based on the offset of the first record entry within the metafile 120. A check of the regenerated file hierarchy 335 may be made at step 580 to determine whether a file node record 525 or directory node record 520 corresponding to the file identifier is already present. If a corresponding record 520 or 525 is not present, a new record file may be created in the regenerated file hierarchy 335. The format of the newly created record depends on whether the file identifier corresponds to a file entry or directory entry in metafile 120. The file system software 47 will make this determination and apply the proper record format 520 or 525.

[0077] At step 585, the fields for the newly created record are populated using the attributes for the file/directory that are found in the metafile 120. If the newly created record corresponds to a directory node, the parent identifier field 540 and directory name field 550 are populated using the data in the parent file identifier and short name fields of the corresponding record in metafile 120. If the newly created record corresponds to a file node, the name sum field 570 may be

populated using data that is directly stored or derived from the file name data of the corresponding record in metafile 120. The extents index field 565 may be populated using the data found in the extents index field 487 of the corresponding file node record 450 (see Figure 11).

- 5 **[0078]** If the newly created file corresponds to a directory node, a search through the regenerated file hierarchy 335 may be undertaken at step 590 to determine whether the parent node exists. If the parent node does not exist, a directory record corresponding to the parent node may be added to the regenerated file hierarchy 335.
- 10 **[0079]** At step 595, the newly generated file/directory record may be linked into the tree structure for the parent directory node. If the child field 545 of the newly generated file/directory record indicates that the parent directory has no children, the value of the child field 545 of the parent directory record may be reset to point to the newly generated file/directory record and the sibling field 555 or 530
 15 of the newly generated file/directory record may be set to indicate that the newly generated file/directory record does not have any siblings. If the child field 545 of the parent node record indicates that the parent directory node has children, the sibling field 565 or 530 of the newly generated file/directory record may be set to point to the existing child of the parent directory and the child field 545 of the
 20 parent directory may be set to point to the newly generated file/directory record. If the newly generated file/directory record corresponds to a directory node, the parent identifier field 540 of the newly generated directory record may be set to point to the parent directory node.
- [0080]** At step 600, the file system software 47 recursively ascends the
 25 parent nodes, beginning with the parent directory of the newly generated file/directory record, and executes a series of processing steps until the root node is reached. At this point, the parent directory node of the newly generated file/directory record may be referred to as the current directory node. In the exemplary process shown in Figure 14, the file system software 47 checks the
 30 regenerated file hierarchy 335 to determine whether a directory node record

corresponding to the parent node of the current directory exists. This process may be executed at steps 605 and 610. If such a directory record does not exist in the regenerated file hierarchy 335, a new directory record may be generated at step 615. The child field 545 of the newly generated directory record may be then set to point to the current directory node record as the only child of the new directory record. At step 620, the parent identifier field 540 of the current directory node record may be set to point to the newly generated directory record. The sibling field 530 of the current directory node record may be set to indicate that there are no siblings for the current directory node record at step 625.

10 **[0081]** If the check executed at steps 605 and 610 indicate that there is a directory record in the regenerated file hierarchy 335 that corresponds to parent node of the current directory, then the current directory node may be linked into the generalized tree structure of the parent directory node at step 630. To this end, the parent identifier field 540 of the current node may be set to point to the location of the parent node record in the regenerated file hierarchy 335. The sibling field 530 of the current directory node may be set to point to the same record as pointed to by the child field 545 of the parent node record. Finally, the child field 545 of the parent directory node may be set to point to the location of the current directory node.

20 **[0082]** At step 635, the file system software 47 checks to determine whether the recursive directory processing is completed. In this example, the recursive directory processing is completed when the processing a sends to the root node, which has a unique and recognizable file identifier. If the root node has been reached at step 635, processing of the next file record entry in metafile 120 may be begun at step 640, which returns control of the processing back to step 575. If the root node has not been reached at step 635, then processing of the next parent node in the ascending file/directory hierarchy may be repeated beginning at step 605.

[0083] Figure 15 is a logical representation of the reconstructed file system 330 and corresponds to the application of the processing steps of Figures 8 and 14

to a file system having the file hierarchy shown in Figure 2. In this exemplary representation, lines 665, 670, 675, and 680 represent pointers that correspond to the content of the parent identifier fields 540 for the directory node records representing directories 105, 100, 80 and 85, respectively. Lines 645, 650, 660, 5 655 and 652 represent pointers that correspond to the content of the child identifier fields 545 for the directory node records representing directories 110, 100, 105, 80 and 85, respectively. Lines 685, 690, 695 and 705 represent pointers that correspond to the content of the sibling identifier fields 530 for the directory node records corresponding directories 100, 105 and 80, respectively. Lines 700, 705, 10 710 and 715 represent pointers that correspond to the content of the sibling identifier fields 555 for the file node records corresponding to files 90, 55, 60 and 70, respectively.

[0084] One manner of accessing data in the transaction file 220 of persistent storage unit 30 using the reconstructed file system 330 is also illustrated in Figure 15. As shown, the file system software 47 provides a file identifier 730 for the file node record that the software is to access. In this example, the file identifier 730 points to the file node record representing file 55. The file system software 47 then uses the contents of the extents index 565 of the file node record as an index into extents pool 340 to locate the data for the file in the transaction file 220. It will be 20 recognized, however, that the file system software 47 may use the contents of the reconstructed file system 330 in a variety of different manners other than the one illustrated in Figure 15.

[0085] As noted above, complete verification of the integrity of a file is not performed during startup so that startup processing may be expedited. Instead, 25 the file system software 47 may defer complete verification of the file until the first time that the file may be accessed. To this end, the file system software 47 may maintain a table indicating whether or not the integrity of each file has been completely verified. Alternatively, the file system software 47 may use one or more bits of each file node record in the regenerated file hierarchy 335 to indicate 30 whether the integrity of the file has been completely verified. This indicator may be checked by the file system software 47 at least the first time that a file may be

accessed after startup. If the indicator shows that the file has not been completely verified, a complete verification of the file may be executed at that time.

Alternatively, since the headers of the transactions for the file have already been checked, the file system software need only verify the integrity of the data portions
5 of each transaction for the file. The verification processes may include one or more CRC processes, one or more ECC processes, and the like.

[0086] As shown in Figures 5, 6 and 7, a number of different fields in each of the transaction record headers may be dedicated to verifying the integrity of the entire transaction record. If the integrity checks fail and an application using the
10 relevant error-correcting codes cannot correct the error, then a program error may be reported back to the application or system that made the request to access the file contents.

[0087] While various embodiments of the invention have been described, it will be apparent to those of ordinary skill in the art that many more embodiments
15 and implementations are possible within the scope of the invention. Accordingly, the invention is not to be restricted except in light of the attached claims and their equivalents.

Claims:

1. A computer-implemented method comprising:

providing a file system structure of files on a persistent data storage device;

maintaining a transaction file on the persistent data storage device;

including a plurality of transaction records in the transaction file with a processor, each of the transaction records representing a file transaction that affects at least one of the files stored on the persistent data storage device; and

generating a reconstructed file system in a random access memory with the processor in a startup process from the transaction records and the file system structure of the files, wherein generating the reconstructed file system comprises generating a regenerated file hierarchy in the random access memory that includes a file node record for each file affected by the file transactions represented by the transaction records, wherein the file node record for each affected file identifies a corresponding physical location in the persistent data storage device that includes a transaction record representing at least one of the file transactions that affects a file identified by the file node record; and

error checking a data portion of the transaction record representing the at least one of the file transactions that affect the file identified by the file node record, the transaction record comprising the data portion and a metadata portion, wherein the transaction record is identified by the file node record and the data portion is error checked when the file identified by the file node record is first accessed after completion of the startup process.

2. The method of claim 1 wherein generating the reconstructed file system further comprises generating a plurality of transaction list records in the random access memory corresponding to the transaction records stored in the persistent data storage device, each of the transaction list records identifying a physical location in the persistent data storage device at which a corresponding one of the transaction records is stored.

3. The method of claim 2 wherein generating the reconstructed file system further comprises sorting the transaction list records in the random access memory in an order in which the corresponding file transactions are to be applied to the files.

4. The method of claim 3 wherein generating the reconstructed file system further comprises linking the transaction list records together that are related to a respective one of the affected files.

5. The method of claim 4 wherein error checking further comprises determining, for each of the affected files, whether the last occurring file transaction for the affected file was committed, and, in response to a determination that the last occurring file transaction was committed, removing or ignoring the corresponding transaction list record without checking whether the other file transactions related to the affected file were committed.

6. The method of claim 1 wherein generating the reconstructed file system further comprises generating a plurality of block records in the random access memory from a scan of the transaction records, each one of the block records identifying a corresponding device block of the persistent data storage device that includes at least one of the transaction records.

7. The method of claim 6 wherein generating the reconstructed file system further comprises generating, for each one of the block records, at least one transaction list record in the random access memory for each transaction record stored in the corresponding device block of the persistent data storage device, each of the at least one transaction list record identifying a physical location within the corresponding device block of the persistent data storage device at which each transaction record is stored, the at least one transaction list record included in a plurality of transaction list records.

8. The method of claim 7 wherein generating the regenerated file hierarchy further comprises generating the file node record in the random access memory for each file affected by the file transactions identified by the transaction list records,

each file node record identifying an affected file within the file system, each file node record identifying a first of the transaction list records that relates to the affected file, the file node record identifying the corresponding physical location in the persistent data storage device that includes the transaction record corresponding to the first of the transaction list records through inclusion of a pointer in the file node record to the first of the transaction list records.

9. An apparatus comprising:

a persistent data storage device comprising a file system structure of files and a transaction file on the persistent data storage device, the transaction file comprising including a plurality of transaction records, each of the transaction records representing a file transaction that affects at least one of the files stored on the persistent data storage device;

a random access memory; and

a processor configured to generate a reconstructed file system in the random access memory in a startup process from the transaction records and the file system structure of the files, wherein the reconstructed file system in the random access memory comprises a regenerated file hierarchy that includes a file node record for each file affected by the file transactions represented in the transaction records, wherein the file node record for each affected file identifies a corresponding physical location in the persistent data storage device that includes a transaction record representing at least one of the file transactions that affects a file identified by the file node record, wherein the processor is further configured to error check a data portion of the transaction record representing the at least one of the file transactions affecting the file identified by the file node record, the transaction record comprising a metadata portion and the data portion, wherein the transaction record is identified by the file node record and the data portion is error checked in response to a request to access the file identified by the file node record after the startup process is completed.

10. A non-transitory machine readable medium having tangibly stored thereon executable instructions that, when executed by a processor, cause the processor to perform the method of any one of claims 1-8.

11. An apparatus, comprising:

a processor;

a memory coupled to the processor, the memory storing executable instructions that, when executed by the processor, cause the processor, to perform the method of any one of claims 1-8.

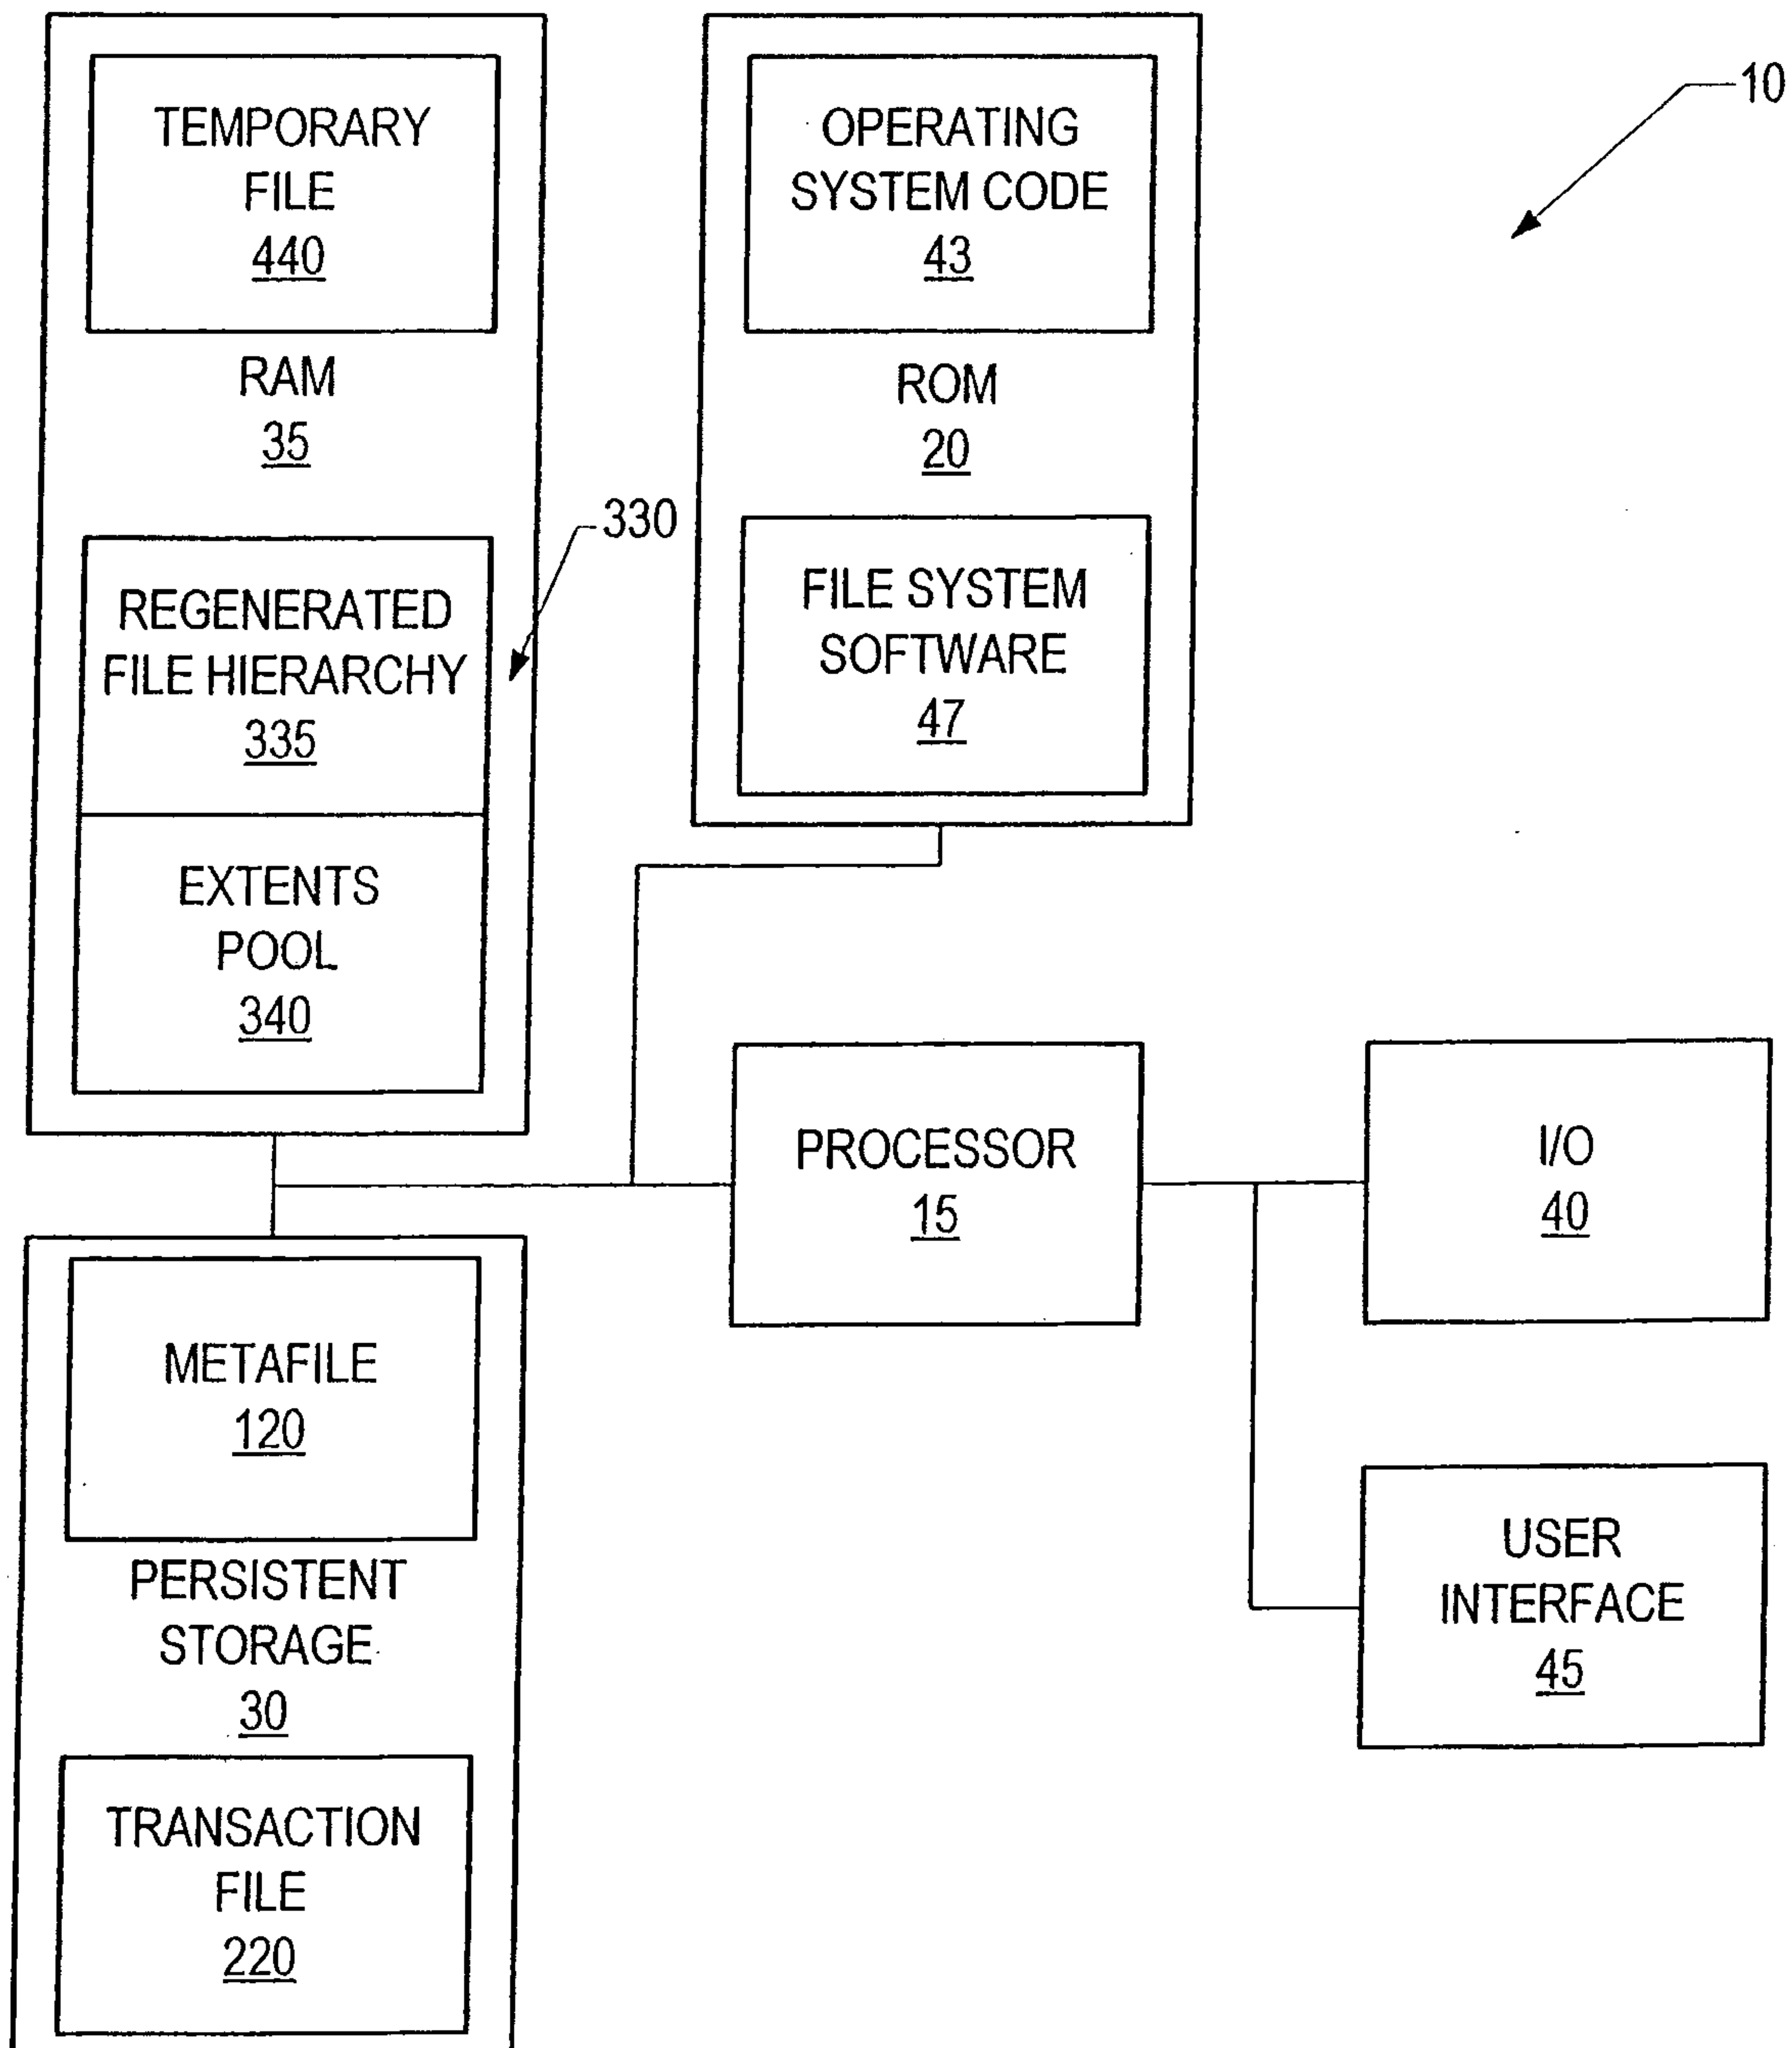


Fig. 1

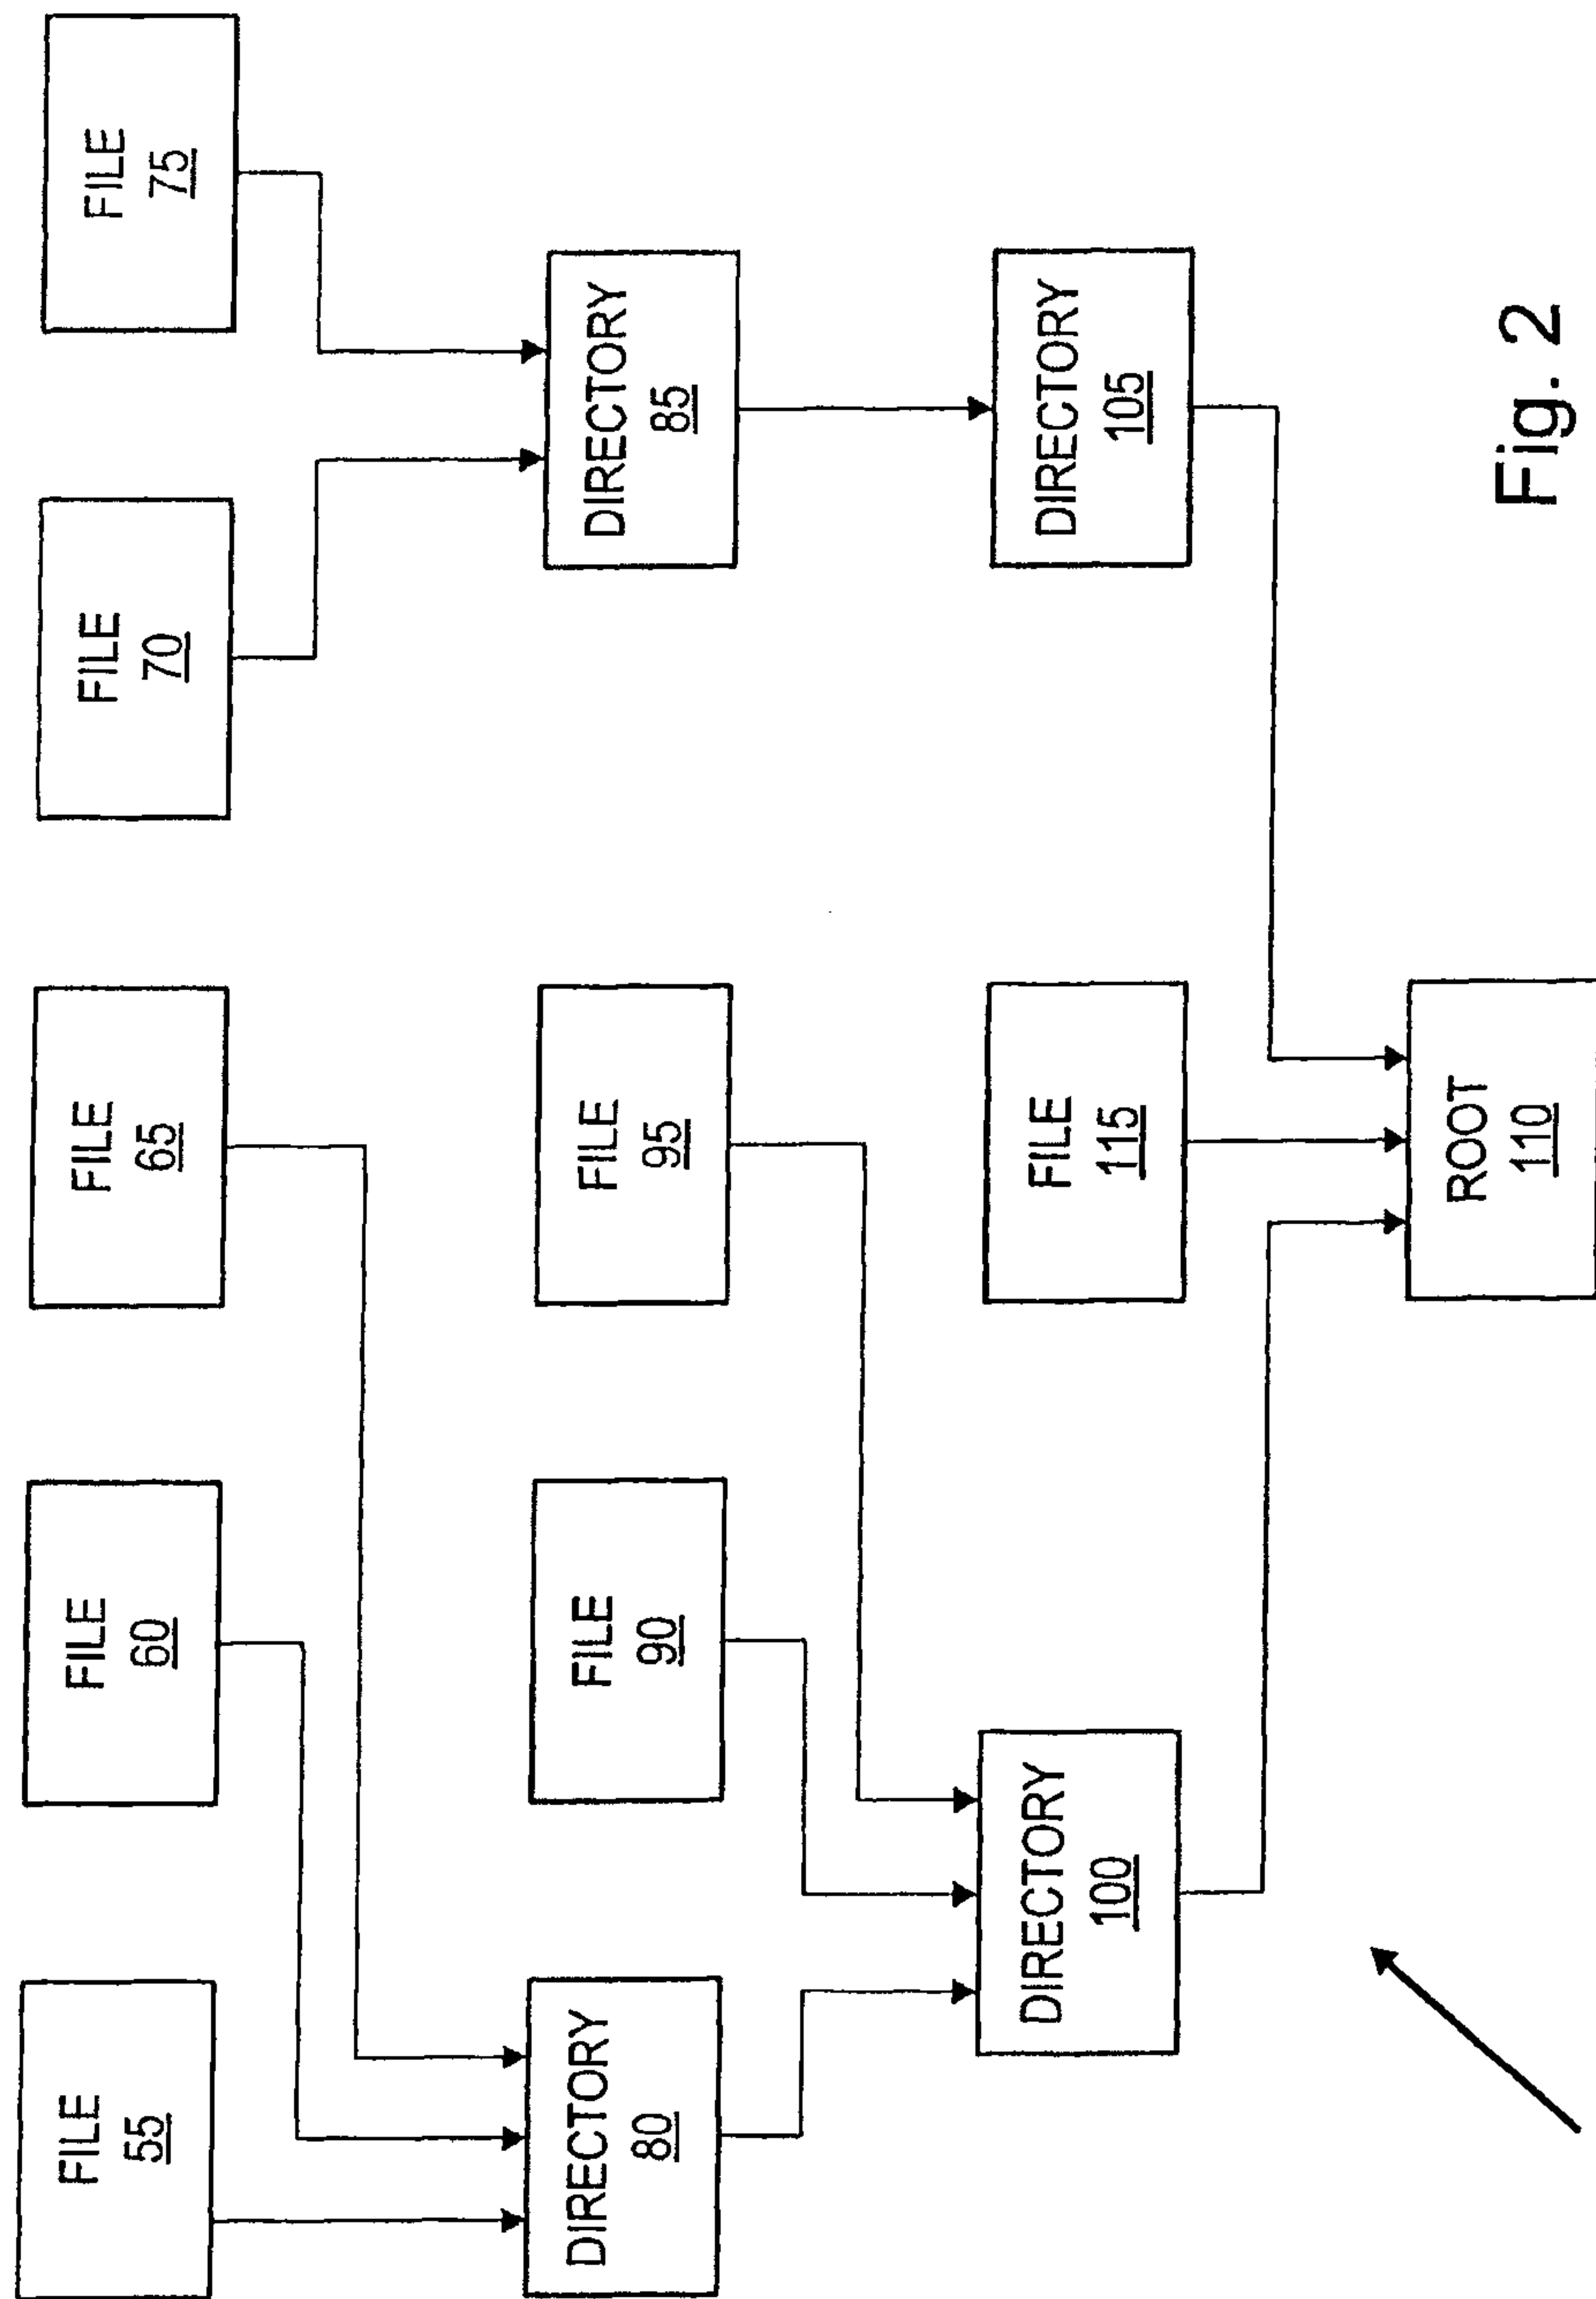


Fig. 2

50




Fig. 3




Fig. 4




Fig. 5




Fig. 6




Fig. 7

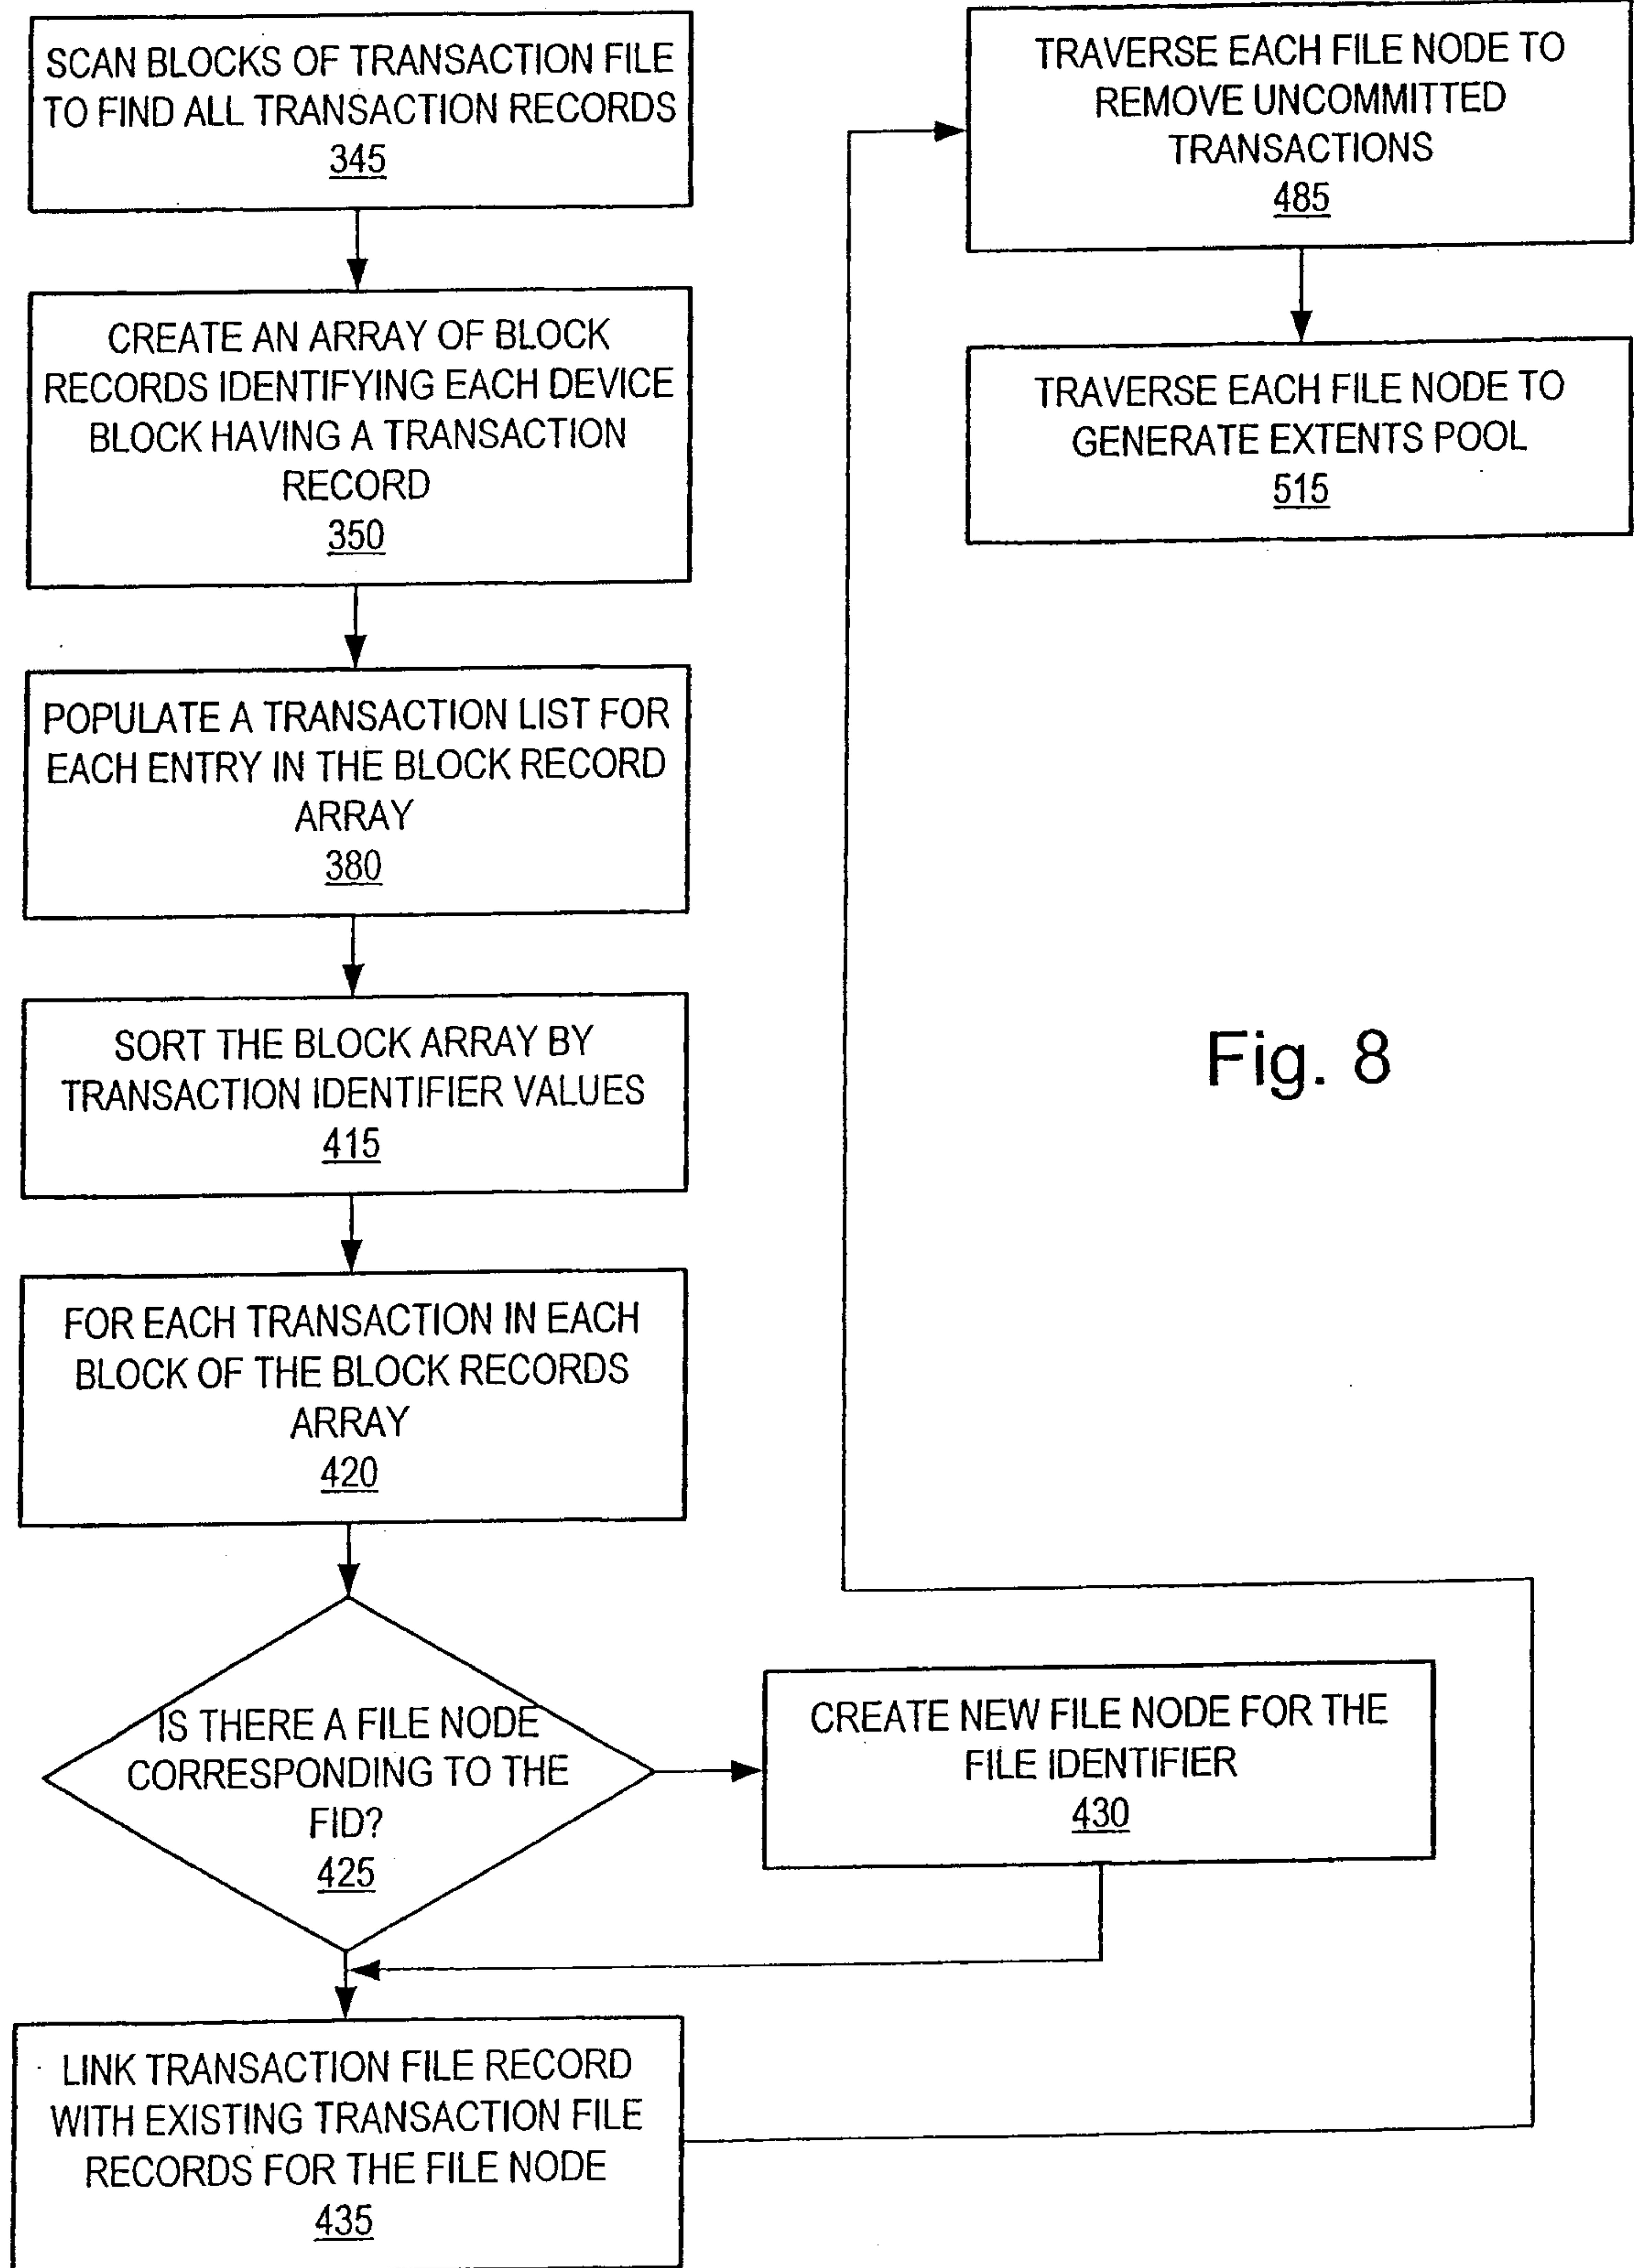


Fig. 8




Fig. 9




Fig. 10




Fig. 11

DIRECTORY NODE



SIBLING <u>530</u>
FID <u>535</u>
PARENT ID <u>540</u>
CHILD <u>545</u>
DIRECTORY NAME <u>550</u>

Fig. 12




Fig. 13

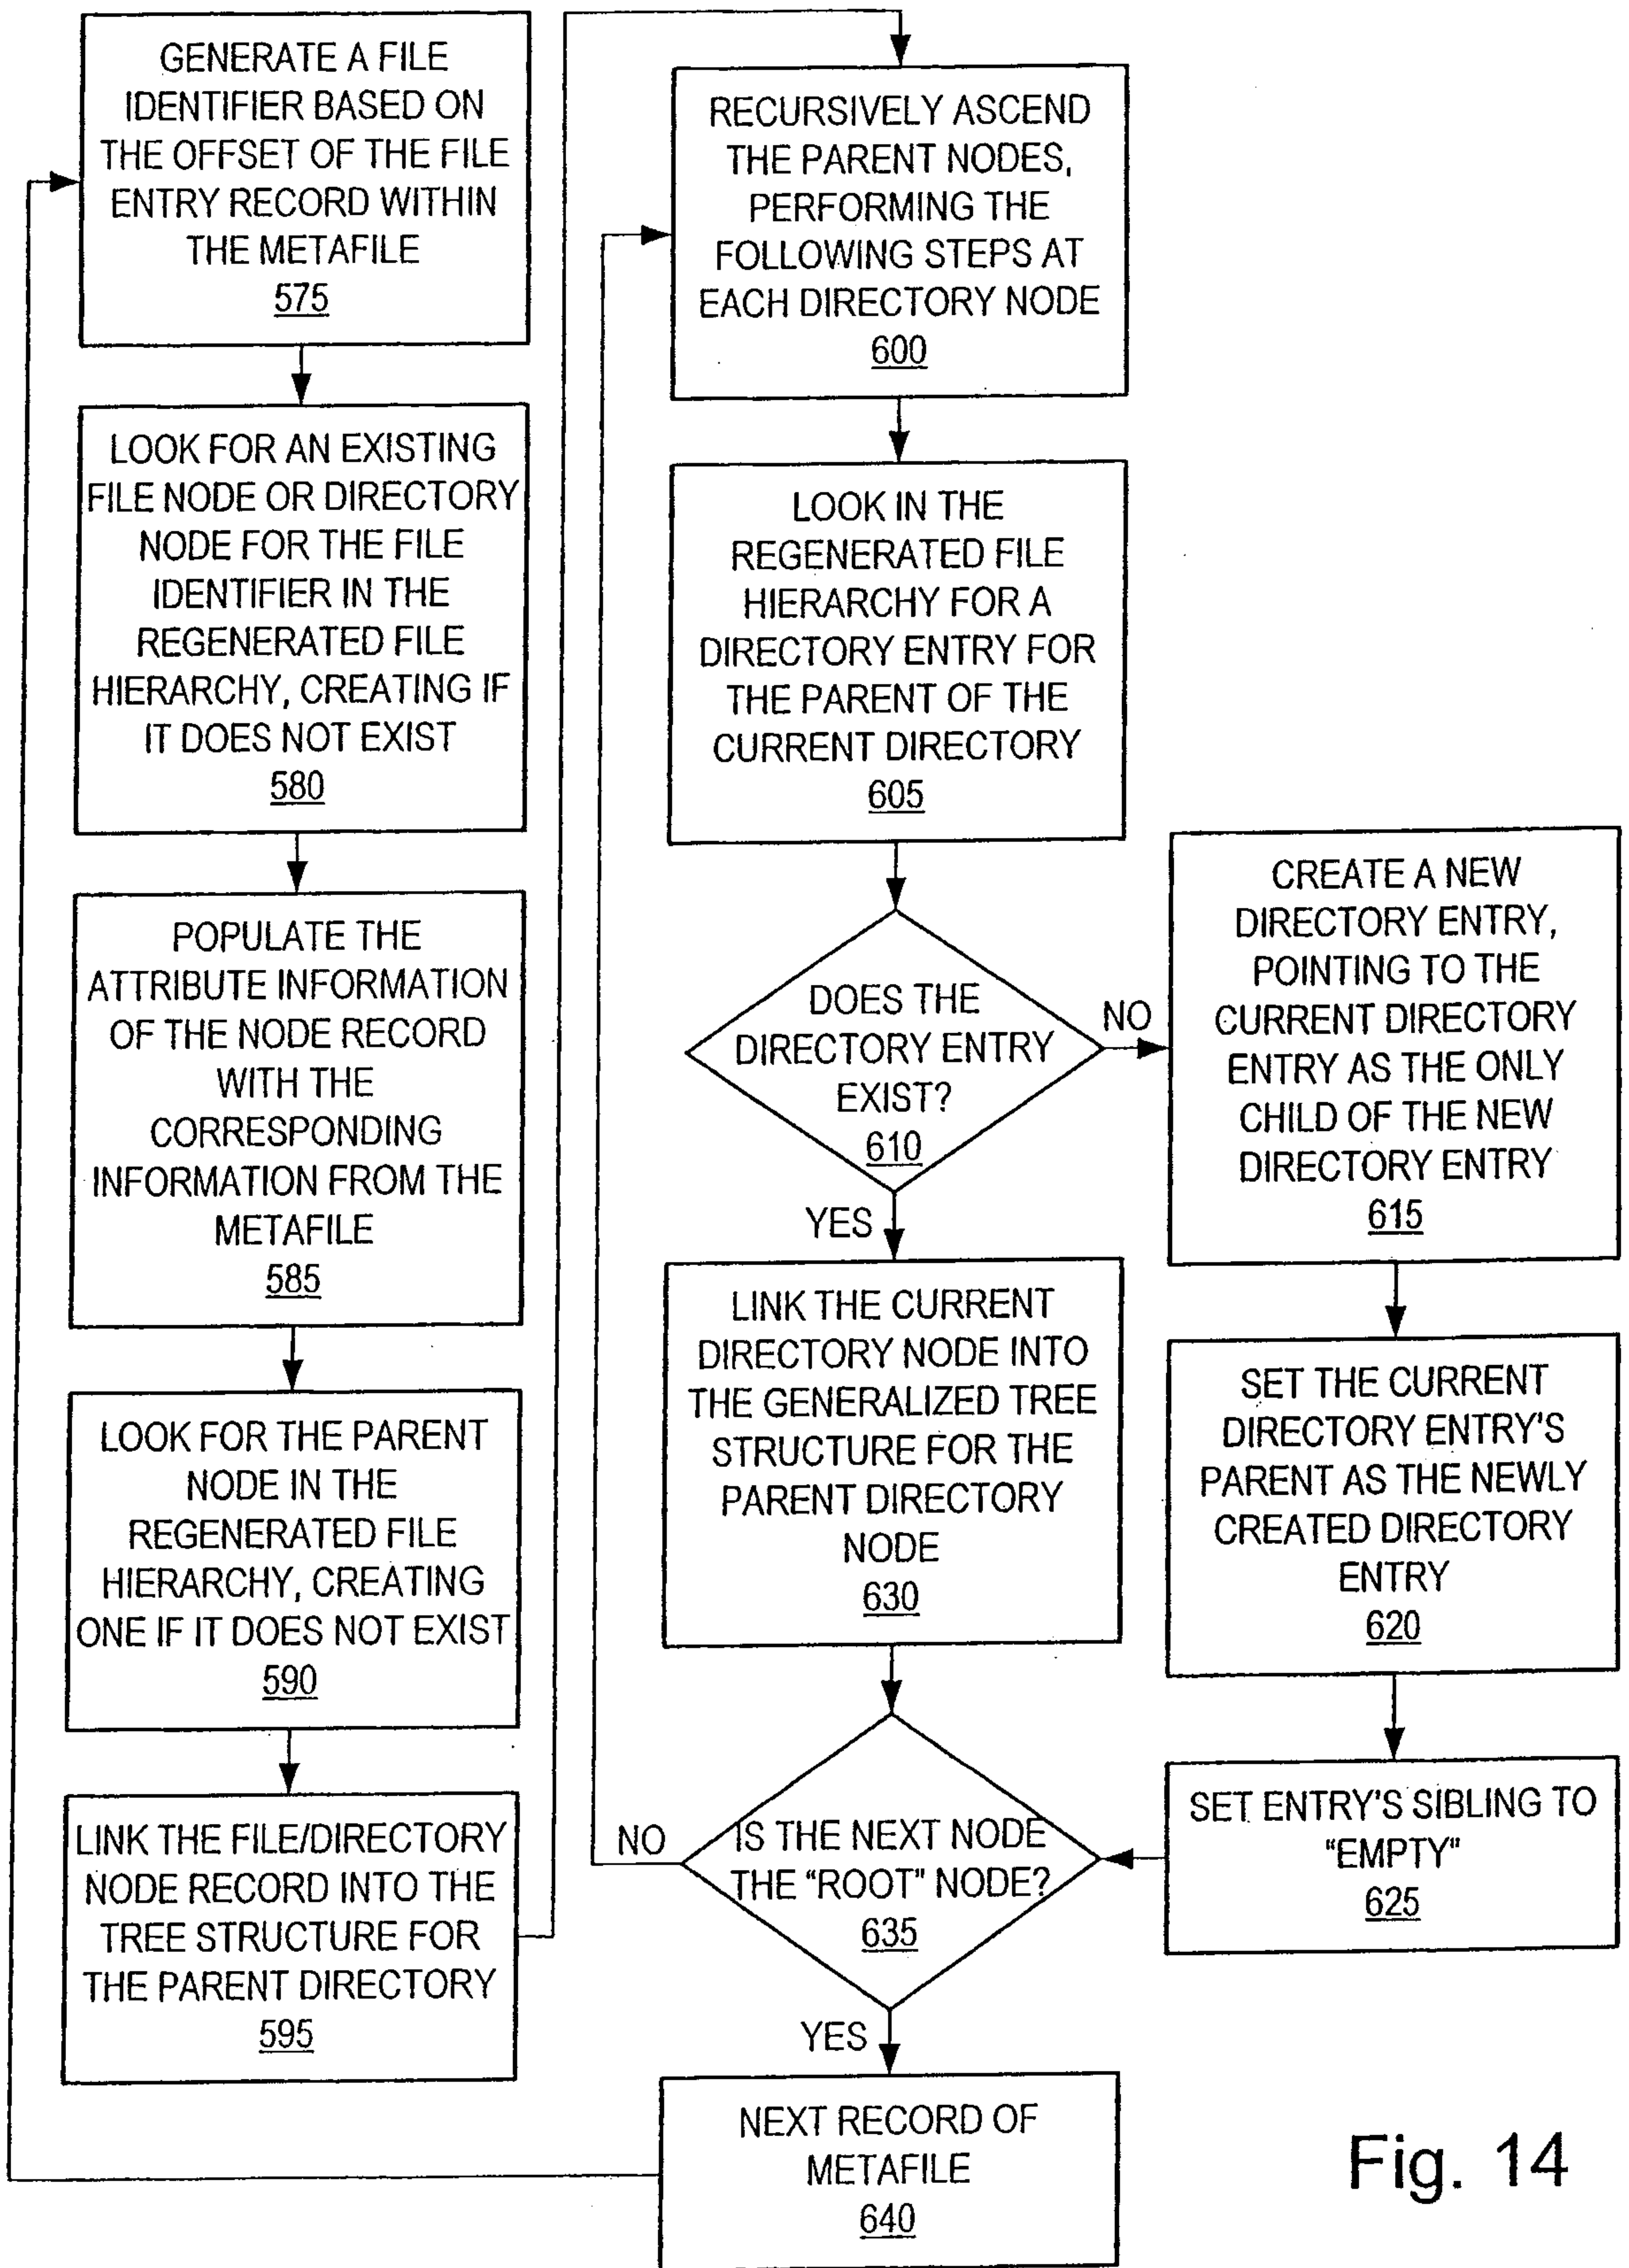


Fig. 14

Fig. 15

