



(51) International Patent Classification:
G06F 3/06 (2006.01) **G06F 17/30** (2006.01)

(21) International Application Number:
PCT/US2010/059298

(22) International Filing Date:
7 December 2010 (07.12.2010)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/645,364 22 December 2009 (22.12.2009) US

(71) Applicant (for all designated States except US): **APPLE INC.** [US/US]; 1 Infinite Loop, Cupertino, CA 95014 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **TAMURA, Eric, Brandon** [US/US]; 1 Infinite Loop, M/s:301-2cos, Cupertino, CA 95014 (US). **MAJNEMER, David, Alexander** [US/US]; 1 Infinite Loop, Cupertino, CA 95014 (US).

(74) Agents: **VINCENT, Lester, J.** et al.; Blakely, Sokoloff, Taylor & Zafman LLP, 1279 Oakmead Parkway, Sunnyvale, CA 94085-4040 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

(54) Title: METHODS AND APPARATUSES TO ALLOCATE FILE STORAGE VIA TREE REPRESENTATIONS OF A BITMAP

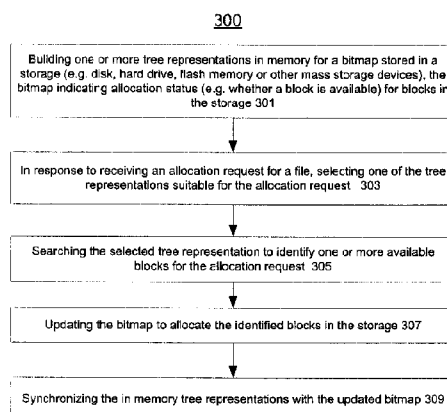


Fig. 3

(57) Abstract: Methods and apparatuses that search tree representations of a bitmap for available blocks to allocate in storage devices are described. An allocation request for a file may be received to initiate the search. In one embodiment, the bitmap may include an array of bits corresponding to blocks in the storage devices. Each bit may indicate whether one of the blocks is available. The tree representations may include at least one red-black tree having nodes corresponding to one or more consecutive bits in the bitmap indicating an extent of available blocks. One of the tree representations may be selected according to a file associated with an allocation request to identify an extent of available block matching the allocation request. The tree representations may be synchronized as the bitmap is updated with changes of block allocations in the storage devices.

METHODS AND APPARATUSES TO ALLOCATE FILE STORAGE VIA TREE REPRESENTATIONS OF A BITMAP

FIELD OF INVENTION

[0001] The present invention relates generally to file systems. More particularly, this invention relates to allocating blocks in storage devices for files based on a bitmap.

BACKGROUND

[0002] One major requirement for a file system is to keep track of available free space of storage devices in units of blocks. Traditionally, file systems may use a bitmap to represent free space. A bitmap is simply an array of bits, with the Nth bit indicating whether the Nth block is allocated or free. Thus, the overhead of a bitmap may be relatively low, such as, about 0.003% for one bit per 4K size block. For a 1GB file system, the corresponding bitmap is about 32KB in size, which may easily fit in memory for quick scanning to identify free space.

[0003] However, as sizes of file systems continue to grow faster than the growth of memory sizes, loading a bitmap used in a file system to scan may become non-trivial in either size or time. For example, a 32GB size bitmap for a 1 PB file system may not fit in memory on most data processing systems or machines. As a result, scanning a bitmap may require reading the bitmap from disk each time (e.g. paged in and out of a disk), which can significantly slow down a file system.

[0004] Furthermore, loading a large size bitmap into memory may directly compete with other kernel tasks for finite resources available in a data processing system. For example, if the bitmap cannot be completely cached, a number of buffers may be needed to manage paging operations. Because each buffer may require allocation of one of a finite number of buffer headers in a system, loading a large size bitmap can further degrade system performance.

[0005] Thus, traditional file systems that use a bitmap to allocate free space in storage devices do not scale with the growth of modern file systems.

SUMMARY OF THE DESCRIPTION

[0006] An embodiment of the present invention can include methods and apparatuses that search tree representations of a bitmap (such as a red-black tree) for available blocks to allocate in storage devices. An allocation request for a file may be received to initiate the search. In one embodiment, the bitmap may include an array of bits corresponding to blocks in the storage devices. Each bit may indicate whether one of the blocks is available. The tree representations may include at least one red-black tree having nodes corresponding to one or more consecutive bits in the bitmap indicating an extent (e.g. including offset and length) of available blocks. One of the tree representations may be selected based on the nature of the allocation request for a given file (e.g. allocation size request, whether the file exists or not etc.). The tree representations may be synchronized with the bitmap as it is updated with changes of block allocations in the storage devices.

[0007] In an alternative embodiment, a location tree and an extent tree may be maintained in memory (e.g. RAM) to represent a bitmap in storage devices (e.g. a hard drive) indicating allocation status for blocks in the storage devices. The location tree may be a red-black tree keyed according to locations of extents of available blocks. The extent tree may be a red-black tree keyed according to sizes of extents of available blocks. One or more nodes from the location tree and/or the extent tree may be identified based on searching the specified tree for extents of unallocated space that matches criteria based upon an allocation request for a file. The bitmap may be traversed if the location tree and/or the extent tree do not include matching nodes for search criteria. The location tree and the extent tree may be built up by traversing the bitmap and inserting nodes that represent unallocated regions of disk space.

[0008] Other features of the present invention will be apparent from the accompanying drawings and from the detailed description that follows.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] The present invention is illustrated by way of example and not limitation in the figures of the accompanying drawings, in which like references indicate similar elements and in which:

[0010] Figure 1 is a block diagram illustrating one embodiment of a system to search tree representations of a bitmap to allocate blocks;

[0011] Figure 2 is a sample diagram illustrating exemplary red-black trees for representing a bitmap;

[0012] Figure 3 is a flow diagram illustrating one embodiment of a process to identify available blocks for allocation without searching a bitmap;

[0013] Figure 4 is a flow diagram illustrating one embodiment of a process to search tree representations of a bitmap for block allocation;

[0014] Figure 5 illustrates one example of a data processing system such as a computer system, which may be used in conjunction with the embodiments described herein.

DETAILED DESCRIPTION

[0015] Methods and apparatuses for allocating blocks based on tree representations of a bitmap are described herein. In the following description, numerous specific details are set forth to provide thorough explanation of embodiments of the present invention. It will be apparent, however, to one skilled in the art, that embodiments of the present invention may be practiced without these specific details. In other instances, well-known components, structures, and techniques have not been shown in detail in order not to obscure the understanding of this description.

[0016] Reference in the specification to “one embodiment” or “an embodiment” means that a particular feature, structure, or characteristic described in connection with the embodiment can be included in at least one embodiment of the invention. The appearances of the phrase “in one embodiment” in various places in the specification do not necessarily all refer to the same embodiment.

[0017] The processes depicted in the figures that follow, are performed by processing logic that comprises hardware (e.g., circuitry, dedicated logic, etc.), software (such as is run on a general-purpose computer system or a dedicated machine), or a combination of both. Although the processes are described below in terms of some sequential operations, it should be appreciated that some of the operations described may be performed in different order. Moreover, some operations may be performed in parallel rather than sequentially.

[0018] In one embodiment, at least two red-black trees in memory (e.g. main memory such as DRAM) are used to represent a bitmap in disk (or other mass storage devices, such as flash memory, hard drives, etc.) mapping block allocation status in a file system. High level structures among bits in a bitmap may be explicitly described in tree representations (or tree structures) of the bitmap. For example, a section of 200 unused blocks starting at the 10th block in a bitmap may be readily identified in a tree node including an offset parameter of 10 and a size parameter of 200. Consequently, a perfect match for allocating 200 blocks, e.g. for a giant movie file, can be found within a bounded time frame based on an efficient red-black tree search.

[0019] In one embodiment, separate red-black trees representing a bitmap may enable searching for optimal allocation using different search criteria. For example, to locate a section of contiguous blocks to append to a file, a location based tree representation of the bitmap may be searched to find unallocated blocks nearest to a specified location. Alternatively, to locate a chunk of space for a new file, a size based tree representation of the bitmap may be searched to find a contiguous range of blocks for the chunk of space.

[0020] In the event a search cannot identify a perfect match, in one embodiment, a second search using different search criteria may be performed on the same tree or a different tree. For example, if searching a location based tree for a chunk of space fails to yield a match, subsequent searches may be performed to identify one or many sections of blocks based on a requested size, e.g. a minimum number of blocks having a capacity to accommodate the requested size.

[0021] In one embodiment, another search may be performed according to a search criteria based on a nearby location relative to an existing location of allocated blocks for a file within a configured threshold (e.g. 4 blocks). The search criteria may include the minimum number of blocks to identify extents of available blocks located close to the file and large enough for the requested size. Multiple search results may be combined together for optimal block allocation that minimizes disk fragmentation.

[0022] In some embodiments, one or more red-black trees maintained in memory to represent a bitmap kept in disk may be built on demand according to a lazy mechanism to conserve memory and/or provide faster allocation response. Thus, each tree may grow incrementally as a file system performs operations to allocate/free storage space. When a tree is fully populated, the tree may represent a bitmap in disk completely.

[0023] In one embodiment, the red-black trees may be built during system boot time without incurring additional cost, so that a bitmap representation may be loaded in memory when system starts. Typically, tree representations of a bitmap may be kept in memory without being stored in a disk. Updating tree representation of a bitmap may be performed in a file system prior or subsequent to changes of the bitmap in disk.

[0024] Figure 1 is a block diagram illustrating one embodiment of a system to search tree representations of a bitmap to allocate blocks. In one embodiment, system 100 may include a computer operating environment 105 having a file system 107 in an operating system. Storage 101 may be one or more storage devices, such as hard disks, flash memories or other mass storage media, locally or remotely coupled to system 101. In one embodiment, files in system 100 may be stored in one or more blocks allocated in storage 101. Storage 101 may include bitmap 103 representing availability of each allocated block using a bit (e.g. using 0/1 as status values to indicate if a block is free/allocated).

[0025] In one embodiment file system 107 may include interface module 119 to receive file access (read/write) requests from runtime programs 121, e.g. via API (application programming interface). File management module 117 may

determine file access operations, such as file read, file write, file creation, file deletion, etc. for the received file access requests. For example, file management module 117 may send a storage allocation request to disk block allocation module 115 for creating new file and/or storing/updating data to a file.

[0026] In one embodiment, file system 107 may include bitmap representations 109 for one or more in memory representations, such as location tree 111 or extent tree 113, of bitmap 103. Each in memory representation may correspond to at least a portion of bitmap 103. For example, each node in location tree 111 and/or extent tree 113 may include location and/or size information corresponding to one or more consecutive blocks (or extents) which are available as indicated in bitmap 103.

[0027] In one embodiment, bitmap representation management module 123 can build and/or maintain bitmap representations 109 in memory in sync with bitmap 103 in storage 101. For example, bitmap representation management module 123 can access a portion of bitmap 103 loaded in memory when a search is performed directly on bitmap 103. Alternatively, bitmap representation management module 123 may receive an update notification when bitmap 103 is updated.

[0028] File system 107 may include disk block allocation module 115 to allocate blocks in storage 101 for storing file content according to bitmap representations 109 without accessing bitmap 103 in storage 101. Disk block allocation module 115 may determine which blocks are available in storage 103 and select one or more available blocks large enough to accommodate an amount of data to be stored, e.g. in response to an allocation request for a file. In one embodiment, disk block allocation module 115 may access one or more trees in bitmap representations 109 in memory and/or bitmap 103 in storage 101 to identify available blocks to allocate.

[0029] For example, disk block allocation module 115 can dynamically determine an order (e.g. starting with location tree 111 for writing data to an existing file) to search separate trees in bitmap representations 109 for identifying available blocks for allocation in response to an allocation request, according to, e.g. one or more characteristics of the request, such as whether a file associated

with the request exists or not. Disk block allocation module 115 may directly search bitmap 103 when necessary, for example, if no desired blocks are found via bitmap representations 109.

[0030] In one embodiment, bitmap representation management module 123 can build bitmap representations 109 concurrent to disk block allocation module 115 accessing bitmap 103. For example, disk block allocation module 115 may load a portion of bitmap 103 in memory to update bitmap 103 for block allocation or to search for available blocks to allocate. Bit map representation management module 123 may build into bitmap representations 109 with block allocation status from a portion of bitmap 103 already loaded in memory. In one embodiment disk block allocation module 115 may determine when to access bitmap 103 for bitmap representation management module 123 to update and/or build bitmap representations 109.

[0031] Figure 2 is a sample diagram illustrating exemplary red-black trees for representing a bitmap. Example 200 may be based on a file system maintaining one or more trees to represent a bitmap for allocating blocks of file storage space, such as in system 100 of Figure 1. In one embodiment, bitmap 201 may include an array of bits to indicate an allocation status for blocks in storage devices, such as bitmap 103 in storage 101 of Figure 1. Each bit may indicate whether a corresponding block is available or not. For example, bit 207 may have value 0 for a block which has already been occupied.

[0032] In one embodiment, location tree 203 and extent tree 205 may be stored in memory, such in bitmap representations 109 of Figure 1, to represent portions of bitmap 201. Trees 203, 205 may be based on a self-balancing binary search tree as the underlying data structure to allow efficient update operations on the trees, such as search, delete, insert etc. For example, trees 203, 205 may be red-black trees using red-black tree update operations for ensuring the trees are reasonably balanced.

[0033] Location tree 203 may be configured with nodes representing available blocks indicated in bitmap 201 based on a red-black tree structure (or tree representation) including keys for the nodes. In one embodiment, location tree

203 may be keyed according to starting locations of consecutive available blocks or available block extents. For example, node 209 may include a location key 3 corresponding to a location of a starting block, e.g. indicated by bit 207, for 5 available consecutive blocks, e.g. indicated by 5 bits between bit 207 and bit 213. Node 209 may include a size key indicating the size of the consecutive blocks.

[0034] In one embodiment, extent tree 205 can be configured with nodes representing available blocks indicated in bitmap 201 using sizes of consecutive available blocks as keys based on a red-black tree structure. For example, node 211 may include a size key 5 corresponding to 5 consecutive available blocks, e.g. indicated by 5 bits between bit 207 and bit 213. Node 211 may include a location key, e.g. 3, to denote a starting block corresponding to bit 207 (the third bit) of bitmap 201.

[0035] Figure 3 is a flow diagram illustrating one embodiment of a process to identify available blocks for allocation without searching a bitmap. Exemplary process 300 may be performed by a processing logic that may comprise hardware (circuitry, dedicated logic, etc.), software (such as is run on a dedicated machine), or a combination of both. For example, process 300 may be performed by some components of system 100 of Figure 1. At block 301, the processing logic of process 300 may build one or more tree representations in memory, such as trees 111, 113 of Figure 1, to represent a bitmap maintaining status of block allocation in a storage system, such as bitmap 103 of Figure 1. For example, the processing logic of process 300 may start building tree representations for a bitmap when a system, such as system 100 of Figure 1, is booting up. A tree representation built in a system may stay in memory during the lifetime of the running system.

[0036] In one embodiment, the processing logic of process 300 may load a portion of a large bitmap (e.g. when compared with available memory resources) to build tree representations for the bitmap incrementally. In some embodiments, the processing logic of process 300 can generate tree representations for a portion of a bitmap loaded (e.g. read in from a disk storing the bitmap) when a system starts. A tree representation may correspond to (or cover) a portion of a bitmap.

The processing logic of process 300 may keep track of which portion of a bitmap is covered in a tree representation.

[0037] The processing logic of process 300 may identify extents of available blocks from a bitmap. Each extent may correspond to maximal consecutive available blocks (e.g. bounded by two blocks which are not available) indicated in a bitmap (e.g. based on bits of value 0) between two allocated blocks (e.g. based on bits of value 1). An extent may include a starting location indicated in the bit map, such as an index of a bit for a starting block of the extent. Alternatively, an extent may include a size representing, for example, number of blocks included in the extent. A node in a tree representation of a bit map may correspond to an extent of available blocks.

[0038] In one embodiment, when mounting a file system (e.g. including setting up where files should appear in a directory tree), the processing logic of process 300 may retrieve tree representations for a bitmap of the file system directly from a storage device, such as storage 101 of Figure 1, without loading (or scanning) the bitmap to rebuild the tree representations. The tree representations for the bitmap may be maintained in memory if the corresponding file system is mounted. In one embodiment, the processing logic of process 300 may store or write out the tree representations in a storage device, e.g. when a corresponding file system is unmounted (e.g. removed from an operating system), after a complete scan of a bitmap to update the tree representations for the bitmap. The processing logic of process 300 may determine to skip a full scan of a bitmap if the bitmap represented by available tree representations stored in a storage device has not been modified. For example, the processing logic of process 300 may compare time stamps of the bitmap and tree representations stored in the storage device to determine if the tree representations still represent the bitmap.

[0039] At block 303, in one embodiment, the processing logic of process 300 may receive an allocation request for a file, such as a file write request from runtime programs 121 of Figure 1. In response, the processing logic of process 300 may select a tree representation for a bitmap to identify available blocks most suitable for the received request. For example, the processing logic of process 300

may select a location tree, such as location tree 113 of Figure, for a request to allocate storage space for a file which has already existed in a file system.

Alternatively, the processing logic of process 300 may search a location tree, such as location tree 113 of Figure 1, to allocate storage space for a new file.

[0040] At block 305, in one embodiment, the processing logic of process 300 may search a selected tree representation of a bitmap to identify one or more available blocks to allocate for an allocation request based on one or more search criteria. In one embodiment, search criteria may include considerations to reduce fragmentation effects when allocating storage space (e.g. blocks).

[0041] For example, the processing logic of process 300 may perform a red-black tree search against a location tree to identify a node matching an allocation request specifying an allocation size and/or a starting location for an extent of available blocks. In one embodiment, a node in a location tree may be a match for a tree search if a corresponding extent of available blocks is large enough (e.g. based on a size key) to accommodate a requested allocation size and located (e.g. based on a location key) at the starting location.

[0042] Similarly, the processing logic of process 300 may perform a red-black tree search against an extent tree to identify a node matching an allocation request specifying an allocation size. For example, search criteria may include a minimum number of blocks to accommodate the requested allocation size. A node in an extent tree may be a match for the tree search if the node corresponds to an extent of the minimum number of available blocks. In some embodiments, a matching node may correspond to an extent of available blocks with a capacity large enough (e.g. based on a size key) to accommodate requested allocation size.

[0043] In another embodiment, the processing logic of process 300 may start searching a location tree in response to a request to allocate additional storage space for an existing file. The processing logic of process 300 may use search criteria to locate an extent of available blocks next to the last block allocated for the file. Alternatively, the processing logic of process 300 may start searching an extent tree to allocate storage space for creating a new file. The processing logic of process 300 may specify search criteria to identify an extent of available blocks

best matching the request for the new file (e.g. by comparing the size of an extent of available blocks and the size of storage space requested).

[0044] If no matched node is identified from a tree representation during a tree search, in one embodiment, the processing logic of process 300 may proceed to search another tree representation using similar or different search criteria.

Alternatively, if multiple nodes are matched during a tree search, the processing logic of process 300 may select a best one among the matched nodes to, for example, minimize fragmentation in storage space. In some embodiments, the processing logic of process 300 may directly search a portion of bitmap not yet represented by in memory tree representations, for example if matched nodes cannot be identified during the tree search. Optionally, the processing logic of process 300 may determine to allocate a plurality of blocks, which may not be consecutively located all together, for an allocation request based on extents of available blocks corresponding to multiple nodes in a tree representation of a bitmap.

[0045] At block 307, the processing logic of process 300 may update a bitmap to allocate identified blocks in a storage system (e.g. storage devices for file systems) based on results of searching in memory tree representations of the bitmap. The processing logic of process 300 may synchronize the in memory tree representations with the updated bitmap at block 309, e.g. via operations such as search, insert and/or delete performed according to red-black tree structures. In other embodiments, the processing logic of process 300 may update in memory tree representations of a bitmap to indicate allocation of identified blocks prior to update the bitmap in storage system.

[0046] Figure 4 is a flow diagram illustrating one embodiment of a process to search tree representations of a bitmap for block allocation. Exemplary process 400 may be performed by a processing logic that may comprise hardware (circuitry, dedicated logic, etc.), software (such as is run on a dedicated machine), or a combination of both. For example, process 400 may be performed by some components of system 100 of Figure 1. In one embodiment, at block 401, the processing logic of process 400 can maintain tree representations for at least a

portion of a bitmap indicating block allocation status in a storage system, such as storage 101 of Figure 1. For example, each bit in the bitmap may indicate whether a corresponding one of the blocks represented by the bitmap is available or not.

[0047] Tree representations may include, for example, a location tree and an extent tree, such as location tree 111 and extent tree 113 of Figure 1. Each tree representation may be a red-black tree with nodes corresponding to extents of available blocks in the storage system. The processing logic of process 400 can traverse the bitmap to build the tree representations whenever a portion of the bitmap is loaded in memory, e.g. when a system boots up. Optionally, the processing logic of process 400 can track which portions of a bitmap has been incorporated in each tree representation maintained, for example, based on indices of bits in the bitmap.

[0048] At block 403, the processing logic of process 400 may receive an allocation request to allocate an amount of space in a storage system for a file or other system components. In response, at block 405, the processing logic of process 400 may determine whether an allocation request is to create a new file, e.g. based on whether there are blocks already allocated for the file in a storage system. For an existing file, at block 407, the processing logic of process 400 can perform a red-black tree search in the location tree to identify a matching node based on the size of allocation space requested and/or locations of existing blocks allocated for the file.

[0049] A matching node may be perfect if the corresponding extent of available blocks has a capacity larger than the size of allocation space requested and located next to a block allocated for the file. In some embodiments, a search criteria for a matching node may locate an extent of available blocks close to (e.g. within a certain number of blocks) existing blocks allocated for the file. If a matching node is found, at block 409, the processing logic of process 400 may proceed to allocate storage space based on an extent of available blocks corresponding to a matching node at block 425. Otherwise, in one embodiment, at the processing logic of process 400 may proceed to search the extent tree for matching nodes at block 411.

[0050] In one embodiment, in response to an allocation request for a new file, the processing logic of process 400 may, at block 411, perform a red-black tree search in the extent tree to identify one or more matching nodes based on the size of allocation space requested. If more than one matching nodes are identified during the search, the processing logic of process 400 may select a best matching node (e.g. according to how closely the size of allocation space requested matches the space available in the extent corresponding to a matching node). In some embodiment, the processing logic of process 400 may select the first matched node during the tree search.

[0051] If matching nodes are found at block 413, the processing logic of process 400 may proceed to allocate identified blocks at block 425. Otherwise, at block 415, the processing logic of process 400 may determine if in memory tree representations incorporate all block allocation status in a bitmap. For example, the processing logic of process 400 can track which portion of the bitmap (e.g. via an index to bit locations in the bitmap) have been incorporated (or built) in the tree representations.

[0052] If the bitmap has not been completely built in the tree representations, at block 417, the processing logic of process 400 can traverse or search the bitmap directly to identify an extent of available blocks matching an allocation request. In one embodiment, the processing logic of process 400 may load a portion of the bitmap not yet covered in the tree representations to search for matching extents of available blocks. The portion of the bitmap may be loaded according to a size limitation of available memory space. In one embodiment, the processing logic of process 400 may search remaining portions of the bitmap, e.g. until a matching extent of available blocks is identified.

[0053] In the meanwhile when performing a search (e.g. a linear search bit by bit) directly in a bitmap, for example, at block 419, the processing logic of process 400 may continue building tree representations of the bitmap, such as a location tree and an extent tree, with loaded portions of the bitmap. If one or more matching extents of available blocks are found directly from the bitmap, the processing logic of process 400 may proceed to allocate blocks at 425.

[0054] If no matching node or extents of available blocks are found from tree representations of a bitmap and/or directly from the bitmap, at block 423, the processing logic of process 400 may identify one or more extents of available blocks satisfying an allocation request based on, for example, results of traversing tree representations (e.g. based on red-black tree search) of a bitmap and/or directly from the bitmap. The processing logic of process 400 may determine an optimal combination of available blocks to allocate for an allocation request considering, for example, minimizing fragmentation effects, maximizing storage usage, and/or fine tuning other file system performance parameters. Subsequently, the processing logic of process 400 may proceed to allocate identified available blocks at block 425. In one embodiment, at block 427, the processing logic of process 400 may update tree representations, such as a location tree and/or an extent tree, of a bitmap and the bitmap itself according to the allocation of identified blocks.

[0055] Figure 5 shows one example of another data processing system such as a computer system, which may be used with one embodiment the present invention. For example, the system 500 may be implemented as a part of the system shown in Figure 1. Note that while Figure 5 illustrates various components of a computer system, it is not intended to represent any particular architecture or manner of interconnecting the components as such details are not germane to the present invention. It will also be appreciated that network computers and other data processing systems which have fewer components or perhaps more components may also be used with the present invention.

[0056] As shown in Figure 5, the computer system 500, which is a form of a data processing system, includes a bus 503 which is coupled to a microprocessor(s) 505 and a ROM (Read Only Memory) 507 and volatile RAM 509 and a non-volatile memory 511. The microprocessor 505 may retrieve the instructions from the memories 507, 509, 511 and execute the instructions to perform operations described above. The bus 503 interconnects these various components together and also interconnects these components 505, 507, 509, and 511 to a display controller and display device 513 and to peripheral devices such

as input/output (I/O) devices which may be mice, keyboards, modems, network interfaces, printers and other devices which are well known in the art. Typically, the input/output devices 515 are coupled to the system through input/output controllers 517. The volatile RAM (Random Access Memory) 509 is typically implemented as dynamic RAM (DRAM) which requires power continually in order to refresh or maintain the data in the memory.

[0057] The mass storage 511 is typically a magnetic hard drive or a magnetic optical drive or an optical drive or a DVD RAM or a flash memory or other types of memory systems which maintain data (e.g. large amounts of data) even after power is removed from the system. Typically, the mass storage 511 will also be a random access memory although this is not required. While Figure 5 shows that the mass storage 511 is a local device coupled directly to the rest of the components in the data processing system, it will be appreciated that the present invention may utilize a non-volatile memory which is remote from the system, such as a network storage device which is coupled to the data processing system through a network interface such as a modem or Ethernet interface or wireless networking interface. The bus 503 may include one or more buses connected to each other through various bridges, controllers and/or adapters as is well known in the art.

[0058] Portions of what was described above may be implemented with logic circuitry such as a dedicated logic circuit or with a microcontroller or other form of processing core that executes program code instructions. Thus processes taught by the discussion above may be performed with program code such as machine-executable instructions that cause a machine that executes these instructions to perform certain functions. In this context, a “machine” may be a machine that converts intermediate form (or “abstract”) instructions into processor specific instructions (e.g., an abstract execution environment such as a “virtual machine” (e.g., a Java Virtual Machine), an interpreter, a Common Language Runtime, a high-level language virtual machine, etc.), and/or, electronic circuitry disposed on a semiconductor chip (e.g., “logic circuitry” implemented with transistors) designed to execute instructions such as a general-purpose

processor and/or a special-purpose processor. Processes taught by the discussion above may also be performed by (in the alternative to a machine or in combination with a machine) electronic circuitry designed to perform the processes (or a portion thereof) without the execution of program code.

[0059] An article of manufacture may be used to store program code. An article of manufacture that stores program code may be embodied as, but is not limited to, one or more memories (e.g., one or more flash memories, random access memories (static, dynamic or other)), optical disks, CD-ROMs, DVD ROMs, EPROMs, EEPROMs, magnetic or optical cards or other type of machine-readable media suitable for storing electronic instructions. Program code may also be downloaded from a remote computer (e.g., a server) to a requesting computer (e.g., a client) by way of data signals embodied in a propagation medium (e.g., via a communication link (e.g., a network connection)).

[0060] The preceding detailed descriptions are presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the tools used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of operations leading to a desired result. The operations are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

[0061] It should be kept in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the above discussion, it is appreciated that throughout the description, discussions utilizing terms such as "processing" or "computing" or "calculating" or "determining" or "displaying" or the like, refer to the action and

processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

[0062] The present invention also relates to an apparatus for performing the operations described herein. This apparatus may be specially constructed for the required purpose, or it may comprise a general-purpose computer selectively activated or reconfigured by a computer program stored in the computer. Such a computer program may be stored in a computer readable storage medium, such as, but is not limited to, any type of disk including floppy disks, optical disks, CD-ROMs, and magnetic-optical disks, read-only memories (ROMs), RAMs, EPROMs, EEPROMs, magnetic or optical cards, or any type of media suitable for storing electronic instructions, and each coupled to a computer system bus.

[0063] The processes and displays presented herein are not inherently related to any particular computer or other apparatus. Various general-purpose systems may be used with programs in accordance with the teachings herein, or it may prove convenient to construct a more specialized apparatus to perform the operations described. The required structure for a variety of these systems will be evident from the description below. In addition, the present invention is not described with reference to any particular programming language. It will be appreciated that a variety of programming languages may be used to implement the teachings of the invention as described herein.

[0064] The foregoing discussion merely describes some exemplary embodiments of the present invention. One skilled in the art will readily recognize from such discussion, the accompanying drawings and the claims that various modifications can be made without departing from the spirit and scope of the invention.

CLAIMS

What is claimed is:

1. A computer implemented method for a file system, the method comprising:
selecting, in response to a request to allocate space in storage devices for a file, one of a plurality of tree representations of a bitmap, each bit in the bitmap indicating availability of a block in the storage devices, each node in the tree representations corresponding to one or more consecutive bits in the bitmap;
searching the selected tree representation for one or more nodes
identifying blocks available for the request;
updating the bitmap to allocate the identified blocks in the storage devices;
and
synchronizing the plurality of trees with the updated bitmap.
2. The method of claim 1, wherein the plurality of tree representations are maintained in memory if the file system is mounted.
3. The method of claim 2, wherein the storage devices are mass storage devices storing files for the file system.
4. The method of claim 1, wherein the node includes a location of the consecutive bits in the bitmap and a size indicating a number of the consecutive bits.
5. The method of claim 4, wherein the selection comprises:
determining if the file is new in the file system.
6. The method of claim 5, wherein the plurality of tree representations include a location tree including nodes keyed according to the location and wherein the selected tree representation is the location tree if the file is not new.

7. The method of claim 6, wherein the location tree is a red-black tree.
8. The method of claim 6, wherein the file has an existing location, and wherein the search comprises:
 - matching the existing location with the nodes based on the location.
9. The method of claim 8, wherein the space has an allocation size, wherein the search further comprises:
 - comparing the allocation size with the nodes based on the size, wherein the identified available blocks have a capacity large enough for the allocation size.
10. The method of claim 5, wherein the plurality of tree representations include an extent tree including nodes keyed according to the size and wherein the selected tree representation is the extent tree if the file is new.
11. The method of claim 10, wherein the space has an allocation size, and wherein the search comprises:
 - comparing the allocation size with the nodes based on the size; and
 - selecting the one or more nodes to reduce fragmentation effect in allocating the identified available blocks in the storage devices.
12. A computer implemented method for a file system, the method comprising:
 - maintaining tree representations for a bitmap stored in storage devices having blocks for storing data, each bit in the bitmap indicating availability of one of the blocks, each node in the tree representations corresponding to an extent of available blocks;
 - in response to receiving an allocation request for a file, identifying one or more matching nodes in the tree representations according to a search criteria for the allocation request;

traversing the bitmap from the storage devices if the matching nodes cannot be identified; and
building up the tree representations based on the traverse of the bitmap.

13. The method of claim 12, wherein the tree representations correspond to extends of available blocks, and wherein the bitmap includes at least one bit indicating an available block not covered in the extends of available blocks.

14. The method of claim 12, wherein the maintenance comprises:
loading a portion of the bitmap from the storage devices; and
generating the tree representation in memory from the portion of the bitmap.

15. The method of claim 12, wherein the storage devices include available tree representations stored for the bitmap, wherein the maintenance comprises:
determining if the available tree representations represent the bitmap; and
retrieving the available tree representations from the storage devices as the tree representations in memory if the available tree representations represent the bitmap.

16. The method of claim 12, wherein the bitmap includes an array of bits, wherein the extent of available blocks corresponds to consecutive bits in the array indicating the available blocks, and wherein the consecutive bits are bounded by two bits in the array indicating two blocks which are not available.

17. The method of claim 12, wherein the node includes a location for a first bit of the consecutive bits in the bitmap, and wherein the node includes a number of the consecutive bits.

18. The method of claim 17, wherein the allocation request includes an amount of space for allocation, wherein the search criteria is based on a minimum number

of blocks having a capacity large enough to accommodate the amount of space, and wherein each matched node is associated with the minimum number.

19. The method of claim 18, wherein the allocation request includes an existing location for the file, wherein the search criteria is based on the existing location and the minimum number, and wherein the matched node is associated with a matching location and a matching number, the matching location being near the existing location and the matching number being greater than with the minimum number.

20. The method of claim 17, wherein the identification comprises:
selecting one of the tree representations; and
performing a red-black tree search on the selected tree representation.

21. The method of claim 20, wherein the red-black tree search is based on the location.

22. The method of claim 20, wherein at least a portion of the bitmap is not covered in the selected tree representation.

23. The method of claim 12, wherein building up the tree representations comprises:
determining if the tree representations have been fully populated by the bitmap for the storage devices, wherein each available block belongs to one of the extents of available blocks for the tree representations.

24. A machine-readable storage medium having instructions, when executed by a machine, cause the machine to perform a method for a file system, the method comprising:
selecting, in response to a request to allocate space in storage devices for a file, one of a plurality of tree representations of a bitmap, each bit

in the bitmap indicating availability of a block in the storage devices, each node in the tree representations corresponding to one or more consecutive bits in the bitmap;
searching the selected tree representation for one or more nodes identifying blocks available for the request;
updating the bitmap to allocate the identified blocks in the storage devices;
and
synchronizing the plurality of trees with the updated bitmap.

25. An apparatus, comprising:

a memory storing executable instructions including a file system;
a storage device having blocks for storing data, the storage storing a bitmap, each bit in the bitmap indicating availability of one of the blocks;
a processor coupled to the storage and the memory to execute the executable instructions from the memory for the file system, the processor being configured to:
 maintain tree representations for the bitmap, each node in the tree representations corresponding to an extent of available blocks;
 in response to receiving an allocation request for a file, identify one or more matching nodes in the tree representations according to a search criteria for the allocation request;
 traverse the bitmap from the storage devices if the matching nodes cannot be identified; and
 build up the tree representations based on the traverse of the bitmap.

1/5

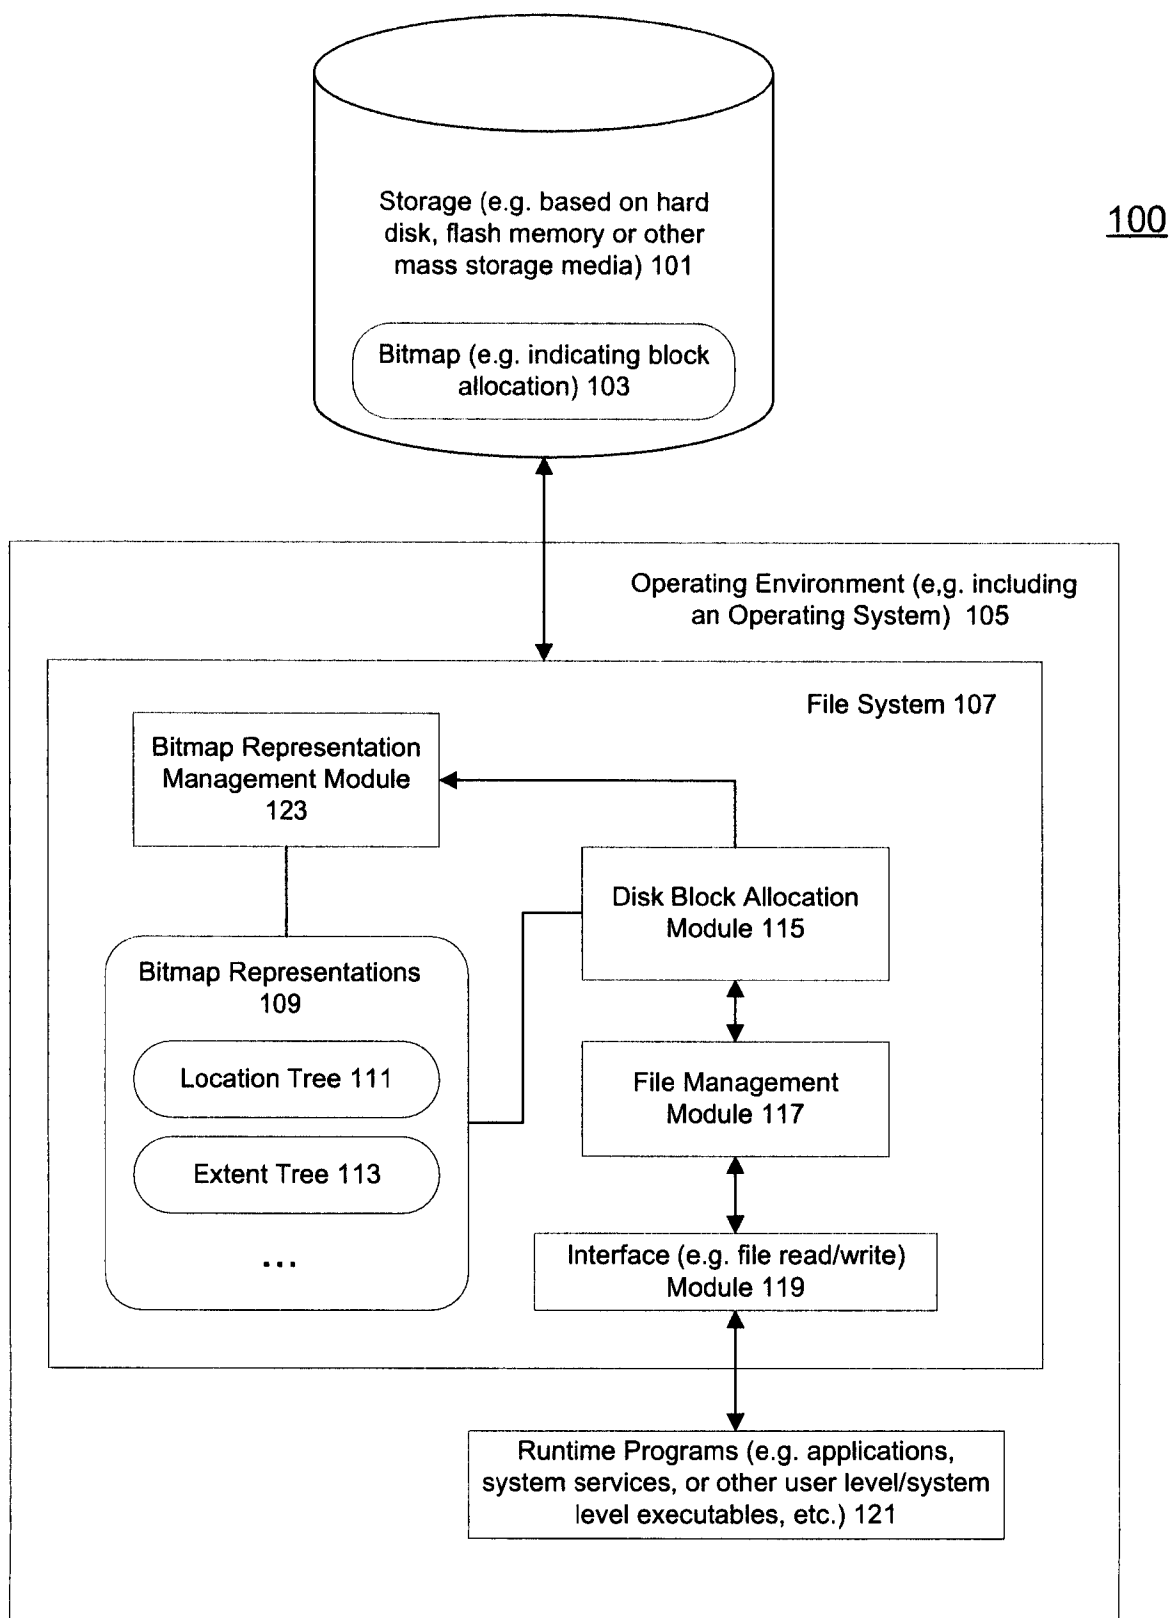


Fig. 1

2/5

200

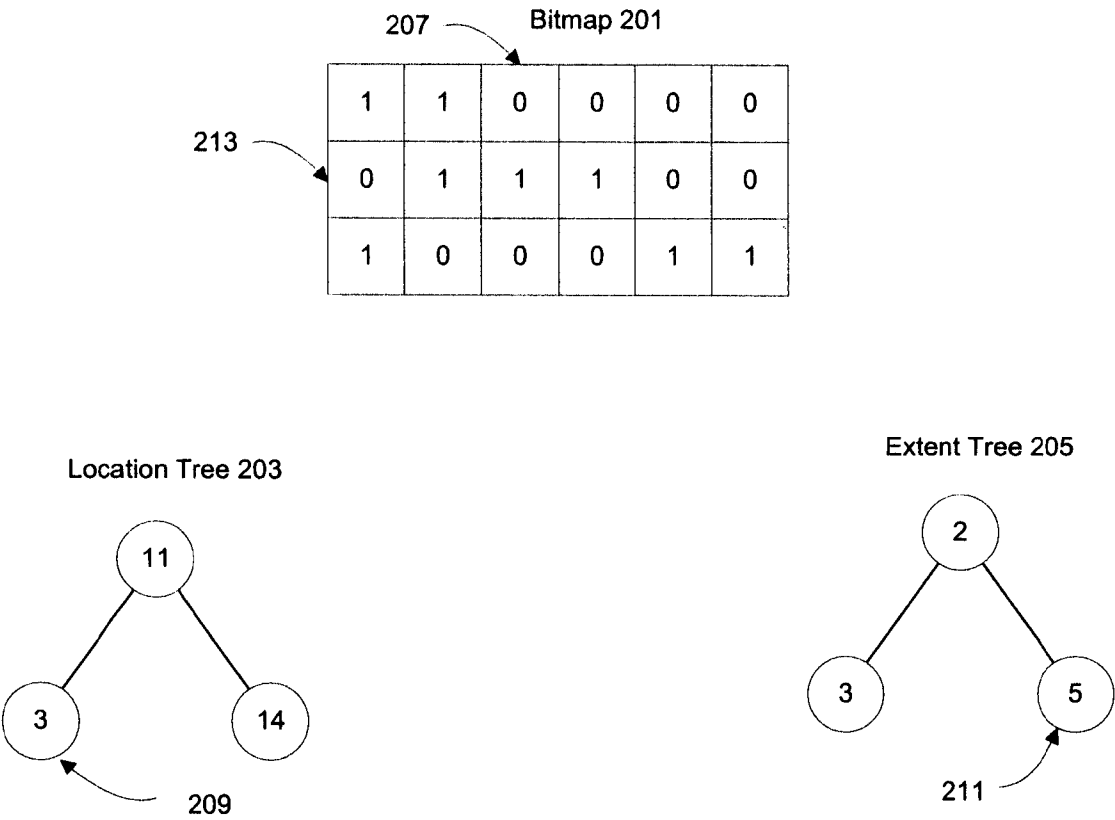
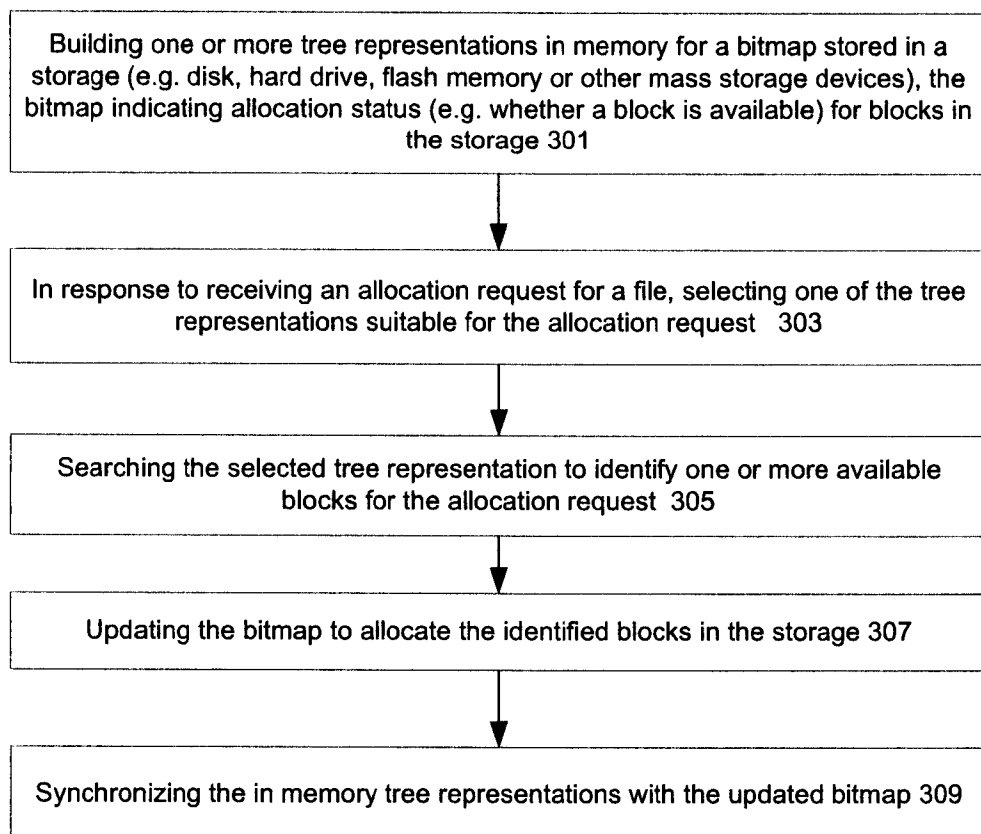
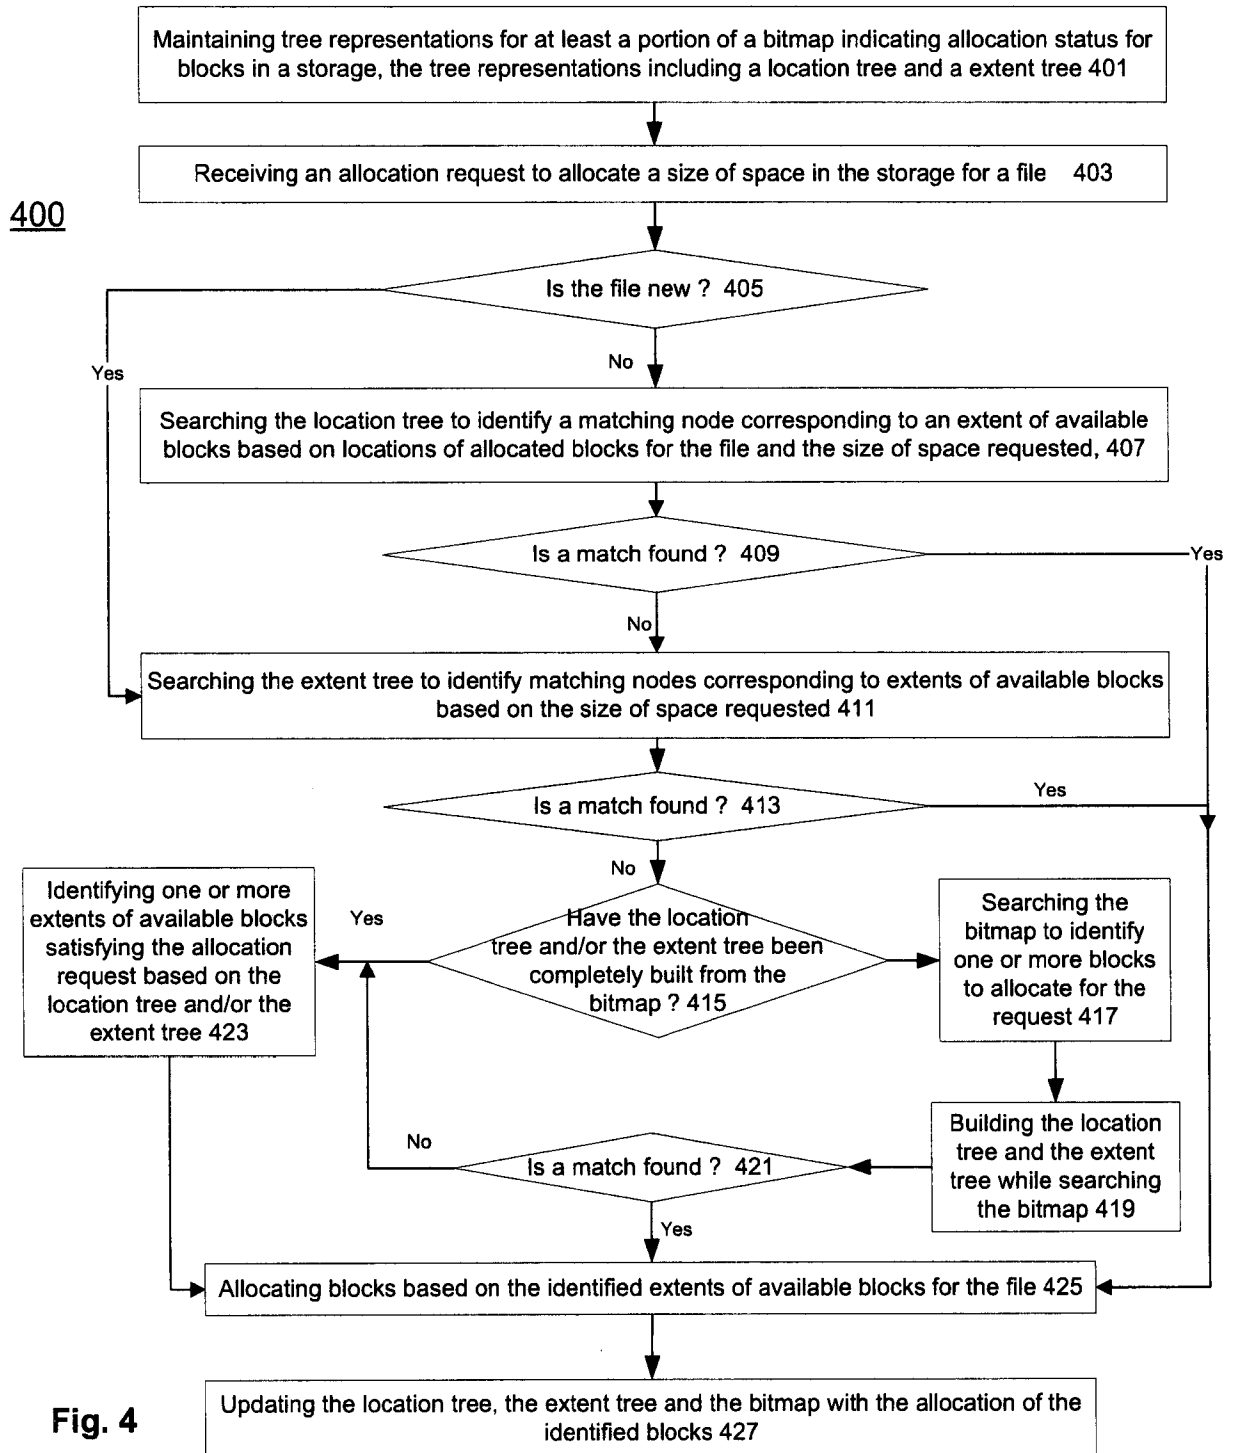


Fig. 2

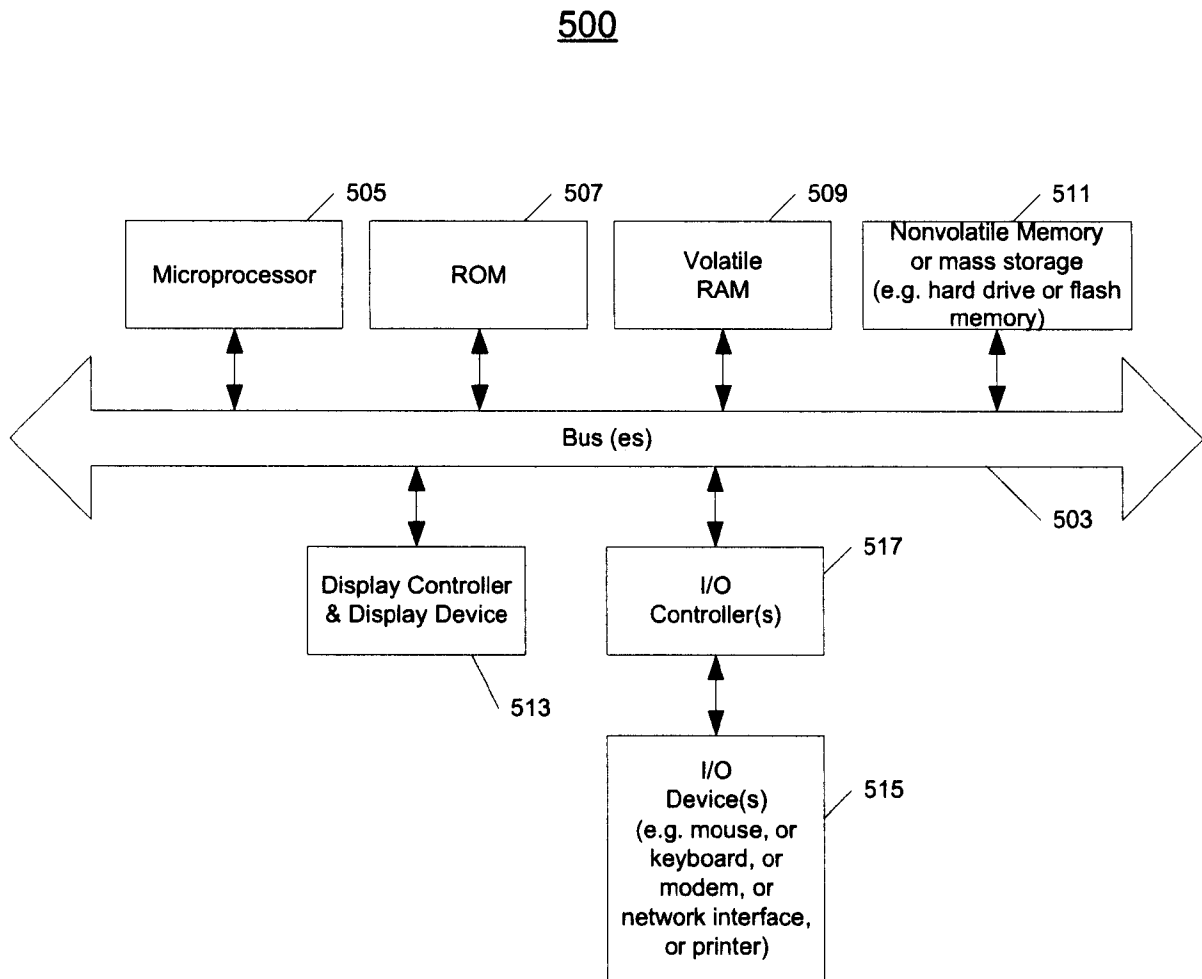
3/5

300**Fig. 3**

4/5



5/5

**Fig. 5**

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2010/059298A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F3/06 G06F17/30
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F

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 practical, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5 375 233 A (KIMBER SUSAN P [GB] ET AL) 20 December 1994 (1994-12-20) figures 1-4 column 1, line 9 - column 4, line 68 -----	1-25



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"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)

"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

"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

"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

"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.

"&" document member of the same patent family

Date of the actual completion of the international search

11 February 2011

Date of mailing of the international search report

18/02/2011

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Andlauer, J

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2010/059298

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5375233	A	20-12-1994	NONE
