

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
6 September 2002 (06.09.2002)

PCT

(10) International Publication Number  
**WO 02/069200 A1**

(51) International Patent Classification<sup>7</sup>: **G06F 17/30**

(21) International Application Number: PCT/US02/05613

(22) International Filing Date: 25 February 2002 (25.02.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
60/271,124 24 February 2001 (24.02.2001) US

(71) Applicants and

(72) Inventors: **GARA, Alan, G.** [US/US]; 38 Marion Avenue, Mount Kisco, NY 10549 (US). **GIAMPAPA, Mark, E.** [US/US]; 140 North Broadway, Apt. GT-14, Irvington, NA 10533 (US). **STEINMACHER-BUROW, Burkhard, D.** [CA/US]; 28 Manchester Drive, Mount Kisco, NY 10549 (US).

(74) Agents: **GROLZ, Edward** et al.; Scully, Scott, Murphy & Presser, 400 Garden City Plaza, Garden City, NY 11530 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.

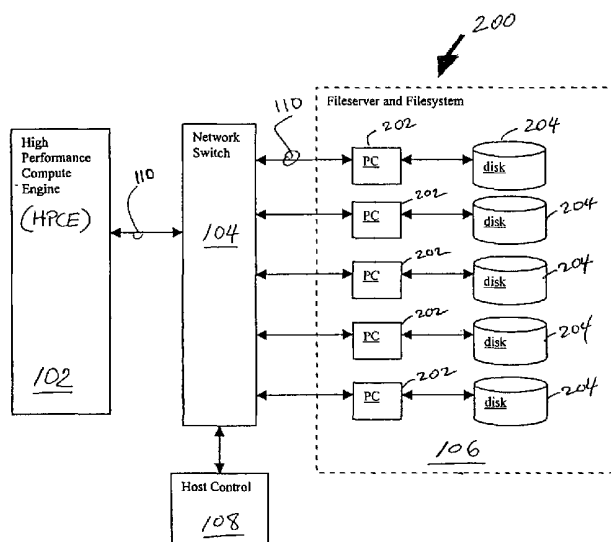
(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

**Published:**

- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

[Continued on next page]

(54) Title: CHECKPOINTING FILESYSTEM



(57) **Abstract:** The present invention is directed to a checkpointing filesystem of a distributed memory parallel supercomputer (102) comprising a node (104) that accesses user data (110) on the filesystem (200) the filesystem (110) comprising an interface (202) that is associated with a disk (204) for storing the user data. The checkpointing filesystem provides for taking and checkpoint of the filesystem and rolling back to a previously taken checkpoint, as well as for writing user data to and deleting user data from the checkpointing filesystem. The checkpointing filesystem provides a recently written file allocation table (WFAT) for maintaining information regarding the user data written since a previously taken checkpoint and a recently deleted file allocation table (DFAT) for maintaining information regarding user data deleted from since the previously taken checkpoint, both of which are utilized by the checkpointing filesystem to take a checkpoint of the filesystem and rollback the filesystem to a previously taken checkpoint, as well as to write and delete user data from the checkpointing filesystem.



WO 02/069200 A1



---

*For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

**CHECKPOINTING FILESYSTEM**

5

**CROSS REFERENCE**

The present invention claims the benefit of commonly-owned, co-pending United States Provisional Patent Application Serial Number 60/271,124 filed February 24, 2001 entitled MASSIVELY PARALLEL SUPERCOMPUTER, the whole contents and disclosure of which is expressly incorporated by reference herein as if fully set forth herein. This patent application is additionally related to the following commonly-owned, co-pending United States Patent Applications filed on even date herewith, the entire contents and disclosure of each of which is expressly incorporated by reference herein as if fully set forth herein. U.S. patent application Serial No. (YOR920020027US1, YOR920020044US1 (15270)), for "Class Networking Routing"; U.S. patent application Serial No. (YOR920020028US1 (15271)), for "A Global Tree Network for Computing Structures"; U.S. patent application Serial No. (YOR920020029US1 (15272)), for 'Global Interrupt and Barrier Networks'; U.S. patent application Serial No. (YOR920020030US1 (15273)), for 'Optimized Scalable Network Switch'; U.S. patent application Serial No. (YOR920020031US1, YOR920020032US1 (15258)), for "Arithmetic Functions in Torus and Tree Networks"; U.S. patent application Serial No. (YOR920020033US1, YOR920020034US1 (15259)), for 'Data Capture Technique for High Speed Signaling'; U.S. patent application Serial No. (YOR920020035US1 (15260)), for 'Managing Coherence Via Put/Get Windows'; U.S. patent application Serial No. (YOR920020036US1, YOR920020037US1 (15261)), for "Low Latency Memory Access And Synchronization"; U.S. patent application Serial No. (YOR920020038US1 (15276), for 'Twin-Tailed Fail-Over for Fileservers Maintaining Full Performance in the Presence of Failure'; U.S. patent application Serial No. (YOR920020039US1 (15277)), for "Fault Isolation Through No-Overhead Link Level Checksums"; U.S. patent application Serial No. (YOR920020040US1 (15278)), for "Ethernet Addressing Via Physical Location for Massively Parallel Systems"; U.S. patent application Serial No. (YOR920020041US1 (15274)), for "Fault Tolerance in a Supercomputer Through

Dynamic Repartitioning”; U.S. patent application Serial No. (YOR920020042US1 (15279)), for “Checkpointing Filesystem”; U.S. patent application Serial No. (YOR920020043US1 (15262)), for “Efficient Implementation of Multidimensional Fast Fourier Transform on a Distributed-Memory Parallel Multi-Node Computer”;  
5 U.S. patent application Serial No. (YOR9-20010211US2 (15275)), for “A Novel Massively Parallel Supercomputer”; and U.S. patent application Serial No. (YOR920020045US1 (15263)), for “Smart Fan Modules and System”.

## **BACKGROUND OF THE INVENTION**

10

### **Technical Field of the Invention**

The present invention generally relates to a field of filesystems in computers. More particularly, the present invention is directed to efficiently checkpointing a  
15 filesystem on a distributed-memory parallel supercomputer, thereby facilitating faster execution of applications.

### **Description of the Prior Art**

20 In large computing systems, such as a distributed-memory parallel supercomputer, it is standard to save a state of a system at regular intervals such that a application can be rolled back and rerun from a last saved state of the system, thereby saving time and computing resources. This is necessary because the large computing systems do not have the reliability of small computing systems and the applications  
25 that utilize the large computing systems often run for hours, days or weeks. More particularly, the large computing systems may crash or be brought down for maintenance, exceptions may be encountered while an application executes, or programmer-defined conditions may be met which terminate the application. Since the application manipulates disk files, the roll back restores the manipulated  
30 disk files to a previous clean state, i.e., to a previous checkpoint. Therefore, checkpointing the application mitigates the rerunning of the application since the application need only be rolled back to the previous checkpoint rather than be rerun from the start. Thus, checkpointing of the filesystem is a critical aspect of

checkpointing in the large computing systems, such as the distributed-memory parallel supercomputer.

5 A main motivation for a distributed-memory parallel supercomputer is a fast execution of the application. Thus, during the execution of an application on the distributed-memory parallel supercomputer, there is a need for a checkpointing filesystem, which is not significantly slower than a filesystem without checkpointing. Similarly, there is a need for a checkpointing filesystem in which the act of checkpointing is fast, since during checkpointing of the filesystem the application is not executing. There also is a need for a checkpointing filesystem, which appears to the application as a normal filesystem and does not complicate the implementation of the application.

10 Therefore, there is a need in the art for providing a checkpointing filesystem on the distributed-memory parallel supercomputer that facilitates faster execution of an application executing on the distributed-memory parallel supercomputer.

### **SUMMARY OF THE INVENTION**

20 It is therefore an object of the present invention to provide a checkpointing filesystem that facilitates fast execution of an application on the distributed-memory parallel supercomputer.

25 It is another object of the present invention to provide a checkpointing filesystem, which is not significantly slower than a filesystem without checkpointing on the distributed-memory parallel supercomputer.

30 It is a yet another object of the present invention to provide a checkpointing filesystem, which appears to the application as a conventional filesystem and does significantly effect an application on the distributed-memory parallel supercomputer.

It is a further object of the present invention to provide a checkpointing filesystem in which the taking of a checkpoint does not significantly effect execution of an application on the distributed-memory parallel supercomputer.

- 5 It is yet a further object of the present invention to provide a checkpointing filesystem in which the rolling back of a previously taken checkpoint is efficient as possible, thereby not significantly effecting the application on the on the distributed-memory. parallel supercomputer.
- 10 According to an embodiment of the present invention, there is provided a method for checkpointing a filesystem of a distributed-memory parallel supercomputer comprising a node that accesses user data on the filesystem, the filesystem comprising an interface that is associated with a disk for storing the user data, the method comprising the steps of: informing the node and the interface to complete  
15 their access to the filesystem in order to take a checkpoint of the filesystem; directing the interface to take the checkpoint of the associated disk, the taking of the checkpoint comprising a step of clearing: (i) a recently written file allocation table (WFAT) for maintaining information regarding the user data written to the disk since a previously taken checkpoint; (ii) a recently deleted file allocation table  
20 (DFAT) for maintaining information regarding the user data deleted from the disk since the previously taken checkpoint; and informing the node and the interface that they may resume their access to the filesystem, wherein the WFAT and DFAT are utilized to rollback the filesystem to a clean state.
- 25 According to another embodiment of the present invention, there is provided a method for rolling back a checkpointing filesystem of a distributed-memory parallel supercomputer comprising a node that accesses user data on the filesystem to a previously taken checkpoint of the filesystem, the filesystem comprising an interface that is associated with a disk for storing the user data, the method  
30 comprising the steps of: informing the node and the interface to complete their access to the filesystem in order rollback to the previously taken checkpoint of the filesystem; directing the interface to rollback to the previously taken checkpoint of the associated disk, the rolling back to the previously taken checkpoint comprising:

- (i) removing user data from a file allocation table (FAT) utilizing a recently written file allocation table (WFAT) that maintains information regarding the user data written to the disk since the previously taken checkpoint; (ii) adding user data to the FAT from a recently deleted file allocation table (DFAT) that
- 5 maintains information regarding user data deleted from the disk since the previously taken checkpoint; and informing the node and the interface that they may resume their access to the filesystem, wherein the WFAT and DFAT are utilized to rollback the filesystem to a clean state.
- 10 According to yet another embodiment of the present invention, there is provided a method for writing user data to a checkpointing filesystem of a distributed-memory parallel supercomputer comprising a node that accesses the user data on the filesystem, the filesystem comprising an interface that is associated with a disk for storing the user data, the method comprising the steps of: selecting sectors that are
- 15 marked as free in a file allocation table (FAT) and not marked as deleted in a recently deleted file allocation table (DFAT) that maintains information regarding user data deleted from the disk since a previously taken checkpoint; writing the user data to the selected sectors on the disk; updating the FAT to reflect the written user data; and updating a recently written file allocation table (WFAT) that
- 20 maintains information regarding the user data written to the disk since the previously taken checkpoint, wherein the WFAT and DFAT are utilized to rollback the filesystem to a clean state.

- According to a further embodiment of the present invention, there is provided a
- 25 method for deleting user data from a checkpointing filesystem of a distributed-memory parallel supercomputer comprising a node that accesses the user data on the filesystem, the filesystem comprising an interface that is associated with a disk for storing the user data, the method comprising the steps of: deleting sectors associated with the user data to be deleted from the disk from a file allocation table
- 30 (FAT); determining whether the sectors to be deleted are marked in a recently written file allocation table (WFAT) that maintains information regarding user data written to the disk since a previously taken checkpoint; and updating a recently deleted file allocation table (DFAT) if the sectors are not marked in WFAT at step

(b), the DFAT maintaining information regarding the user data deleted from the disk since the previously taken checkpoint, wherein the WFAT and DFAT are utilized to rollback the filesystem to a clean state.

- 5 According to yet a further embodiment, there is provided a checkpointing filesystem of a distributed-memory parallel supercomputer comprising a node that accesses user data on the filesystem, the filesystem comprising an interface that is associated with a disk for storing the user data, comprising: a host controller comprising means for informing the node and the interface to complete their access
- 10 to the filesystem in order to take a checkpoint of the filesystem, directing the interface to take a checkpoint of the associated disk, and informing the node and the interface that they may resume their access to the filesystem after the checkpoint it taken; and the interface comprising a means for clearing: a recently written file allocation table (WFAT) that maintains information regarding the user
- 15 data written to the disk since a previously taken checkpoint; and a recently deleted file allocation table (DFAT) for maintaining information regarding the user data deleted from the disk since the previously taken checkpoint, wherein the WFAT and DFAT are utilized to rollback the filesystem to a clean state.
- 20 The checkpointing filesystem of the distributed-memory parallel supercomputer further comprises rolling back to the previously taken checkpoint, wherein: the host controller further comprises means for informing the node and the interface to complete their access to the filesystem in order to rollback to the previously taken checkpoint of the filesystem, for directing the interface to rollback to the
- 25 previously taken checkpoint of the associated disk, and for informing the node and the interface that they may resume their access to the filesystem; and the interface further comprises means for rolling back to the previously taken checkpoint, the means comprising removing user data from a file allocation table (FAT) utilizing the WFAT and adding user data to the FAT from DFAT.
- 30 The checkpointing filesystem of the distributed-memory parallel supercomputer further comprising writing user data to the checkpointing filesystem, wherein: the interface further comprises means for selecting sectors that are marked as free in

the FAT and not marked for deletion in DFAT, for writing the user data to the selected sectors on the disk; for updating the FAT to reflect the written user data, and for updating WFAT.

- 5 The checkpointing filesystem of a distributed-memory parallel supercomputer further comprising deleting user data from the checkpointing filesystem, wherein: the interface further comprises means for deleting sectors associated with the user data to be deleted from the disk from the FAT, for determining whether the sectors to be deleted are marked in the WFAT, and for updating the DFAT if the sectors  
10 are not marked in WFAT.

### **BRIEF DESCRIPTION OF THE DRAWINGS**

- The objects, features and advantages of the present invention will become apparent  
15 to one skilled in the art, in view of the following detailed description taken in combination with the attached drawings, in which:

Figure 1 illustrates an exemplary high-level system diagram for performing checkpointing of the filesystem according to the present invention;  
20

Figure 2 illustrates is an exemplary more-detailed system diagram of Figure 1 for performing checkpointing of the filesystem according to the present invention;

Figure 3 illustrates a conventional filesystem without checkpointing;  
25

Figure 4 illustrates an exemplary checkpointing filesystem according to the present invention;

Figure 5 illustrates an exemplary method flowchart that depicts taking a checkpoint  
30 in the checkpointing filesystem of Figure 4 according to the present invention;

Figure 6 illustrates an exemplary method flowchart that depicts rolling back to a previously taken checkpoint in the checkpointing filesystem of Figure 4 according to the present invention;

- 5 Figure 7 illustrates an exemplary method flowchart that depicts writing of user data to the checkpointing filesystem 400 of Figure 4 according to the present invention;

Figure 8 illustrates an exemplary method flowchart that depicts deleting of user data from the checkpointing filesystem 400 according to the present invention; and

10

Figure 9 illustrates an exemplary flowchart that depicts reading of user data in the checkpointing filesystem 400 according to the present invention.

15

#### **DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT OF THE INVENTION**

- The present invention is directed a checkpointing filesystem on the distributed-memory parallel supercomputer that facilitates faster execution of an application  
20 executing on the distributed-memory parallel supercomputer.

- Figure 1 is an exemplary high-level system diagram 100 for performing checkpointing of the filesystem according to the present invention. The high performance computer engine (i.e., "HPCE") 102 represents the distributed-memory parallel supercomputer. The HPCE is a distributed-memory parallel  
25 supercomputer that comprises a plurality of nodes. The HPCE 102 communicates via the network switch 104 with the fileserver/filesystem 106 and host control server 108. The HPCE communicates with the foregoing system components over communication connections 110, which may represent standard networking  
30 connections, such as Ethernet and the like. The host control server 108 controls the HPCE 102 and the fileserver/filesystem 106, and determines when to perform a checkpoint of the filesystem and when to restore the filesystem from a previous checkpoint as will be described below in greater detail.

Figure 2 is an exemplary more-detailed system diagram 200 of Figure 1 for performing checkpointing of the filesystem according to the present invention. The fileserver/filesystem 106 comprises a plurality of personal computers (i.e., "PCs") 202. The distributed-memory parallel supercomputer comprises

5 approximately 1024 PCs. In the distributed-memory parallel supercomputer, each PC 202 serves a Redundant Array of Independent Disks (i.e., "RAID") cage, i.e., a plurality of disk drives. For brevity and clarity of the present invention, the following description will be in terms of each PC 202 serving a disk 204 instead of a RAID cage. More particularly, each PC 202 serves as an interface for the HPCE

10 102 to an associated disk 204 via the network switch 104. The HPCE 102 communicates with the fileserver/filesystem 106 via messages that represent basic file operations, such as read, write and delete. A node of the HPCE 102 may access any file stored on any disk 204 via associated PC interface 202.

Further with reference to Figure 2, in order to perform a checkpoint, the host

15 control 108 informs each of the nodes of the HPCE 102 and the PCs 202 to complete any disk activity (i.e., access) and not to start any new disk activity. The host control 108 directs each PC 202 to make a checkpoint of its associated disk 204. This checkpointing activity is completely local to each PC 202 and its associated disk 204. The checkpointing activity is the same with regard to all PCs,

20 i.e., as if a single PC 202 performs a checkpoint of its associated disk 204. Similarly, when a checkpoint is restored, the restoration is completely local to each PC 202 and its associated disk 204. Therefore, in the following description performing a checkpoint and restoring a checkpoint will be described in detail with regard to a single PC 202 and associated drive 204. The host control 108 is a

25 complete computer including its own internal filesystem, which maintains the software for driving the host control 108, including the taking of a checkpoint and rolling back to previous checkpoint as described below in greater detail. In fileserver/filesystem 106, each PC 202 is a complete computer including its own internal filesystem, in addition to the associated disk 204. The internal filesystem

30 of the PC 202 maintains the software for driving the PC 202, including the taking of a checkpoint and rolling back to a previous checkpoint as described below in greater detail.

Figure 3 depicts a conventional filesystem 300 without checkpointing. The filesystem 300 occupies an entire disk or a portion of the disk. The filesystem comprises two parts. The first part is user data 304 and the free space 306 in the filesystem. The second part is a File Allocation Table (i.e., "FAT"), also known as metadata, which includes a description of the user data 304 and the free space 306. Thus, the FAT 302 includes filenames, file data start sectors, file sizes, as well as free sectors. An example of such a FAT 302 is the industry standard FAT32.

Figure 4 is an exemplary checkpointing filesystem 400 according to the present invention. The filesystem comprises five parts. The first part is the user data 304 and the free space 306 in the checkpointing filesystem 400, which are analogous to the filesystem without checkpointing 300 in Figure 3. The second part is the File Allocation Table (i.e., "FAT") 302, which is analogous to that of the filesystem without checkpointing 300 in Figure 3. The third part of the checkpointing filesystem 400 is a recently Written FAT (i.e., "WFAT") 402, which is structured like the FAT 302, but only maintains information on the user data 304 written since a previous checkpoint. The fourth part of the checkpointing filesystem 400 is a recently Deleted FAT (i.e., "DFAT") 404, which is structured like the FAT 302, but only maintains information on the user data 304 deleted since the previous checkpoint. The fifth and final part the checkpointing filesystem 400 is the recently Written or Deleted list (i.e., "WDLIST") of user data, which includes the same information as the WFAT 402 and the DFAT 404. The WFAT 402 and DFAT 404 are optimal for fast random access, since they provide indexed by sector access to information about recently written or deleted user data 304. Like the FAT 302, the WFAT 402 includes an entry for every sector in the checkpointing filesystem 400. Typically only a small fraction, if any, of sectors includes user data 304 written since the previous checkpoint, so most of the WFAT 402 entries are empty. Similarly, most of the DFAT 404 entries are empty. By contrast to the WFAT 402 and DFAT 404, the WDLIST 406 is optimal for fast sequential access to all the information about the recently written or deleted user data 304. More particularly, the WDLIST 406 provides a list of sectors in the WFAT 402 marking data written since the previous checkpoint and of sectors in the DFAT 404 marking data deleted since the previous checkpoint.

Figure 5 is an exemplary method flowchart 500 that depicts taking a checkpoint in the checkpointing filesystem 400 according to the present invention. The taking of the checkpoint starts at step 502. At step 504, the WFAT 402 is cleared, since at the checkpoint the user data 304 written since a previous checkpoint becomes part of the checkpoint. Since the WFAT 402, just like the FAT 302, is a large data structure, clearing the WFAT 402 by iterating through the entire data structure is inefficient. Because the WDLIST 406 is a much smaller data structure than the WFAT 402, it is faster to clear the WFAT 402 using entries in the WDLIST 406. At step 506, the DFAT 404 is cleared since at the checkpoint the user data 304 deleted since the previous checkpoint must be deleted from the checkpoint. For efficiency, the DFAT 404 is cleared using WDLIST 406. More particularly, the WDLIST 406 is a list of all sectors in the WFAT 402 and DFAT 404 marking user data 304 written or deleted since the previous checkpoint. The WFAT 404 and DFAT 404 are cleared by unmarking each sector listed in the WDLIST 406. After the WFAT 402 and the DFAT 404 are respectively cleared at steps 504 and 506, the WDLIST 406 is also cleared by deleting its list of marked sectors. At this point the filesystem is in a clean or consistent state. The exemplary taking of the checkpoint 500 is fast as required by the distributed-memory parallel supercomputer illustrated in Figures 1 and 2, and minimally impacts the application.

Figure 6 is an exemplary method flowchart 600 that depicts rolling back to a previously taken checkpoint according to the present invention. The rolling back to the previously taken checkpoint starts at step 602. At step 604, the FAT 302 is returned to the previously taken checkpoint using WDLIST 406. More particularly, at the rollback to the previously taken checkpoint, the user data 304 written since the previous checkpoint is removed from the FAT 302, and the user data 304 deleted since the previous checkpoint is returned to the FAT 302. Since the WDLIST 406 includes all the information on the user data 304 written and deleted since the previously taken checkpoint, the WDLIST is used to return the FAT to the previously taken checkpoint. At step 606, the WFAT 402 is cleared utilizing WDLIST 406 for efficiency. Since, the WDLIST 406 is a list of all

- sectors in the WFAT 402 and DFAT 404 marking user data 304 written or deleted since the previous checkpoint, the WFAT 402 and DFAT 404 are cleared by unmarking each sector listed in the WDLIST 406. At step 608, the DFAT 404 is cleared utilizing the WDLIST 406. At step 610, the WDLIST 406 is then cleared.
- 5 The situation immediately after the rollback of the FAT is the same as after taking a checkpoint in Figure 5 above, i.e., the filesystem is in a clean or consistent state. Therefore, the application may resume execution from the rolled back checkpoint without restarting the application.
- 10 Figure 7 is an exemplary method flowchart 700 that depicts a write of user data 304 to the checkpointing filesystem 400 of Figure 4 according to the present invention. It is to be noted that the steps of Figure 7 described below are behind the scenes and are not visible to the application. The writing of the user data 304 to the checkpointing filesystem 400 starts at step 702. At step 704, sectors marked as
- 15 free in the FAT 302 and not marked for deletion in the DFAT 404 are chosen for the write. More particularly, as in the case of the filesystem without checkpointing 300 of Figure 3, free sectors to contain the user data 304 that is to be written have to be found first, except that in the checkpointing filesystem 400 the sectors cannot be marked as deleted in the DFAT 404 because they are sectors marked as deleted
- 20 since the previously taken checkpoint and must be preserved in case of a rollback to the previously taken checkpoint. Given that the free sectors are not in the DFAT 402, the user data 304 may be written to the free sectors at step 706. At step 708, the FAT 302 is updated to reflect the write of user data 304. Since the write is after a previous checkpoint, WFAT 402 and WDLIST 406 are respectively updated
- 25 at steps 710 and 712. At step 714, the writing of user data 304 to the checkpointing filesystem 400 ends. In the conventional filesystem 300 without checkpointing, the filesystem 300 also performs a set of steps for writing the user data to the filesystem 300. The conventional steps include: choosing sectors marked free in the FAT 302; writing the user data 304 to the free sectors; and
- 30 updating the FAT to reflect the user data written. Over and above the conventional filesystem 300, the following steps are additional due to the checkpointing of the filesystem according to the present invention: ensuring that the free sectors are not in DFAT 404; updating WFAT 402 and updating WDLIST 406 as particularly

described above. The additional steps take little time as compared to the writing the user data 304. For example, ensuring that the sectors are not in DFAT is a fast random access lookup. Thus the checkpointing filesystem 400 of Figure 4 according to the present invention is almost as fast as the filesystem without  
5 checkpointing 300 of Figure 3. This fact satisfies the speed required by the distributed-memory parallel supercomputer illustrated in Figures 1 and 2.

Figure 8 is an exemplary method flowchart 800 describing how a delete of user data 304 from the checkpointing filesystem 400 according to the present invention.  
10 The deleting of the user data 304 from the checkpointing filesystem 400 starts at step 802. It is to be noted that the steps of Figure 8 described below are behind the scenes and are not visible to the application. At step 804, sectors associated with user data 304 to be deleted are deleted from the FAT 302. At step 806, it is determined whether the sectors to be deleted from the FAT 302 are marked in the  
15 WFAT 402. If the sectors to be deleted are marked in the WFAT 402, then the sectors are deleted from the WFAT 402 at step 808. Otherwise, if the sectors to be deleted are not marked in the WFAT 402, the DFAT 404 is updated at step 810 by marking in the DFAT 404 the deleted sectors. More particularly with regard to steps 806 and 808, if the sectors to be deleted are marked in the WFAT 402, then  
20 the entries are unmarked in the WFAT 402 and the DFAT 404 is not updated since the sectors deleted were created after the previous checkpoint and thus should not be recovered in a rollback. With regard to steps 806 and 810, the sectors to be deleted are not marked in the WFAT 402, and thus the sectors were created before the previous checkpoint and thus should be recovered in a rollback. Therefore, the  
25 sectors are marked in the DFAT 404. In either case, whether or not the sectors to be deleted are marked in the WFAT 402, the WDLIST 406 is updated at step 812 with the sectors unmarked in the WFAT 402 or marked in the DFAT 404. At step 814, the deleting of the user data from the checkpointing filesystem 400 ends. In the conventional filesystem 300 without checkpointing, the filesystem 300 also  
30 updates the FAT 302 with information concerning the delete. However, the remaining steps of Figure 8 are unique to the checkpointing filesystem 400 according to the present invention. Since the remaining steps take time that is

comparable to updating the FAT 302, the checkpointing filesystem according to the present invention is fast.

Figure 9 is an exemplary flowchart 900 that depicts how a read of user data 304 in the checkpointing filesystem 400 according to the present invention. The reading of the user data 304 from the checkpointing filesystem 400 starts at step 902. At step 904, the PC 202 receives a request from a node of the distributed-memory parallel supercomputer to read a file. At step 906, the FAT 302 is examined to determine which sectors to read. At step 908, the PC 202 reads the sectors determined at step 906 from the associated disk 204. At step 910, PC 202 transmits the read sectors to the requesting node of the distributed-memory parallel supercomputer. At step 912, the reading of the user data 304 from the checkpointing filesystem 400 ends. Thus, the reading of user data 304 from the checkpointing filesystem is exactly as it is in a filesystem without checkpointing 300 of Figure 3. More particularly, the reads in the checkpointing filesystem 400 of distributed-memory parallel supercomputer illustrated in Figures 1 and 2 are as fast as the on the supercomputer with a filesystem without checkpointing 300.

While the invention has been particularly shown and described with regard to preferred embodiments thereof, it will be understood by those skilled in the art that the foregoing and other changes in form and details may be made therein without departing from the spirit and scope of the invention.

**CLAIMS:**

1

Having thus described our invention, what we claim as new, and desire to secure by Letters Patent is:

1 1. A method for checkpointing a filesystem of a distributed-memory parallel  
2 supercomputer comprising a node that accesses user data on the filesystem, the  
3 filesystem comprising an interface that is associated with a disk for storing the user  
4 data, the method comprising the steps of:

5 (a) informing the node and the interface to complete their access to  
6 the filesystem in order to take a checkpoint of the filesystem;

7 (b) directing the interface to take the checkpoint of the associated  
8 disk, the taking of the checkpoint comprising a step of clearing:

9 (i) a recently written file allocation table (WFAT) for  
10 maintaining information regarding the user data written to the disk since a  
11 previously taken checkpoint;

12 (ii) a recently deleted file allocation table (DFAT) for  
13 maintaining information regarding the user data deleted from the disk since the  
14 previously taken checkpoint;

15 (c) informing the node and the interface that they may resume their  
16 access to the filesystem, wherein the WFAT and DFAT are utilized to rollback the  
17 filesystem to a clean state.

1 2. The method for checkpointing a filesystem of a distributed-memory parallel  
2 supercomputer according to Claim 1, wherein the WFAT and the DFAT are  
3 accessible in random order.

1 3. The method for checkpointing a filesystem of a distributed-memory parallel  
2 supercomputer according to Claim 1, wherein a written or deleted list (WDLIST) is  
3 provided for maintaining information associated with the WFAT and the DFAT  
4 and is accessible in a sequential order.

- 1 4. The method for checkpointing a filesystem of a distributed-memory parallel  
2 supercomputer according to according to Claim 3, wherein the step of clearing  
3 utilizes the information of the WDLIST to clear the WFAT and DFAT.
- 1 5. The method for checkpointing a filesystem of a distributed-memory parallel  
2 supercomputer according to according to Claim 4, wherein the method further  
3 comprises a step of clearing the WDLIST.
- 1 6. A method for rolling back a checkpointing filesystem of a distributed-memory  
2 parallel supercomputer comprising a node that accesses user data on the filesystem  
3 to a previously taken checkpoint of the filesystem, the filesystem comprising an  
4 interface that is associated with a disk for storing the user data, the method  
5 comprising the steps of:  
6 (a) informing the node and the interface to complete their access to  
7 the filesystem in order rollback to the previously taken checkpoint of the  
8 filesystem;  
9 (b) directing the interface to rollback to the previously taken  
10 checkpoint of the associated disk, the rolling back to the previously taken  
11 checkpoint comprising:  
12 (i) removing user data from a file allocation table (FAT)  
13 utilizing a recently written file allocation table (WFAT) that maintains information  
14 regarding the user data written to the disk since the previously taken checkpoint;  
15 (ii) adding user data to the FAT from a recently deleted file  
16 allocation table (DFAT) that maintains information regarding user data deleted  
17 from the disk since the previously taken checkpoint; and  
18 (c) informing the node and the interface that they may resume their  
19 access to the filesystem, wherein the WFAT and DFAT are utilized to rollback the  
20 filesystem to a clean state.
- 1 7. The method for rolling back a checkpointing filesystem of a distributed-memory  
2 parallel supercomputer according to Claim 6, wherein the WFAT and the DFAT  
3 are accessible in random order.

1 8. The method for rolling back a checkpointing filesystem of a distributed-memory  
2 parallel supercomputer according to Claim 1, wherein a written or deleted list  
3 (WDLIST) is provided for maintaining information associated with the WFAT and  
4 the DFAT and is accessible in a sequential order.

1 9. The method for rolling back a checkpointing filesystem of a distributed-memory  
2 parallel supercomputer according to Claim 8, wherein the steps of removing and  
3 adding user data utilize the information of the WDLIST.

1 10. The method for rolling back a checkpointing filesystem of a distributed-  
2 memory parallel supercomputer according to Claim 9, wherein the method further  
3 comprises the steps of:  
4 clearing the WFAT and the DFAT utilizing information of the  
5 WDLIST; and  
6 clearing the WDLIST.

1 11. A method for writing user data to a checkpointing filesystem of a distributed-  
2 memory parallel supercomputer comprising a node that accesses the user data on  
3 the filesystem, the filesystem comprising an interface that is associated with a disk  
4 for storing the user data, the method comprising the steps of:  
5 (a) selecting sectors that are marked as free in a file allocation table  
6 (FAT) and not marked as deleted in a recently deleted file allocation table (DFAT)  
7 that maintains information regarding user data deleted from the disk since a  
8 previously taken checkpoint;  
9 (b) writing the user data to the selected sectors on the disk  
10 (c) updating the FAT to reflect the written user data; and  
11 (d) updating a recently written file allocation table (WFAT) that  
12 maintains information regarding the user data written to the disk since the  
13 previously taken checkpoint; wherein the WFAT and DFAT are utilized to rollback  
14 the filesystem to a clean state.

1 12. The method for writing user data to a checkpointing filesystem of a distributed-  
2 memory parallel supercomputer according to Claim 11, wherein the WFAT and the  
3 DFAT are accessible in random order.

1 13. A method for writing user data to a checkpointing filesystem of a distributed-  
2 memory parallel supercomputer Claim 11, wherein a written or deleted list  
3 (WDLIST) is provided for maintaining information associated with the WFAT and  
4 the DFAT and is accessible in a sequential order.

1 14. A method for deleting user data from a checkpointing filesystem of a  
2 distributed-memory parallel supercomputer comprising a node that accesses the  
3 user data on the filesystem, the filesystem comprising an interface that is  
4 associated with a disk for storing the user data, the method comprising the steps of:  
5 (a) deleting sectors associated with the user data to be deleted from  
6 the disk from a file allocation table (FAT);  
7 (b) determining whether the sectors to be deleted are marked in a  
8 recently written file allocation table (WFAT) that maintains information regarding  
9 user data written to the disk since a previously taken checkpoint; and  
10 (c) updating a recently deleted file allocation table (DFAT) if the  
11 sectors are not marked in WFAT at step (b), the DFAT maintaining information  
12 regarding the user data deleted from the disk since the previously taken checkpoint,  
13 wherein the WFAT and DFAT are utilized to rollback the filesystem to a clean  
14 state.

1 15. The method for deleting user data from a checkpointing filesystem of a  
2 distributed-memory parallel supercomputer according to Claim 14, wherein the  
3 WFAT and the DFAT are accessible in random order.

1 16. The method for deleting user data from a checkpointing filesystem of a  
2 distributed-memory parallel supercomputer according to Claim 14, wherein a  
3 written or deleted list (WDLIST) is provided for maintaining information  
4 associated with the WFAT and the DFAT and is accessible in a sequential order.

1 17. The method for deleting user data from a checkpointing filesystem of a  
2 distributed-memory parallel supercomputer according to Claim 16, further  
3 comprising a step of updating the WDLIST.

1 18. A checkpointing filesystem of a distributed-memory parallel supercomputer  
2 comprising a node that accesses user data on the filesystem, the filesystem  
3 comprising an interface that is associated with a disk for storing the user data,  
4 comprising:

5 a host controller comprising means for informing the node and the  
6 interface to complete their access to the filesystem in order to take a checkpoint of  
7 the filesystem, directing the interface to take a checkpoint of the associated disk,  
8 and informing the node and the interface that they may resume their access to the  
9 filesystem after the checkpoint is taken; and

10 the interface comprising a means for clearing: a recently written file  
11 allocation table (WFAT) that maintains information regarding the user data written  
12 to the disk since a previously taken checkpoint; and a recently deleted file  
13 allocation table (DFAT) for maintaining information regarding the user data  
14 deleted from the disk since the previously taken checkpoint, wherein the WFAT  
15 and DFAT are utilized to rollback the filesystem to a clean state.

1 19. The checkpointing filesystem of a distributed-memory parallel supercomputer  
2 according to Claim 18, further comprising rolling back to the previously taken  
3 checkpoint, wherein:

4 the host controller further comprising means for informing the node  
5 and the interface to complete their access to the filesystem in order to rollback to the  
6 previously taken checkpoint of the filesystem, for directing the interface to  
7 rollback to the previously taken checkpoint of the associated disk, and for  
8 informing the node and the interface that they may resume their access to the  
9 filesystem; and

10 the interface further comprising means for rolling back to the  
11 previously taken checkpoint, the means comprising: removing user data from a file  
12 allocation table (FAT) utilizing the WFAT and adding user data to the FAT from  
13 DFAT.

1 20. The checkpointing filesystem of a distributed-memory parallel supercomputer  
2 according to Claim 18, further comprising writing user data to the checkpointing  
3 filesystem:

4 the interface further comprising means for selecting sectors that are  
5 marked as free in the FAT and not marked for deletion in DFAT, for writing the  
6 user data to the selected sectors on the disk; for updating the FAT to reflect the  
7 written user data, and for updating WFAT.

1 21. The checkpointing filesystem of a distributed-memory parallel supercomputer  
2 according to Claim 18, further comprising deleting user data from the  
3 checkpointing filesystem:

4 the interface further comprising means for deleting sectors  
5 associated with the user data to be deleted from the disk from the FAT, for  
6 determining whether the sectors to be deleted are marked in the WFAT, and for  
7 updating the DFAT if the sectors are not marked in WFAT.

1 22. The checkpointing filesystem of a distributed-memory parallel supercomputer  
2 according to Claim 18, wherein the WFAT and the DFAT are accessible in random  
3 order.

1 23. The checkpointing filesystem of a distributed-memory parallel supercomputer  
2 according to Claim 18, wherein a written or deleted list (WDLIST) is provided for  
3 maintaining information associated with the WFAT and the DFAT and is  
4 accessible in a sequential order.

5

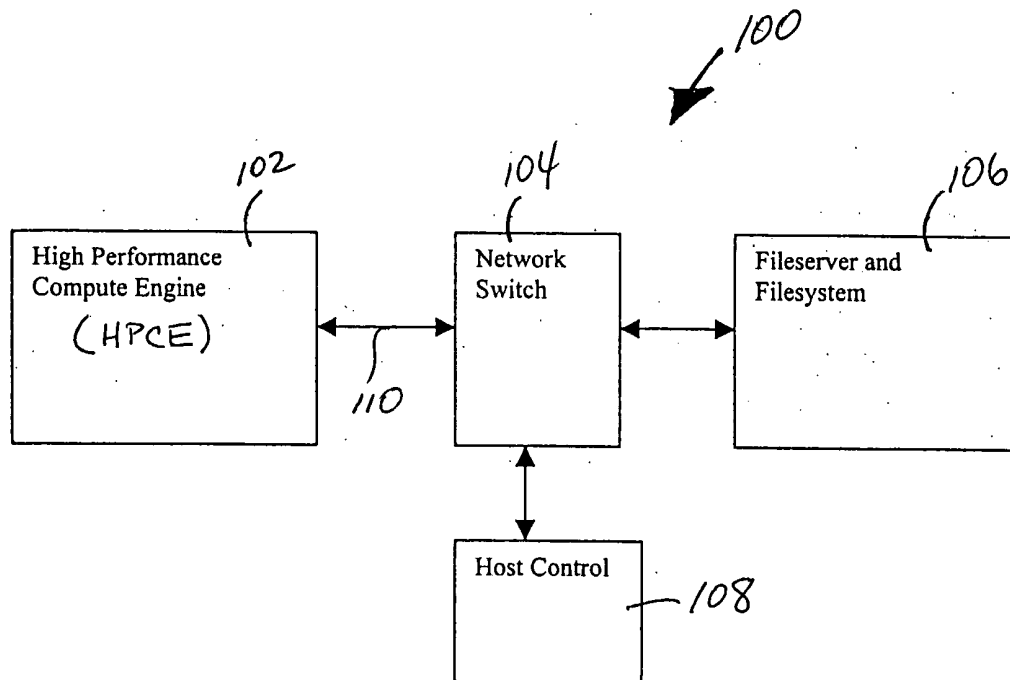
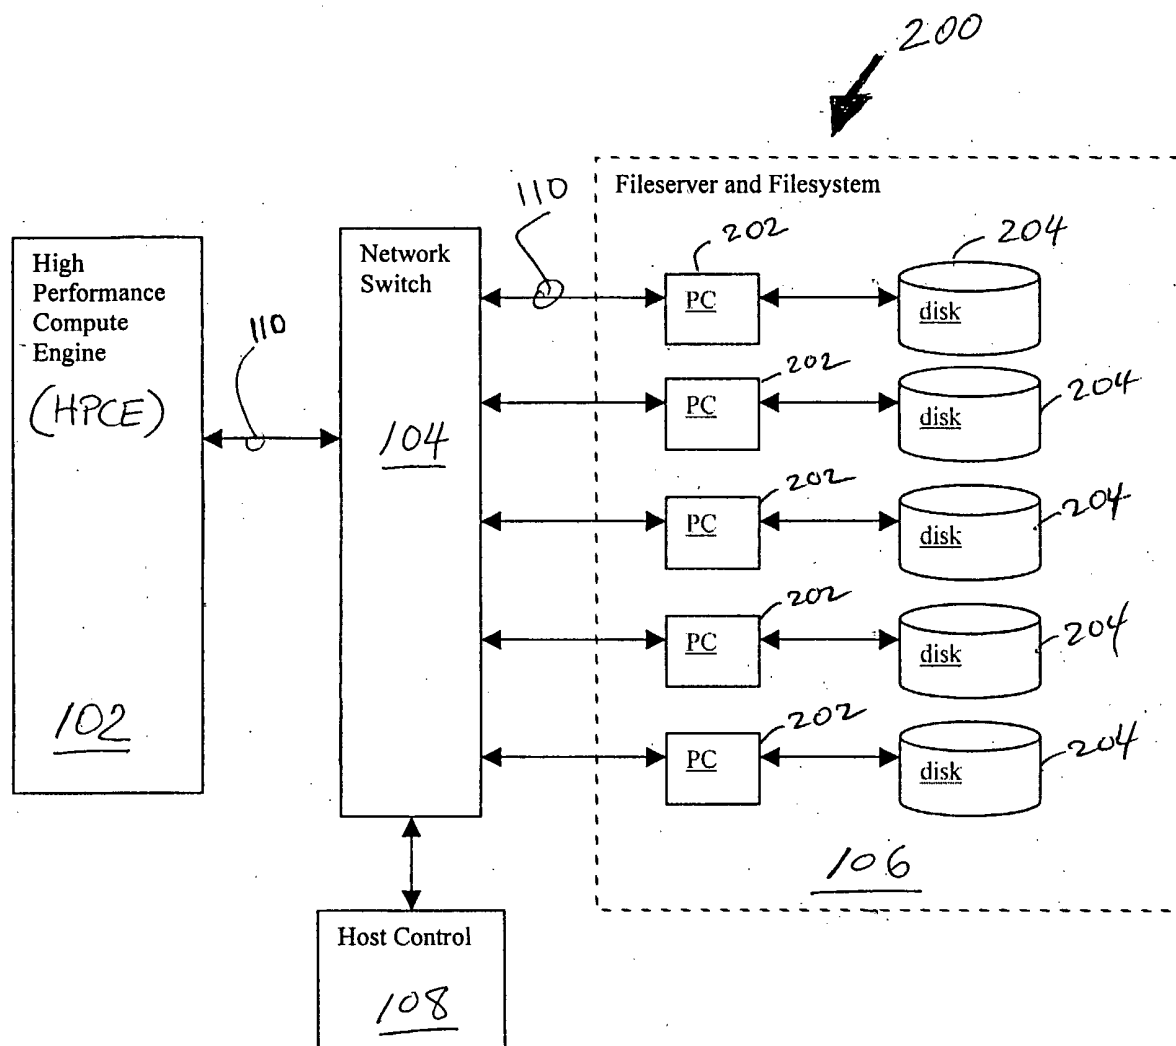


FIGURE 1



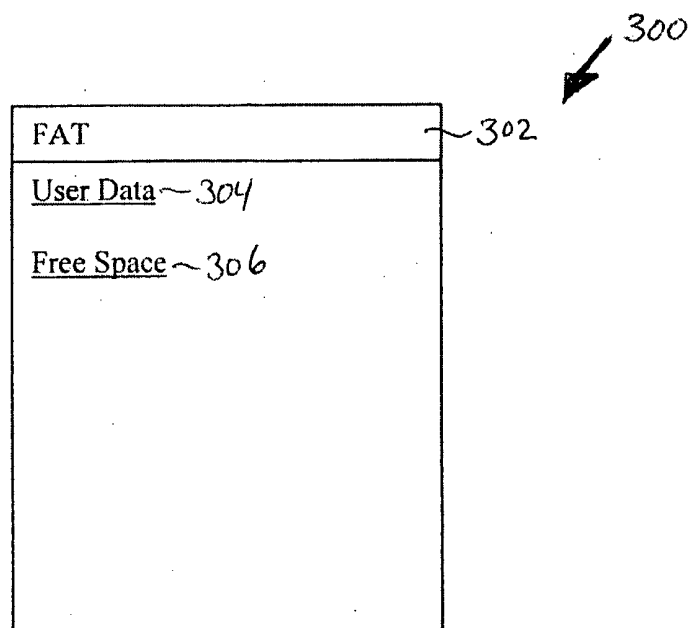


Figure 3

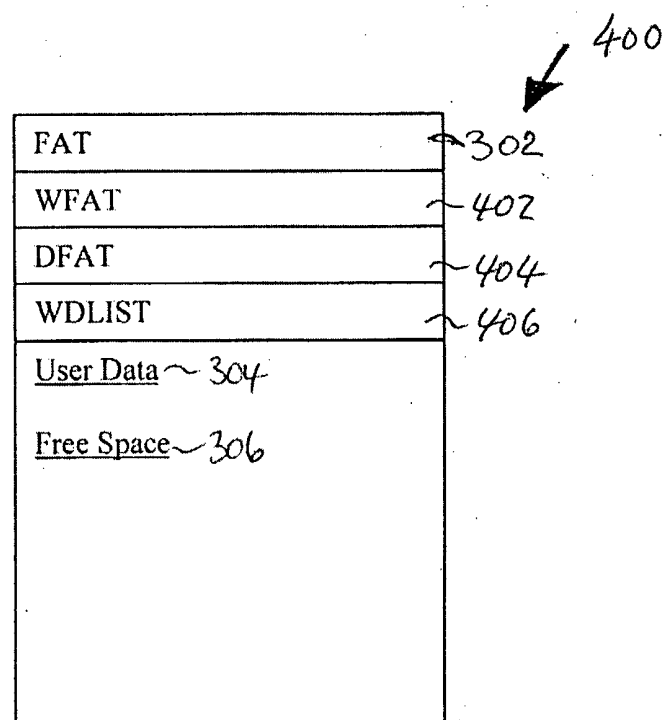


Figure 4

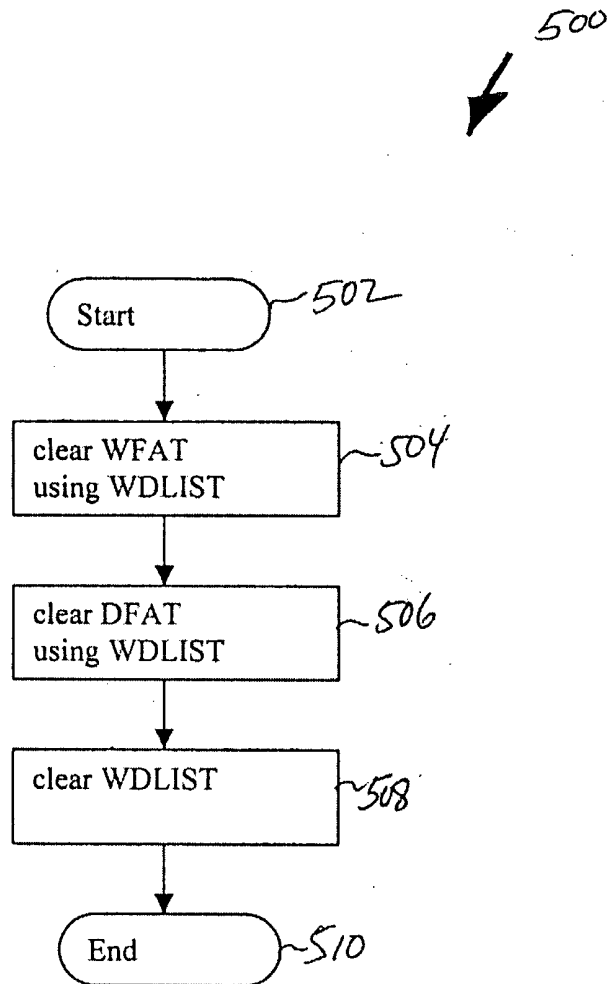


Figure 5

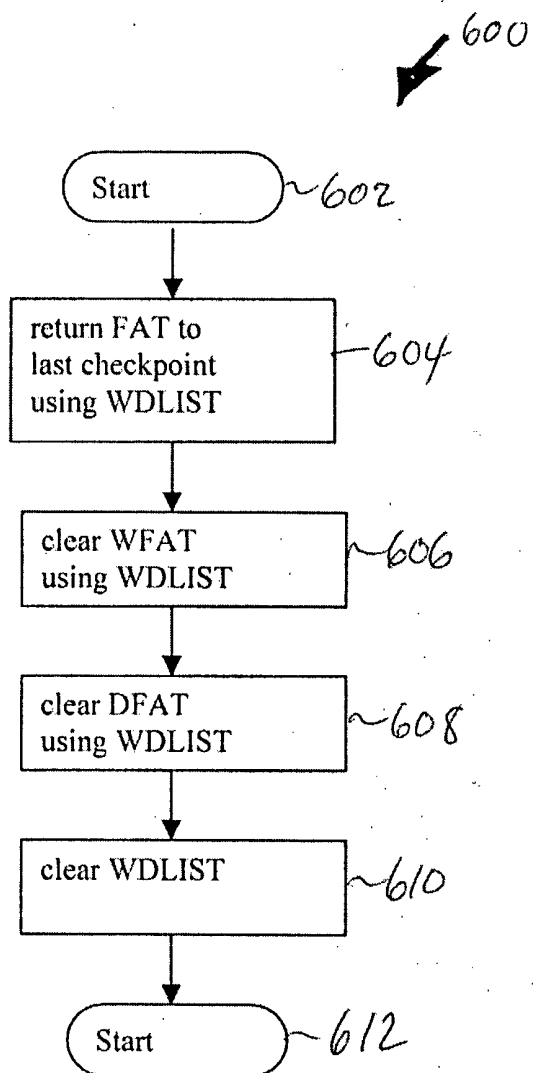
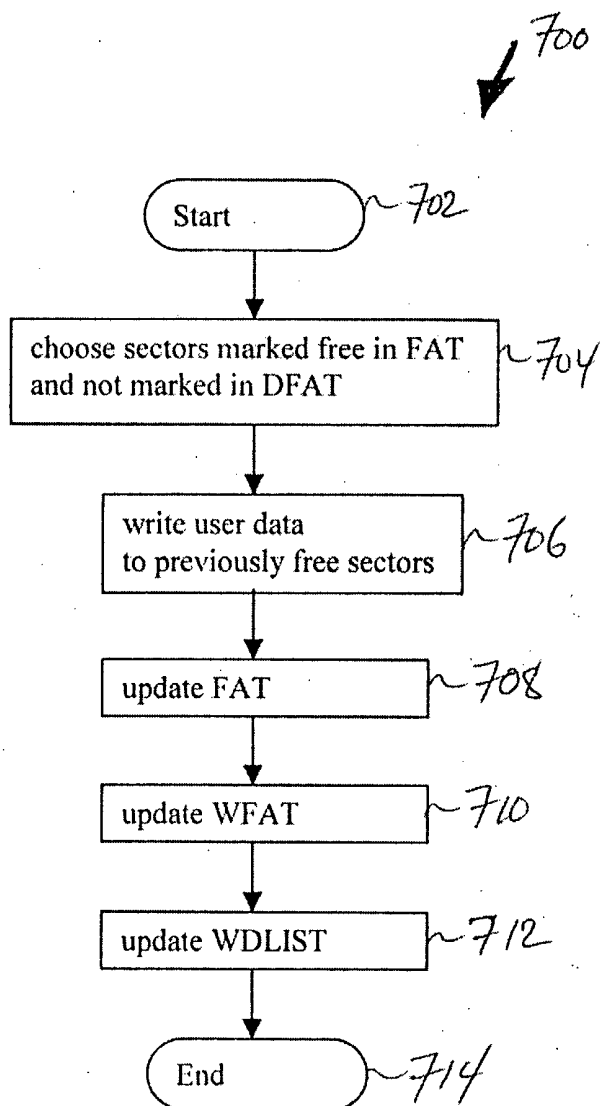


Figure 6

*Figure 7*

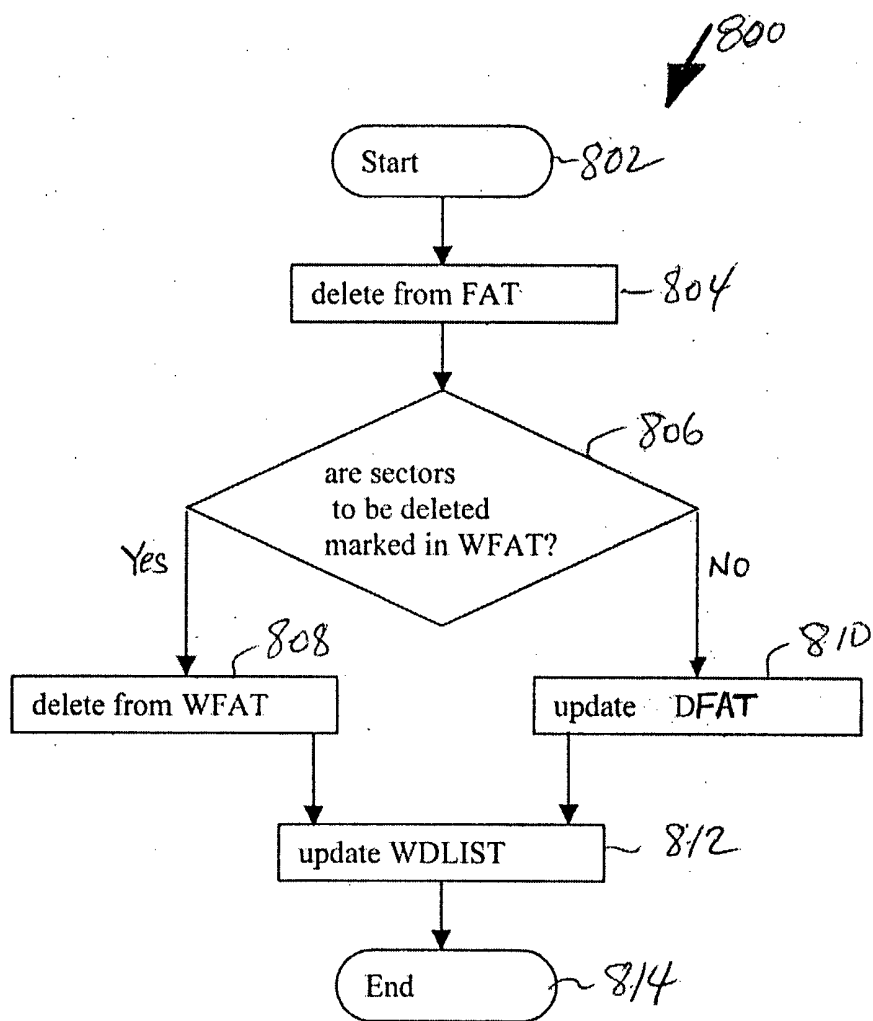


Figure 8

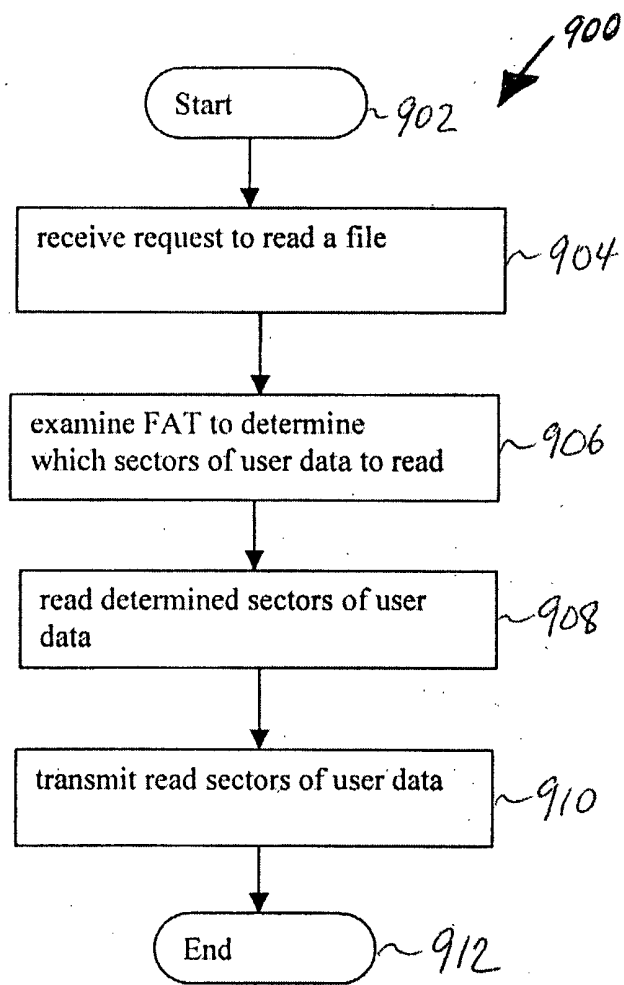


Figure 9

# INTERNATIONAL SEARCH REPORT

International application No.

PCT/US02/05613

<b>A. CLASSIFICATION OF SUBJECT MATTER</b>																				
IPC(7) : G06F 17/30																				
US CL : 707/201, 202, 204																				
According to International Patent Classification (IPC) or to both national classification and IPC																				
<b>B. FIELDS SEARCHED</b>																				
Minimum documentation searched (classification system followed by classification symbols) U.S. : 707/201, 202, 204																				
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched																				
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) Please See Continuation Sheet																				
<b>C. DOCUMENTS CONSIDERED TO BE RELEVANT</b>																				
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.																		
Y	US 6,173,292 B1 (BARBER et al) 09 January 2001 (09.01.2001) whole document.	1-10 ----- 18-23																		
Y	US 5,551,043 A (CRUMP et al) 27 August 1996 (27.08.1996) whole document.	1-10 ----- 18-23																		
X	US 5,864,849 A (BOHANNON et al) 26 January 1999 (26.01.1999) whole document.	11-17																		
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input type="checkbox"/> See patent family annex.																				
<table border="0"> <tr> <td colspan="2">* Special categories of cited documents:</td> <td>"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</td> </tr> <tr> <td>"A" document defining the general state of the art which is not considered to be of particular relevance</td> <td></td> <td>"X" document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone</td> </tr> <tr> <td>"E" earlier application or patent published on or after the international filing date</td> <td></td> <td>"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</td> </tr> <tr> <td>"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)</td> <td></td> <td>"&amp;" document member of the same patent family</td> </tr> <tr> <td>"O" document referring to an oral disclosure, use, exhibition or other means</td> <td></td> <td></td> </tr> <tr> <td>"P" document published prior to the international filing date but later than the priority date claimed</td> <td></td> <td></td> </tr> </table>			* Special categories of cited documents:		"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	"A" document defining the general state of the art which is not considered to be of particular relevance		"X" document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone	"E" earlier application or patent published on or after the international filing date		"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	"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)		"&" document member of the same patent family	"O" document referring to an oral disclosure, use, exhibition or other means			"P" document published prior to the international filing date but later than the priority date claimed		
* Special categories of cited documents:		"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																		
"A" document defining the general state of the art which is not considered to be of particular relevance		"X" document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone																		
"E" earlier application or patent published on or after the international filing date		"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																		
"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)		"&" document member of the same patent family																		
"O" document referring to an oral disclosure, use, exhibition or other means																				
"P" document published prior to the international filing date but later than the priority date claimed																				
Date of the actual completion of the international search 04 June 2002 (04.06.2002)		Date of mailing of the international search report 03 JUL 2002																		
Name and mailing address of the ISA/US Commissioner of Patents and Trademarks Box PCT Washington, D.C. 20231 Facsimile No. (703)305-3230		Authorized officer Haythim J. Alaubaidi <i>James R. Matthews</i> Telephone No. (703) 305-1950																		

# INTERNATIONAL SEARCH REPORT

International application No.

PCT/US02/05613

## Continuation of B. FIELDS SEARCHED Item 3:

WEST Search

Databases Searched: USPT, PGPB, JPAB, EPAB, DWPI, TDBD

Search Terms: checkpoint\$, file\$, rollback, fat, delet\$