

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
23 December 2004 (23.12.2004)

PCT

(10) International Publication Number
WO 2004/111881 A1

(51) International Patent Classification⁷: G06F 17/30, 12/00

(21) International Application Number: PCT/US2004/017259

(22) International Filing Date: 28 May 2004 (28.05.2004)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data: 10/447,205 28 May 2003 (28.05.2003) US

(71) Applicant (for all designated States except US): PERVASIVE SOFTWARE, INC. [US/US]; 12385 Riata Trace Parkway, Building B, Austin, TX 78727 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): REINAUER, Rob [US/US]; 3008 Toro Ring, Austin, TX 78746 (US). WHITE, Ken [US/US]; 104 Oakmont Forest Drive, Cedar Park, TX 78613 (US). SUN, Chunsheng [CN/US]; 8012 Whitworth Lane, Round Rock, TX 78681 (US).

ARNOLD, Richard [US/US]; 4814 Shoal Creek, Austin, TX 78756 (US). JACOB, Sunil [IN/US]; 13109 Troops Trail, Austin, TX 78727 (US). TAN, Desmond [SG/US]; 604 Kai Drive, Cedar Park, TX 78613 (US). LEWIS, Kevin [US/US]; 8541 Copano Drive, Austin, TX 78749 (US).

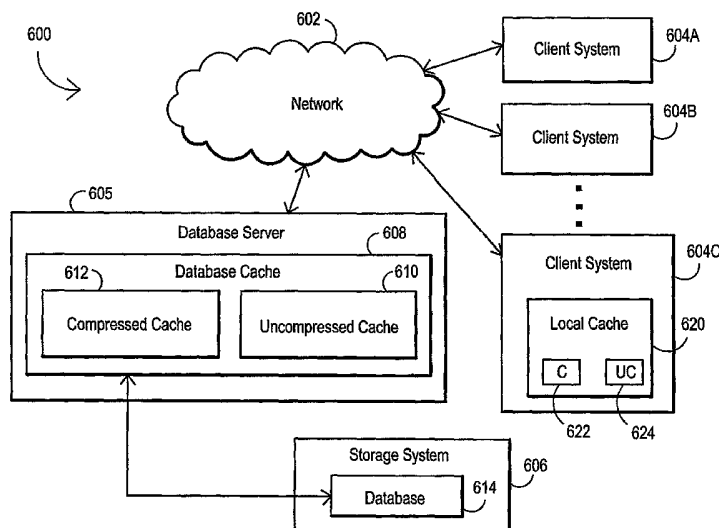
(74) Agent: PETERMAN, Brian, W.; 1101 Capital of Texas Hwy. So., Building C, Suite 200, Austin, TX 78746 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, 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, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR UTILIZING COMPRESSION IN DATABASE CACHES TO FACILITATE ACCESS TO DATABASE INFORMATION



(57) Abstract: A system and method are disclosed for utilizing compression in database caches to facilitate access to database information. In contrast with applying compression to the database that is stored on disk, the present invention achieves performance advantages by using compression within the main memory database cache used by a database management system to manage data transfers to and from a physical database file stored on a storage system or stored on a networked attached device or node. The disclosed system and method thereby provide a significant technical advantage by increasing the effective database cache size. And this effective increase in database cache size can greatly enhance the operations-per-second capability of a database management system by reducing unnecessary disk or network accesses thereby reducing data access times.

WO 2004/111881 A1



European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*

Published:

— *with international search report*

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.

**SYSTEM AND METHOD FOR UTILIZING COMPRESSION IN DATABASE CACHES
TO FACILITATE ACCESS TO DATABASE INFORMATION**

Inventors: Rob Reinauer, Ken White, Chunsheng Sun, Richard Arnold,
Sunil Jacob, Desmond Tan and Kevin Lewis

TECHNICAL FIELD OF THE INVENTION

[0001] The present invention relates to the management of data and more particularly to database management, for example, in networked client-server environments.

BACKGROUND OF THE INVENTION

[0002] In prior systems, cache memory has been used to improve performance of central processing units (CPUs) in computer systems. This cache memory is typically small in size as compared to main memory and is used to hold a segment of main memory for use by the CPU in its operations. To the extent the CPU can use instructions in cache memory without having to pull new information from main memory, CPU performance is typically enhanced. For this reason, it is often desirable to increase the cache memory size. Limitations exist, however, that hinder the ability to add more physical memory. For example, many operating systems limit how much main memory and cache memory a system can physically access. To increase the effective size of CPU cache memory, therefore, prior solutions have proposed the use of cache compression within the cache memory.

[0003] With respect to database environments, a portion of the main memory for the computer system managing the database is often used by the database management application as a database cache for data being read from or written to the stored database file. The database cache provides a buffer between the access, create and modify instructions from the database users and the database file itself. In addition, the database cache can provide for improved access times to client systems to the extent that the database management application can satisfy queries to the database from the data currently sitting in the database cache. The database file is typically stored on some physical media, such as one or more hard disks. With respect to the storage of large database files, prior work has focused on using data compression algorithms to reduce the size of the database files stored on hard drives. In addition, because most existing database access protocols operate on uncompressed data, prior work has also focused on protocols and query methodology for directly accessing the compressed data where the database file is compressed on disk. However, with increases in the speed

of CPUs outpacing the speed of disk access, this disk compression can provide only limited improvement due to the physical limitations related to accessing data from a physical disk.

SUMMARY OF THE INVENTION

[0006] The present invention provides a system and method for utilizing compression in database caches to facilitate access to database information. In contrast with applying compression to the database that is stored on disk, the present invention achieves performance advantages by using compression within the main memory database cache used by a database management system to manage data transfers to and from a physical database file stored on a storage system or stored on a networked attached device or node. As discussed herein, the present invention provides a significant technical advantage by increasing the effective database cache size. And this effective increase in database cache size can greatly enhance the operations-per-second capability of a database management system by reducing unnecessary disk or network accesses thereby reducing data access times.

[0007] In part, the present invention provides a solution that substantially eliminates or reduces disadvantages and problems associated with previously developed database cache management systems. More specifically, the present invention provides systems and methods for managing data in database caches. This system includes a first data storage location. This first data storage location typically comprises a disk or network resource. Additional data storage locations, typically in the form of local or cache memory, allow database management systems to more quickly access frequently used data from a database cache as opposed to disk. To increase the effective size of the database cache, data stored within the local memory may be compressed. This compression typically becomes desirable when decompressing compressed data from the database cache and supplying this data to the data user can occur more quickly than accessing the original data from disk. The database cache can include both data and instructions, and the data may be formatted into pages or other like data structures as known to those skilled in the art. Furthermore, these pages (data structures) may include one or more flags that indicate whether the page has been compressed.

[0008] Other embodiments may further optimize the present invention by utilizing a local uncompressed database cache and a local compressed database cache. In such an embodiment, data pages, such as the least recently used pages within the local database cache, move to the local compressed database cache when the local uncompressed database cache is full. This use of compressed and uncompressed database caches can enhance performance because the performance penalty for compression/decompression is typically significantly less than disk access costs even when the compression penalty is incurred multiple times. By utilizing an uncompressed cache for the

pages most likely to be re-used, the system can trade the fewest number of compression efforts for the highest number of disk reads. In certain embodiments, the pages within the compressed database cache may, in fact, only be compressed if the page actually compresses when compressed and/or when the compressed database cache has become full. This procedure can be used to avoid unnecessary compression and decompression, thus effectively increasing system resources.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] It is noted that the appended drawings illustrate only exemplary embodiments of the invention and are, therefore, not to be considered limiting of its scope, for the invention may admit to other equally effective embodiments.

[0010] FIGURE 1 graphically depicts the cache curve demonstrating the benefit of effectively increasing the size of database cache.

[0011] FIGURE 2 depicts one general embodiment of the system and method provided by the present invention;

[0012] FIGURE 3 illustrates an embodiment of the present invention that utilizes two cache memories;

[0013] FIGURE 4 illustrates other embodiments of the present invention utilizing two cache memories;

[0014] FIGURE 5 depicts various structures within a block of compressed data.

[0015] FIGURE 6 is a block diagram for a client-server database environment in which a database server utilizes compression within its database cache.

DETAILED DESCRIPTION OF THE INVENTION

[0016] The present invention compresses data within a database cache in order to effectively increase the size of database cache. Increased processor speeds make this possible as the time required to compress and decompress data and/or instructions with increased processor speed is less than the time required to access the data from disk. Previously, the time spent to compress and decompress data and/or instructions exceeded the time required to access and retrieve data and/or instructions from disk. The present invention couples increased processor speeds with high

performance compression algorithms as known to those skilled in the art, allowing data and/or instructions to be compressed and decompressed faster than the time required for disk access.

[0017] FIGURE 1 graphically depicts the improved performance of a system based on a percentage of data caches, assuming a standard hard disk drive is used for persistent storage. Cache curve 12 shows that the benefit increases exponentially. This benefit increases exponentially as the number of operations per second on the y axis increase. Improved performance comes from the elimination of disk access times that may be associated with each operation. Cache curve 12 shows a change from 10 to 20 percent does not yield a large increase in the number of operations per second. However, an increase from 80 to 100 percent yields a much greater increase in operations per second. The Y-axis on FIGURE 1 has been left without a scale because the actual numbers are dependent on a number of factors. These factors include the time to acquire a given piece of information from disk, the time to acquire the same piece of information from cache, and the time required to process that piece of information on behalf of the client (where this time is considered the same regardless of the source of the piece of information). Given these values, the Y-axis value for a given percentage of the total information in cache for a given time period can be calculated using the following formula, where "td" represents the time to acquire information from disk, "tm" represents the time to acquire information from memory, "tp" represents the time to process a given piece of information, and "tt" represents the total time taken to run the test:

$$Y = tt / ((X * tm) + ((100 - x) * td) + tp)$$

Assuming current hard drive technology, which is multiple orders of magnitude slower than typical RAM access times and a reasonable "tp" value, as X increases, Y will increase exponentially. The reason for this relationship is the exponential disparity between "tm" and "td."

[0018] The method and system of the present invention add incremental costs to the access of cache through various processes. Primarily, compression, decompression and cache management all add incremental costs. These costs previously exceeded those associated with disk I/O access. However, increased processor speeds have greatly reduced these costs while disk I/O costs remain relatively unchanged.

[0019] The present invention provides a system and method of effectively increasing database cache or local memory for database management systems. This greatly improves the performance of many database applications. FIGURE 2 generally illustrates how the present invention handles data and/or instructions. This data and/or instructions are contained within pages. For example, when user 26 requests access to page 22, cache manager 28 causes page 22 to be accessed from disk 24.

Compression/decompression algorithm 30 compresses page 22 within cache 32. Cache manager 28 directs compression/decompression algorithm 30 to decompress page 22 from cache 32, then page 22 is provided to user 26. Whenever page 22 is accessed by user 26, user 26 accesses page 22 from cache 32 and decompresses page 22 with a compression/decompression algorithm.

[0020] If desired, as shown in FIGURE 3, two or more database caches may be used by a database management system, according to the present invention. Here, page 22, when first accessed, resides within uncompressed cache 34. This embodiment uses two caches: uncompressed cache 34 and compressed cache 36. Users always read data and/or instructions from uncompressed cache 34. Hence, if page 22 has been compressed, the compression/decompression algorithm 40 decompresses page 22 and delivers the page to uncompressed cache 34. Compression/decompression algorithm 40 compresses the least recently used (LRU) pages from uncompressed cache 34 to compressed cache 36. These actions are directed by cache manager 41. Other cache management strategies may be utilized, if desired, including cache management strategies such as a least-recently-used cache management strategy in which relative "ages" of cache information is kept and the information that has remained unused for the longest time is replaced, a least-frequently-used cache management strategy in which the number of times information has been used over some number of uses or period of time and the information that is least used is replaced, and a first-in-first-out cache management strategy in which the first information added to the database cache is the first information to be replaced.

[0021] FIGURE 4, similar to FIGURE 3, again illustrates an embodiment of the present invention containing both an uncompressed cache 46 and a compressed cache 48. Here, page 42 initially resides on disk 44. When needed by user application 50, cache manager 52 directs that page 42 be retrieved from disk 44 and stored within uncompressed cache 46. Additional pages are stored in an uncompressed cache 46 until uncompressed cache 46 has been filled. As uncompressed cache 46 fills, cache manager 48 directs that the LRU page, page 54, stored within uncompressed cache 46, to be compressed via compression/decompression algorithm 56 and stored within compressed cache 48. LRU page 54 remains stored within compressed cache 48 until needed. When needed, compression/decompression algorithm 56 decompresses the page, which is then stored within uncompressed cache 46. When compressed cache 48 is full, cache manager 52 directs that the LRU page, page 58, within compressed cache 48 to be deleted or written over as represented by "trash" block 59. The decision to push the LRU pages from uncompressed cache 46 to compressed cache 48, and to delete the LRU from compressed cache 48 comes from the theory that users are more likely to access a recently accessed pages. This theory avoids the costs associated with repeatedly compressing and uncompressing frequently used pages. Although the present invention uses an LRU cache management technique, any cache management technique known to those skilled in the art may

be used in place of this technique. When a requested page has been deleted from compressed cache 48, that page must be read from disk 44.

[0022] The present invention also provides the ability to compress and uncompress asynchronously. Both compression and decompression provide cache management costs. Moving data and/or instructions from one cache to another involves updating pointers within the cache manager. Furthermore, the compression and decompression require processor time. Compressing asynchronously queues uncompressed pages for compression when processor time becomes available. Asynchronous decompression similarly queues pages for decompression but requires predictive read ahead.

[0023] Operating systems often already use predictive read ahead when accessing files. This type of read ahead assumes the user will request the page following that being viewed. The drawback associated with predictive read ahead occurs when the user does not request what was predicted, thus requiring additional resources to be expended.

[0024] Another embodiment of the present invention addresses the problem with compressed pages that compress longer than the original page (i.e., the page actually expands). This problem occurs within all compression techniques. Compression does not guarantee to compress every page.

[0025] In some instances, the cache management technique may store pages that compress larger than their uncompressed size only within the uncompressed cache. This feature reduces or eliminates wasted memory but still consumes processing resources.

[0026] Cache compression enables cache to store more data pages within the cache. For example, if all data pages can be compressed to at least 1/2 their original size, with the size of the cache being constant, the cache can hold twice as many pages. Although having to compress/decompress pages adds overheads (i.e., CPU utilization), this increase in overhead is small when compared to disk I/O access costs.

[0027] There are a wide variety of different compression algorithms that can be used to compress data including data in a database cache. For example, a Lempel-Ziv (LZ) compression algorithm provide a technique that encodes a streaming byte sequence using a dynamic table. Popular variations on this technique include the LZ78 and LZ77 families. It is noted that other compression algorithms could be used, as desired. Preferably, a high performance compression/decompression algorithm is used so that the processing overhead for the algorithm does not outweigh the access time benefits provided by the compressed database cache approach of the present invention. In addition, any of a number of well-understood variants of the LZ algorithm may

be utilized, as desired, depending upon the particular application, as would be understood by one of skill in the art.

[0028] To improve the compression/decompression algorithms performance, the algorithms may be modified with the following abilities. First, the algorithm does not compress the data page when the compression ratio is less than 2, and the page may be flagged to eliminate future attempts to compress the same page. Second, the algorithm writes out the compressed data in predetermined sizes such as 256 bytes (256B). For example, if a 1K data page compresses to 356B, the algorithm, when compressing, writes the first 256B compressed data in one chunk and the remainder of the 100B in another chunk. It is noted that a trade-off exists with respect to the compression block size, and other compression block sizes may be utilized, as desired. In particular, it is noted that the smaller the compression block size, the better able the system typically is in taking advantage of reduced data size, but the more overhead the system will typically incur in managing the compressed blocks. Third, the algorithm provides a pointer to a compressed data header object that provides information about the compressed data pages.

[0029] FIGURE 5 depicts the various structures that compressed data header object 70 tracks within the pieces of compressed data. As shown in this embodiment, the decompression algorithm gathers related chunks of 256B compressed data and decompresses them into their original size.

[0030] The integration of cache compression into an existing application may require modifying some structures within the existing application. The main part of the integration work occurs within the two caches. The structures identified begin with the page pointer (page_ptr) 72. Page pointer 72 may include an additional pointer 74 returned by the compression algorithm when the compression is successful. If pointer 74 is NULL, the data page is not compressed.

[0031] As mentioned above, the compressed cache receives a pool of memory 76. The different pools in the compressed cache include a pool of 256B objects 78 that hold compressed data. Compressed data headers 70 serve as the map table or data pointer list for finding the different chunks of related compressed data. In the embodiment depicted, each header is limited to four 256B objects. It is noted, however, that this size could be altered, as desired, without departing from the present invention. And in addition to being any fixed number of blocks, this header could also contain a structure containing a variable number of blocks, such as a vector or linked list.

[0032] Synchronous compression and decompression manipulates the data on request. When a thread moves a page from uncompressed cache to compressed cache, the thread invokes the compression function to compress the data page directly into the compressed cache pool of objects.

Decompression occurs in the same manner. During decompression, first the correct size buffer is located in the uncompressed cache, after which pieces of related compressed data located in compressed cache decompress directly into the uncompressed cache. Compressing on demand requires no additional memory copying, reducing the amount of overhead. The disadvantage during heavy paging situations between uncompressed and compressed caches occurs when the data pages do not meet the required compression ratio. Failed compression (compressions thresholds in expansion) is an overhead in addition to the memory copy functions that needs to be invoked.

[0033] Asynchronous operation is typically best suited for compression. Data pages move from uncompressed cache to compressed cache, after which the compression function operates to compress the data. The advantage being that compression can happen any time. If heavy paging situations occurs, the overhead incurred is only the memory copy function. When less busy, a thread in the compressed cache can start the compression of the data pages queued for compression. This situation will typically invoke an additional memory copy function.

[0034] As noted above, compression and decompression adds a certain amount of overhead to the workings of an application. However, the idea trades this overhead for I/O by storing more data pages in cache. Not all database environments will experience a performance boost by having cache compression. Situations, which involve small size databases and CPU-bound systems, actually may experience a negative impact on performance with cache compression. Therefore, a setting for cache compression may be made available for the better-informed user within the application configuration or setup.

[0035] Returning to FIGURE 4, in this and other embodiments of the present invention cache compression does not occur until all the available primary cache has been filled with uncompressed pages. This avoids compression and decompression until needed.

[0036] Compression of the secondary cache may begin when the system becomes I/O bound or may be considered the permanent state of the secondary cache. In this way, if an entire database fits into the total available cache without compression, the processor costs associated with compression and decompression are avoided automatically. This may also imply that an asynchronous thread might try to compress all non-compressed pages before freeing the LRU pages. It would do this in order based on a cache management algorithm, such as MRU/LRU.

[0037] In summary, the present invention provides a system and method that substantially eliminates or reduces disadvantages and problems associated with previously developed database management systems. This system includes a first data storage location. Where this first data storage

location typically is a disk or network resource. Additional data storage locations, typically in the form of local or database cache, allow data users to more quickly access frequently used data from local or database cache as opposed to disk.

[0038] To increase the effective size of the database cache stored in memory, the present invention compresses data stored within the database cache. This compression only now becomes desirable as decompressing compressed data from the local memory and supplying decompressed compressed data to the data user can occur more quickly than accessing the original data from disk. Thus, the present invention further includes a processor and instructions operable to decompress the compressed data more quickly than the time required to access the information or data from non-local memory.

[0039] FIGURE 6 is a block diagram for a client-server database environment 600 in which a database server 605 utilizes compression within its database cache 608 to manage the database. In the embodiment depicted, one or more client systems 604A, 604B ... 604C are connected through network 602 to a server-based database management system 605. The database cache 608 provides an interface between the client systems 604A, 604B ... 604C and the database 614 stored on the storage system 606. In its operations, the database management server 605 utilizes a database cache 608. As discussed above and according to the present invention, this database cache 608 includes a compressed cache 612 and an uncompressed cache 610. It is noted that in a client-server environment, the client systems 604A, 604B ... 604C can also utilize a local database cache. For example, client system 604C could include a local database cache 620 that provides an interface between the database related operations of the client system 604C and the database server 605. As such, the client system 604C could cache database information locally in its local database cache 620 thereby reducing the number of accesses the client system 604C needs to make to database server 605 and also reducing latency caused by network access through network 602. In addition, the client system 604C could also utilize compression with respect to its local database cache 620, according to the present invention. As such, the local database cache 620 would include a compressed cache 622 and an uncompressed cache 624. With respect to the database cache 608 and the local database cache 620, if utilized, the ratio of compressed to uncompressed cache can be selected, as desired, and the entire database cache can be compressed if this implementation is desired. In addition, a fixed ratio or a dynamic ratio could be used, as desired. It is noted that where a local database cache is used in addition to the server database cache, coherence between these two caches can be problematic. Solutions to this cache coherence problem are discussed, for example, in U.S. Patent Application No. 10/144,917, filed May 14, 2002, and entitled "SYSTEM AND METHOD OF MAINTAINING

FUNCTIONAL CLIENT SIDE DATA CACHE COHERENCE,” the entire text and all contents of which are hereby expressly incorporated by reference in its entirety.

[0040] Further modifications and alternative embodiments of this invention will be apparent to those skilled in the art in view of this description. It will be recognized, therefore, that the present invention is not limited by these example arrangements. Accordingly, this description is to be construed as illustrative only and is for the purpose of teaching those skilled in the art the manner of carrying out the invention. It is to be understood that the forms of the invention herein shown and described are to be taken as the presently preferred embodiments. Various changes may be made in the implementations and architectures for database processing. For example, equivalent elements may be substituted for those illustrated and described herein, and certain features of the invention may be utilized independently of the use of other features, all as would be apparent to one skilled in the art after having the benefit of this description of the invention.

CLAIMS

What is claimed is:

1. A system for managing access to database information, comprising:
a first data storage device configured to store a database;
a second data storage device configured to store a database cache, at least a portion of the database cache comprising a compressed cache portion, the database cache also including an uncompressed cache portion where the compressed cache portion is less than the full database cache; and
a database management system configured to control data transfers among the database, the compressed portion of the database cache and the uncompressed portion of the database cache, if any, to manage accesses to database information.
2. The system of claim 1, wherein the first data storage device comprises a disk drive.
3. The system of claim 1, wherein the second data storage device comprises memory within a computer system.
4. The system of claim 3, wherein the memory is configured to store database cache information in pages.
5. The system of claim 4, wherein each database cache page stored within memory are configured to include a flag indicating whether or not the page comprises compressed data.
6. The system of claim 1, wherein the entire database cache comprises compressed data and there is no uncompressed cache portion.
7. The system of claim 1, wherein the database management system comprises a server system within a client-server environment.

8. The system of claim 1, wherein the database management system is a client system within a client-server environment, wherein the first data storage device comprises a remote storage device coupled to a server system, and wherein the database cache is a local database cache utilized by the client system.
9. A database management system for managing access to database information in a client-server environment, comprising:
a plurality of client systems, the client systems configured to access information in a database;
a server system coupled to the client systems through a network, the server system configured to manage transfers of information between the client systems and the database;
a first data storage device coupled to the server system and configured to store the database; and
a second data storage device coupled to the server system and configured to store a database cache, at least a portion of the database cache comprising a compressed cache portion, the database cache also including an uncompressed cache portion where the compressed cache portion is less than the full database cache;
wherein the server system is configured to control data transfers among the database, the compressed portion of the database cache and the uncompressed portion of the database cache, if any, to manage accesses to database information by the client systems.
10. The database management system of claim 9, wherein the first data storage device comprises a disk drive.
11. The database management system of claim 9, wherein the second data storage device comprises memory within a computer system.
12. The database management system of claim 9, wherein the entire database cache comprises compressed data and there is no uncompressed cache portion.
13. The database management system of claim 9, wherein one or more of the client systems utilize a local database cache to manage database information accessed by the client system, the local database cache including at least in part a compressed cache portion.

14. A method for managing access to database information, comprising:
storing a database in a first data storage device;
storing a database cached in a second data storage device, at least a portion of the database cache comprising a compressed cache portion, the database cache also including an uncompressed cache portion where the compressed cache portion is less than the full database cache; and
controlling data transfers among the database, the compressed portion of the database cache and the uncompressed portion of the database cache, if any, to manage accesses to database information.
15. The method of claim 14, wherein the first data storage device comprises a disk drive.
16. The method of claim 14, wherein the second data storage device comprises memory within a computer system.
17. The method of claim 16, configuring the memory to store database cache information in pages.
18. The method of claim 17, further comprising utilizing for each stored database cache page a flag within the page to indicate whether or not the page comprises compressed data.
19. The method of claim 14, storing the entire database cache as compressed data such that there is no uncompressed cache portion.
20. The method of claim 14, wherein a server system within a client-server environment performs the controlling step.




FIG. 1

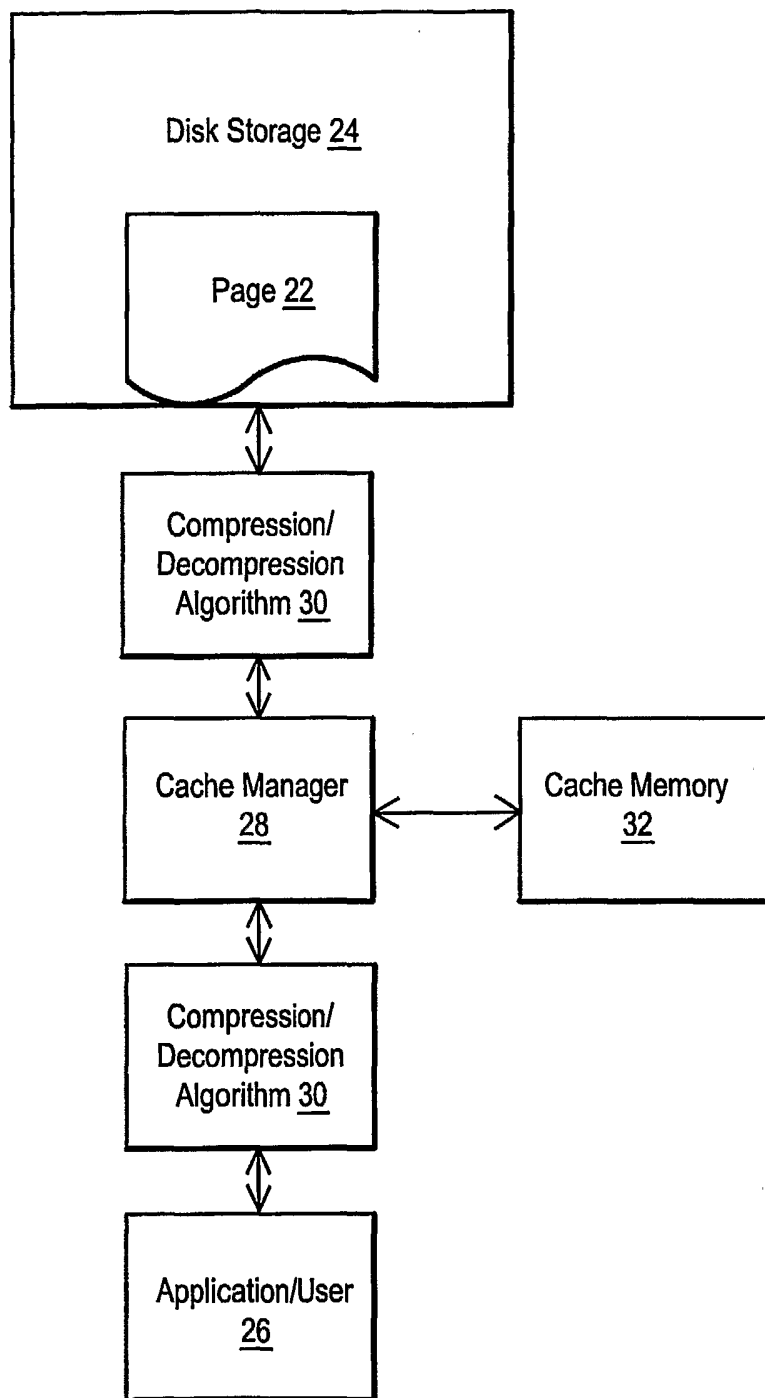


FIG. 2

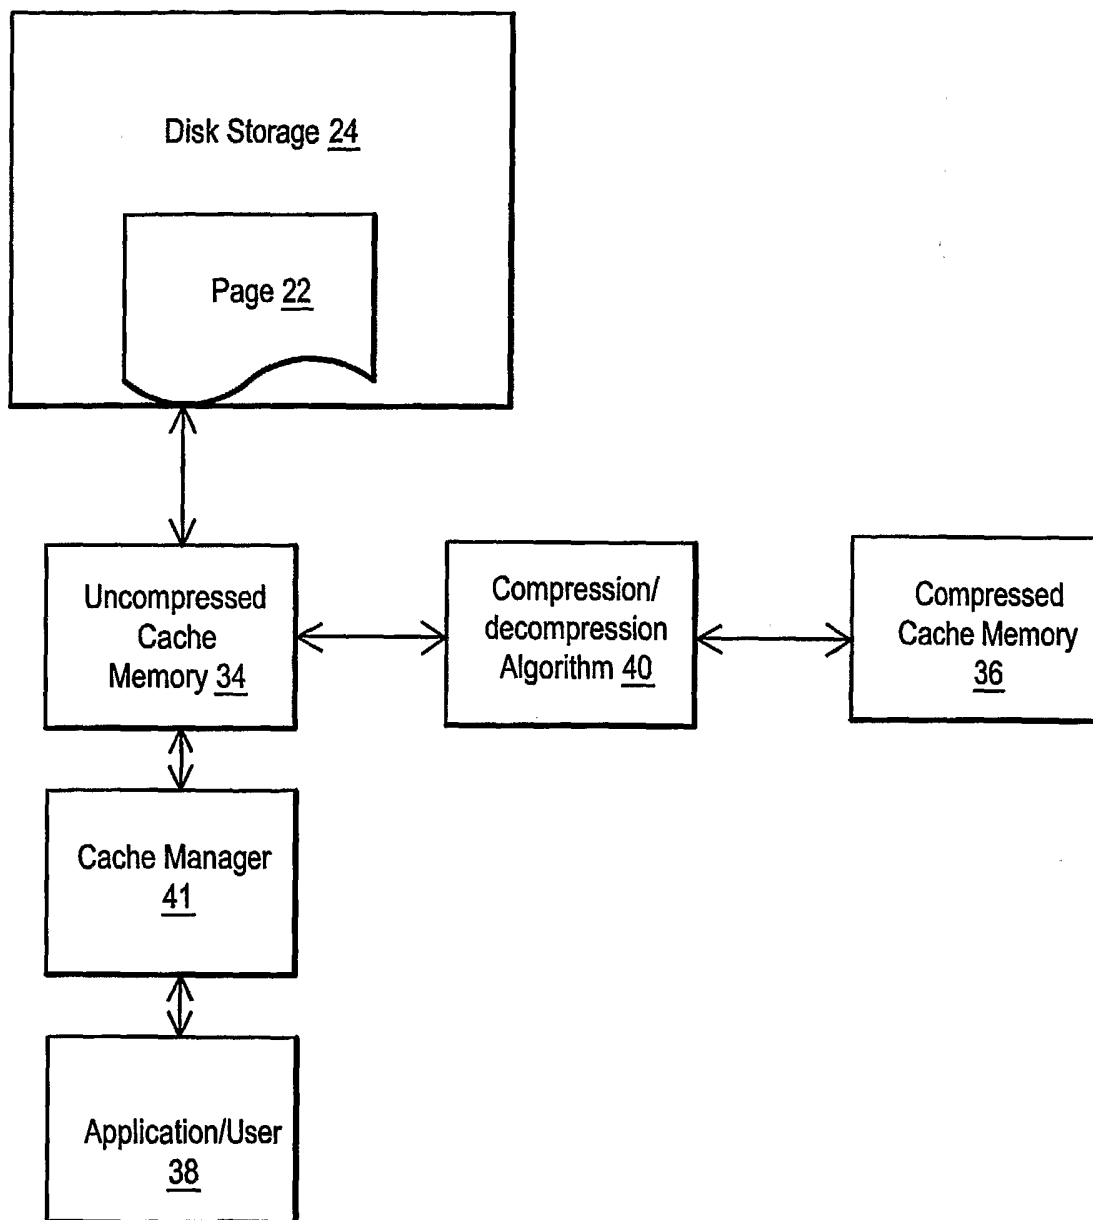


FIG. 3

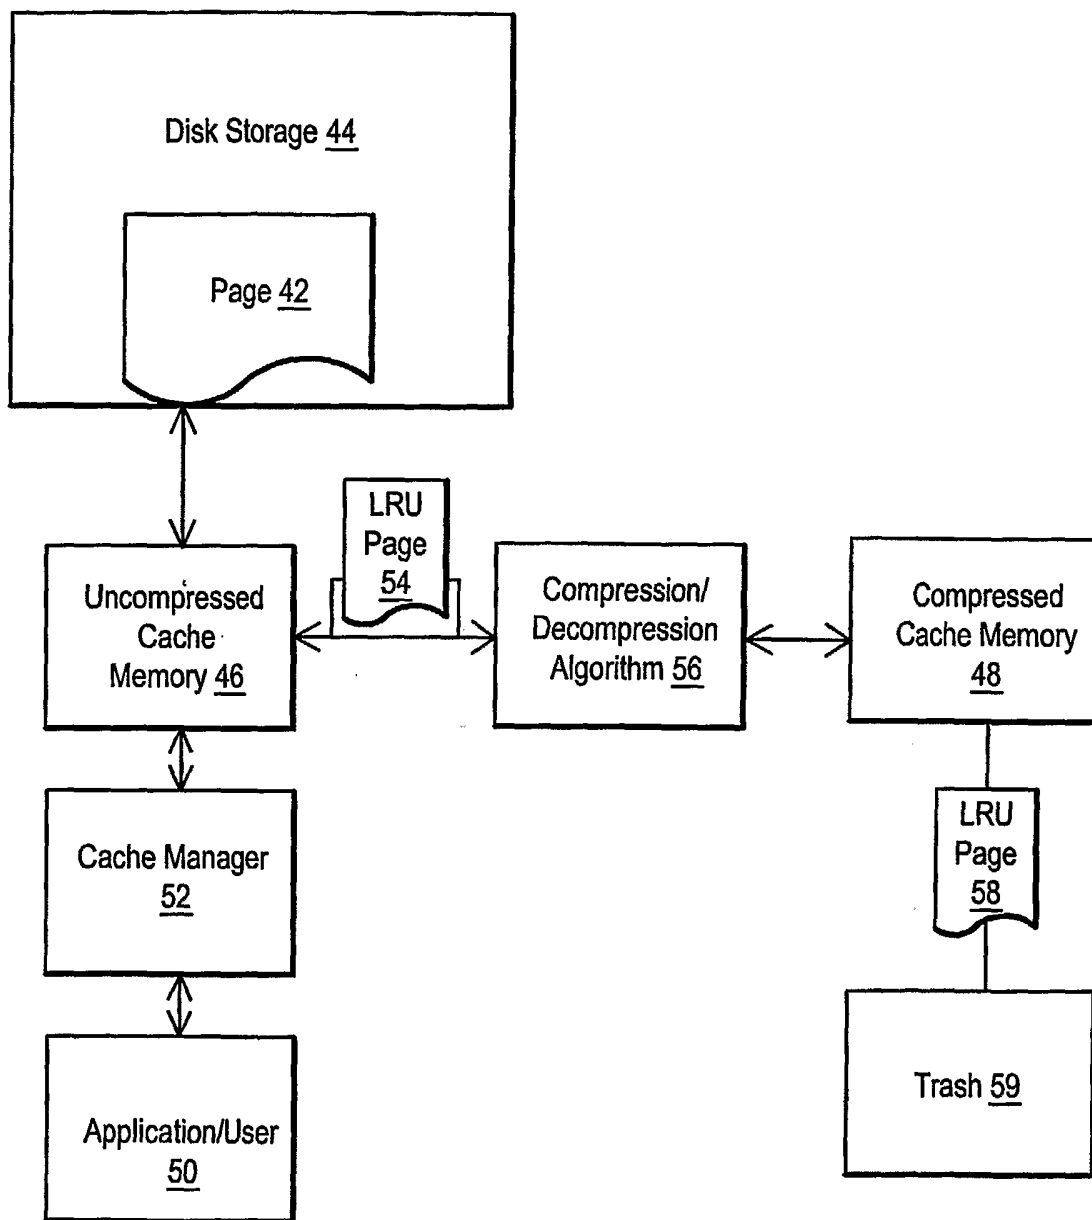


FIG. 4




FIG. 5




FIG. 6

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US2004/017259

A. CLASSIFICATION OF SUBJECT MATTER IPC 7 G06F17/30 G06F12/00				
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) IPC 7 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, INSPEC				
C. DOCUMENTS CONSIDERED TO BE RELEVANT				
Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.		
X	CATE V ET AL: "COMBINING THE CONCEPTS OF COMPRESSION AND CACHING FOR A TWO-LEVEL FILESYSTEM" COMPUTER ARCHITECTURE NEWS, ASSOCIATION FOR COMPUTING MACHINERY, NEW YORK, US, vol. 19, no. 2, 1 April 1991 (1991-04-01), pages 200-211, XP000203262 ISSN: 0163-5964 page 208, left-hand column, paragraph 1 page 201, left-hand column, paragraph 3 abstract page 202, right-hand column, paragraph 2 - paragraph 3 ----- -/--	1-20		
<table style="width: 100%; border: none;"> <tr> <td style="width: 50%; border: none;"> <input checked="" type="checkbox"/> Further documents are listed in the continuation of box C. </td> <td style="width: 50%; border: none;"> <input checked="" type="checkbox"/> Patent family members are listed in annex. </td> </tr> </table>			<input checked="" type="checkbox"/> Further documents are listed in the continuation of box C.	<input checked="" type="checkbox"/> Patent family members are listed in annex.
<input checked="" type="checkbox"/> Further documents are listed in the continuation of box C.	<input checked="" type="checkbox"/> Patent family members are listed in annex.			
° Special categories of cited documents :				
<table style="width: 100%; border: none;"> <tr> <td style="width: 50%; border: none;"> *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 </td> <td style="width: 50%; border: none;"> *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 </td> </tr> </table>			*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
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 <p style="text-align: center;">7 October 2004</p>	Date of mailing of the international search report <p style="text-align: center;">21/10/2004</p>			
Name and mailing address of the ISA European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016	Authorized officer <p style="text-align: center;">Jaedicke, M</p>			

INTERNATIONAL SEARCH REPORT

national Application No
PCT/US2004/017259

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>WILSON P R ET AL: "The case for compressed caching in virtual memory systems" PROCEEDINGS OF THE 1999 USENIX ANNUAL TECHNICAL CONFERENCE USENIX ASSOC BERKELEY, CA, USA, 6 June 1999 (1999-06-06), - 11 June 1999 (1999-06-11) pages 101-116, XP002299640 ISBN: 1-880446-33-2 reprinted without page numbers on 16 pages plus cover sheet abstract page 104, right-hand column, paragraph 3 - page 106, right-hand column, paragraph 4 page 114, left-hand column, last paragraph - page 115, left-hand column, last paragraph</p>	1-20
X	<p>GRAEFE G ET AL: "Data compression and database performance" APPLIED COMPUTING, 1991., YPROCEEDINGS OF THE 1991 SYMPOSIUM ON KANSAS CITY, MO, USA 3-5 APRIL 1991, LOS ALAMITOS, CA, USA, IEEE COMPUT. SOC, US, 3 April 1991 (1991-04-03), pages 22-27, XP010022657 ISBN: 0-8186-2136-2</p>	6,12,19
A	<p>page 22, left-hand column, paragraph 3 page 23, left-hand column, paragraph 2 - paragraph 3</p>	
X	<p>COCKSHOTT W P ET AL: "Data compression in database systems" DATABASE AND EXPERT SYSTEMS APPLICATIONS, 1998. PROCEEDINGS. NINTH INTERNATIONAL WORKSHOP ON VIENNA, AUSTRIA 26-28 AUG. 1998, LOS ALAMITOS, CA, USA, IEEE COMPUT. SOC, US, 26 August 1998 (1998-08-26), pages 981-990, XP010296751 ISBN: 0-8186-8353-8</p>	6,12,19
A	<p>page 988, right-hand column, paragraph 1 page 983, left-hand column, paragraph 4 - right-hand column, paragraph 2</p>	
	----- -/--	

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US2004/017259

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>FRENCH C D: "Teaching an OLTP database kernel advanced datawarehousing techniques" DATA ENGINEERING, 1997. PROCEEDINGS. 13TH INTERNATIONAL CONFERENCE ON BIRMINGHAM, UK 7-11 APRIL 1997, LOS ALAMITOS, CA, USA, IEEE COMPUT. SOC, US, 7 April 1997 (1997-04-07), pages 194-198, XP010218542 ISBN: 0-8186-7807-0 page 197, left-hand column, paragraph 3 - right-hand column, paragraph 5</p>	1-20
A	<p>US 5 794 228 A (WHITE PETER ET AL) 11 August 1998 (1998-08-11) column 3, line 10 - column 4, line 59 column 16, line 3 - line 21</p>	1-20
A	<p>MCDONALD I: "Distributed, configurable memory management in an operating system supporting quality of service" DISTRIBUTED COMPUTING SYSTEMS, 1999. PROCEEDINGS. 7TH IEEE WORKSHOP ON FUTURE TRENDS OF CAPE TOWN, SOUTH AFRICA 20-22 DEC. 1999, LOS ALAMITOS, CA, USA, IEEE COMPUT. SOC, US, 20 December 1999 (1999-12-20), pages 191-196, XP010367838 ISBN: 0-7695-0468-X abstract page 192, left-hand column, paragraph 2 - page 193, left-hand column, paragraph 2</p>	
A	<p>YANG J ET AL: "FREQUENT VALUE COMPRESSION IN DATA CACHES" PROCEEDINGS OF THE ANNUAL INTERNATIONAL SYMPOSIUM ON MICROARCHITECTURE, 10 December 2000 (2000-12-10), pages 258-265, XP000994541 abstract</p>	
A	<p>LEE J-S ET AL: "Performance analysis of a selectively compressed memory system" MICROPROCESSORS AND MICROSYSTEMS, IPC BUSINESS PRESS LTD. LONDON, GB, vol. 26, no. 2, 17 March 2002 (2002-03-17), pages 63-76, XP004339935 ISSN: 0141-9331 page 67, left-hand column, paragraph 2 - paragraph 3</p>	1,9,14

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

.../US2004/017259

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 5794228	A	11-08-1998	US 5794229 A	11-08-1998
			US 5649181 A	15-07-1997
			US 5918225 A	29-06-1999
			US 5852821 A	22-12-1998
