



(86) Date de dépôt PCT/PCT Filing Date: 2005/07/07  
(87) Date publication PCT/PCT Publication Date: 2006/02/09  
(85) Entrée phase nationale/National Entry: 2006/12/29  
(86) N° demande PCT/PCT Application No.: US 2005/024203  
(87) N° publication PCT/PCT Publication No.: 2006/014573  
(30) Priorité/Priority: 2004/07/07 (US60/586,364)

(51) Cl.Int./Int.Cl. *G06F 15/173* (2006.01)

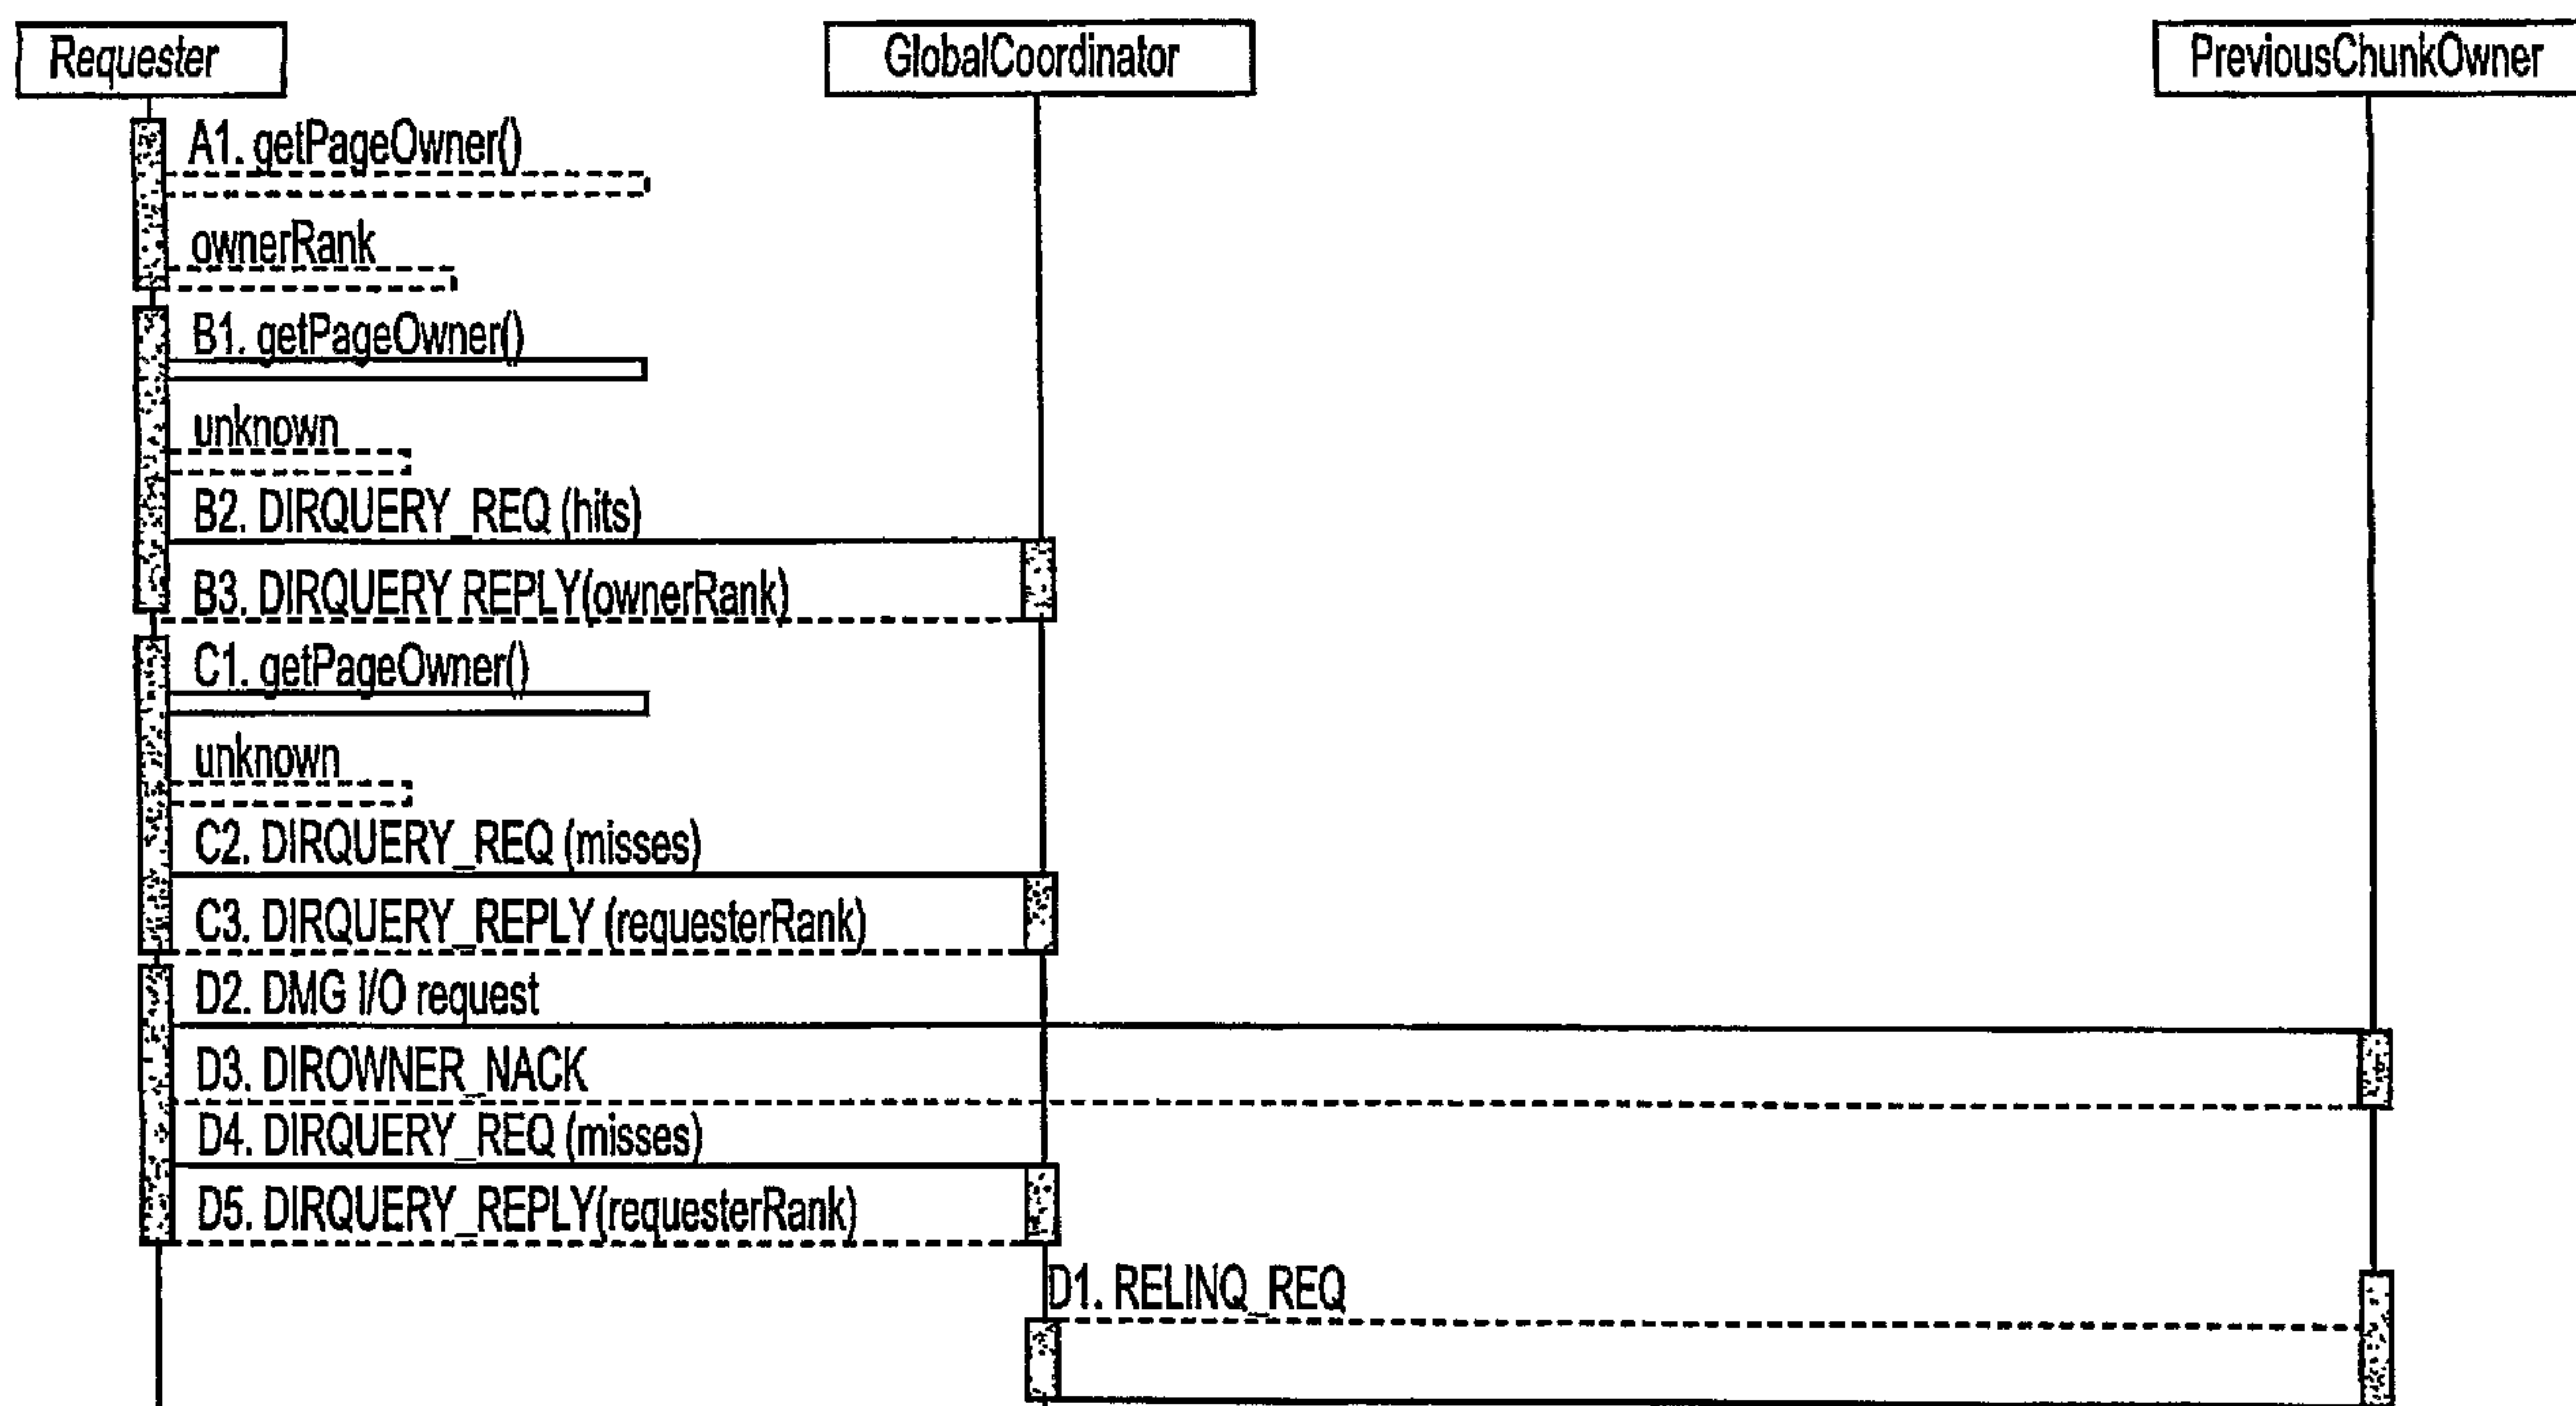
(71) Demandeur/Applicant:  
YOTTAYOTTA, INC., CA

(72) Inventeurs/Inventors:  
BROMLING, STEVEN, CA;  
KARPOFF, WAYNE, CA;  
UNRAU, RON, CA

(74) Agent: SMART & BIGGAR

(54) Titre : SYSTEMES ET PROCEDES POUR METTRE EN OEUVRE UNE COHERENCE DE MEMOIRE CACHE  
REPARTIE

(54) Title: SYSTEMS AND METHODS FOR PROVIDING DISTRIBUTED CACHE COHERENCE



### DMG LCDM Directory Ownership Changes

(57) **Abrégé/Abstract:**

A plurality of access nodes sharing access to data on a storage network implement a directory based cache ownership scheme. One node, designated as a global coordinator, maintains a directory (e.g., table or other data structure) storing information about I/O operations by the access nodes. The other nodes send requests to the global coordinator when an I/O operation is to be performed on identified data. Ownership of that data in the directory is given to the first requesting node. Ownership may transfer to another node if the directory entry is unused or quiescent. The distributed directory-based cache coherency allows for reducing bandwidth requirements between geographically separated access nodes by allowing localized (cached) access to remote data.



## (12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization  
International Bureau



PCT



(43) International Publication Date  
9 February 2006 (09.02.2006)

(10) International Publication Number  
**WO 2006/014573 A2**

(51) International Patent Classification:  
**G06F 15/173** (2006.01)

(21) International Application Number:  
PCT/US2005/024203

(22) International Filing Date: 7 July 2005 (07.07.2005)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
60/586,364 7 July 2004 (07.07.2004) US

(71) Applicant (for all designated States except US): **YOTTA YOTTA, INC.** [CA/CA]; 6020 104th Street, Edmonton, Alberta T6H 5S4 (CA).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **UNRAU, Ron** [CA/CA]; 8934 120 Street, Edmonton, Alberta T6G 1X5 (CA). **BROMLING, Steven** [CA/CA]; 10907 36A Avenue, Edmonton, Alberta T6J 0E3 (CA). **KARPOFF, Wayne** [CA/CA]; 678 Estate Drive, Sherwood Park, Alberta T8B 1M4 (CA).

(74) Agents: **GRAY, Gerald, T.** et al.; Townsend and Townsend and Crew LLP, Two Embarcadero Center, 8th Floor, San Francisco, California 94111 (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, KM, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, 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), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, 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).

**Declaration under Rule 4.17:**

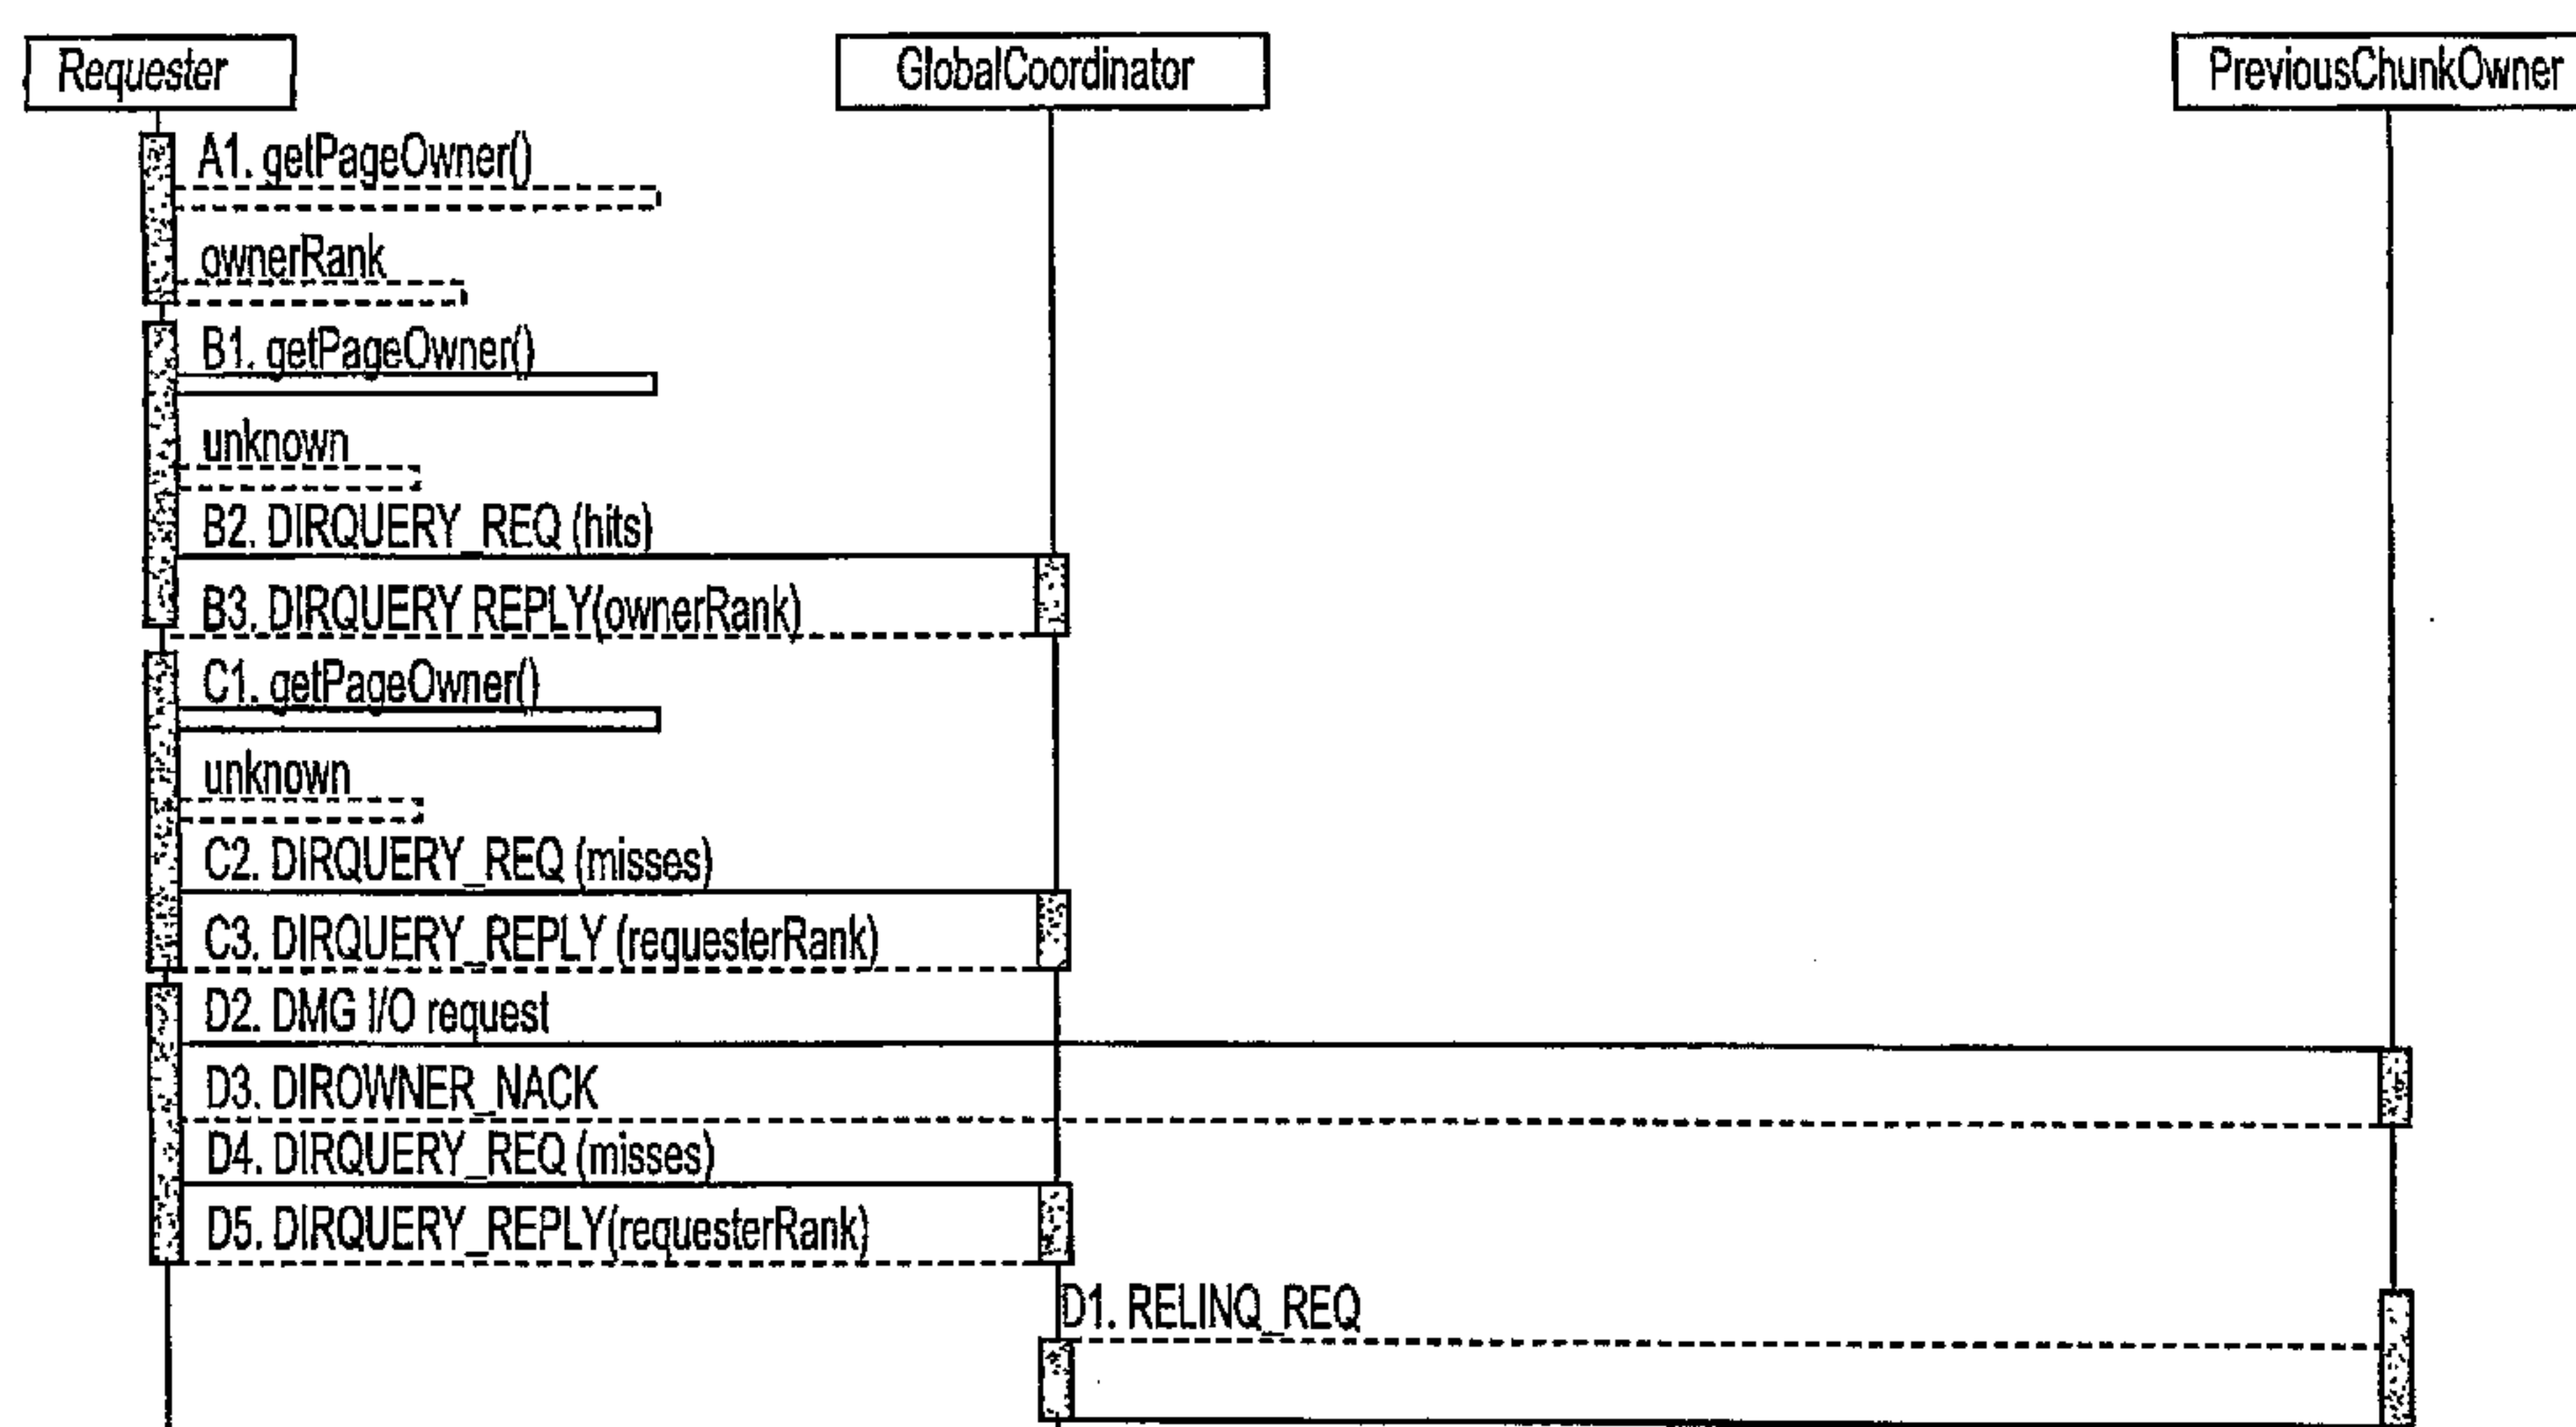
— of inventorship (Rule 4.17(iv)) for US only

**Published:**

— without international search report and to be republished upon receipt of that report

[Continued on next page]

(54) Title: SYSTEMS AND METHODS FOR PROVIDING DISTRIBUTED CACHE COHERENCE



DMG LCDM Directory Ownership Changes

(57) Abstract: A plurality of access nodes sharing access to data on a storage network implement a directory based cache ownership scheme. One node, designated as a global coordinator, maintains a directory (e.g., table or other data structure) storing information about I/O operations by the access nodes. The other nodes send requests to the global coordinator when an I/O operation is to be performed on identified data. Ownership of that data in the directory is given to the first requesting node. Ownership may transfer to another node if the directory entry is unused or quiescent. The distributed directory-based cache coherency allows for reducing bandwidth requirements between geographically separated access nodes by allowing localized (cached) access to remote data.

WO 2006/014573 A2

**WO 2006/014573 A2**



---

*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.*

## SYSTEMS AND METHODS FOR PROVIDING DISTRIBUTED CACHE COHERENCE

### CROSS-REFERENCES TO RELATED APPLICATIONS

- 5 [0001] This application claims the benefit of US provisional application No. 60/586,364, filed July 7, 2004, which is hereby incorporated by reference in its entirety.

### BACKGROUND OF THE INVENTION

10 [0002] The present invention relates generally to sharing data across one or more data storage access nodes on a data storage network, and more particularly to systems and method for providing directory-based cache coherency across a distributed network of data storage access nodes.

15 [0003] In current storage networks, and in particular storage networks including geographically remote access nodes and storage resources, preserving or reducing bandwidth between resources and access nodes is highly desirable. It is therefore also desirable that data access be localized, in part to improve access speed to pages requested by host devices. Caching pages at access nodes provides localization, however, the cached data must be kept coherent with respect to modifications at other access nodes that may be caching the same data. Current storage network access solutions, however, do not provide viable coherency mechanisms for caching pages locally at storage network access nodes.

20 [0004] Accordingly, it is desirable to provide efficient data localization and cache coherency systems and methods that overcome the above and other problems. Such systems and methods should also provide reduced bandwidth usage, or messaging requirements, between storage network access nodes.

25

### BRIEF SUMMARY OF THE INVENTION

[0005] The present invention provides systems and methods for implementing directory-based cache coherency across a distributed network of data storage access nodes.

30 [0006] According to the present invention, a plurality of access nodes sharing access to data on a storage network implement a directory based cache ownership scheme. One node,

designated as a global coordinator, maintains a directory (e.g., table or other data structure) storing information about I/O operations by the access nodes. The other nodes send requests to the global coordinator when an I/O operation is to be performed on identified data.

Ownership of that data in the directory is given to the first requesting node. Ownership may transfer to another node if the directory entry is unused or quiescent. According to the present invention, the distributed directory-based cache coherency allows for reducing bandwidth requirements between geographically separated access nodes by allowing localized (cached) access to remote data.

[0007] According to one aspect of the present invention, a method is provided for reducing the number messages sent between data access nodes sharing access to a data storage network so as to maintain traffic scalability. The method typically includes maintaining a directory of page ownership entries, wherein ownership of an entry is initially granted to the first access node requesting access to a page in the entry, and wherein ownership of an entry automatically transfers to the node that is accessing pages in the entry more often so as to reduce the number of synchronization messages sent between nodes.

[0008] According to another aspect of the present invention, a method is provided for reducing bandwidth between geographically separated access nodes sharing access to data in a data storage network. The method typically includes caching data locally to an access node to provide localized cache access to that data for that node, and maintaining data coherency for cached data between the access nodes using a directory based ownership scheme.

[0009] According to yet another aspect of the present invention, a method of providing cache coherence between caches in a distributed set of data access nodes in a data storage network typically includes maintaining a directory in at least one of a plurality of access nodes sharing access to the data storage network, the directory storing information about data accessed by the plurality of access nodes, and receiving, at a first data access node, a data access request from a host system, the data access request identifying data to be processed. The method also typically includes determining whether the first access node has the identified data stored in cache, and if not, determining, using the directory, whether another node in the plurality of access nodes has a copy of the identified data stored in cache, and if a node has a copy of the identified data stored in cache, sending one of a share request to that node to share the identified data so that the requesting node does not have to access the

identified data from storage or an invalidate request to invalidate the copy of the data stored in that node's cache.

[0010] According to a further aspect of the present invention, a method is provided for reducing the a number messages sent between data access nodes sharing access to a data storage network so as to maintain traffic scalability. The method typically includes maintaining a directory for storing information about data accessed by a plurality of data access nodes, where the directory including entries representing one or more pages of data in the data storage network, and receiving, at a first data access node, a data access request from a host system, where the data access request identifying data to be processed. The method also typically includes determining, using a global directory coordinator, whether a node has ownership of the directory entry for the identified data, and if no node has ownership of the directory entry, granting to the first access node ownership of the directory entry for the identified data, and if a node has ownership of the entry, identifying that node to the first node. The first node can then communicate with the identified node to process an I/O request.

[0011] According to yet a further aspect, a system is provided for maintaining cache coherency between a plurality of data access nodes sharing access to a data storage network. The system typically includes a storage system for storing data, and a plurality of access nodes configured to access data in the storage system in response to host requests. One of the nodes is typically configured to maintain a directory (directory node) for storing information about data accessed by the plurality of data access nodes, where the directory includes entries representing one or more pages of data in the data storage network. In operation, upon receiving a data access request identifying data to be processed from a host, a first access node sends a request to the directory node, wherein the directory node determines whether a node has ownership of the directory entry for the identified data; and if no node has ownership of the directory entry, the directory node grants to the first access node ownership of the directory entry for the identified data, and if a node has ownership of the entry, the directory node identifies that node to the first access node.

[0012] Reference to the remaining portions of the specification, including the drawings and claims, will realize other features and advantages of the present invention. Further features and advantages of the present invention, as well as the structure and operation of various embodiments of the present invention, are described in detail below with respect to the

accompanying drawings. In the drawings, like reference numbers indicate identical or functionally similar elements.

### BRIEF DESCRIPTION OF THE DRAWINGS

5 [0013] FIG. 1 shows an example of a basic network configuration according to the present invention.

[0014] FIG. 2 shows examples of the scenarios relevant to directory ownership changes.

[0015] FIG.3 shows an example of a variety of message sequences that are observed when a read misses in the cache and comes down to the DMG in an attempt to share a page from  
10 another node.

[0016] FIG.4 shows an example of message sequences when a write request comes down to the DMG from a requesting node.

[0017] FIG.5 shows examples of the message sequences for update requests.

### 15 DETAILED DESCRIPTION OF THE INVENTION

[0018] According to one embodiment, a Directory Manager module, or DMG, is provided. The DMG is responsible for providing cache coherence mechanisms for shared data across a distributed set of data access nodes. The set of nodes that are caching data from a shared data volume are called a share group. In general, a DMG module includes software executing on  
20 a processor or other intelligence module (e.g., ASIC) in a node. A DMG module can be implemented in a single node or distributed across multiple intercommunicating nodes. In certain aspects, an access node is embodied as a controller device, or blade, communicably coupled to a storage network, such as a storage area network (SAN), that allows access to data stored on the storage network. However, it will be appreciated that an access node can  
25 also be embodied as an intelligent fabric switch or other network device such as a hub adapter. Because Locality Conscious Directory Migration (LCDM) is applicable to databases, any networked compute node can be configured to operate as an access node with DMG functionality (e.g., a DMG can be run on a desktop computer with a network connection). US Patent No. 6,148,414, which is incorporated by reference in its entirety, discloses  
30 controller devices and nodes for which implementation of aspects of the present invention are particularly useful.

[0019] FIG. 1 shows a basic network configuration according to the present invention. As shown, a plurality of network clients  $10_1$  to  $10_N$  are communicably coupled with a plurality of access node devices  $20_1$  to  $20_N$  (access nodes), each of which includes a processor component 22 such as a microprocessor or other intelligence module, a cache 24 (e.g., RAM cache) and/or other local storage, communication ports (not shown) and an instance of a DMG module 26. (In general, "N" is used herein to indicate an indefinite plurality, so that the number "N" when referred to one component does not necessarily equal the number "N" of a different component. For example, the number of clients 10 does not, but may, equal the number of access nodes 20 in FIG. 1.) Each client 10 is preferably communicably coupled to one or more of nodes 20 over a local network connection 15, however, clients may be communicably coupled with nodes 20 over any of a number of connection schemes as required for the specific application and geographical location relative to a node 20, including, for example, a direct wired or wireless connection, an Internet connection, any local area network (LAN) type connection, any metropolitan area network (MAN) connection, any wide area network (WAN) type connection, a VLAN, any proprietary network connection, etc. Each node 20 also typically includes, or is communicably coupled with, one or more AMFs, and is communicably coupled with one or multiple storage resources 40, each including one or more disk drives, over an network 50, such as a storage area network (SAN), LAN, WAN, MAN, etc. It is preferred that a node 20 be coupled to one or more storage resources over a local network connection. Nodes 20 may be located in close physical proximity to each other or one or more may be remotely located, e.g., geographically remote, from other nodes. Access nodes 20 are also able to intercommunicate with other nodes over network 50 and/or over other communication networks or mediums such as over a PCI bus or backbone or a Fibre channel network, or over network 15.

[0013] According to the present invention, distributed cache coherence is important for reducing bandwidth requirements between geographically separated access nodes by allowing localized (cached) access to remote data. According to one aspect, data access cannot be localized unless the data can be cached, yet it is unsafe to cache the data unless it can be kept coherent with respect to modifications at remote access nodes. Although any embodiment of the DMG can satisfy the correctness requirements of cache coherence, the high overhead of many implementations can outweigh the benefits of localized cache access. The LCDM embodiment of the present invention discussed below has demonstrated low enough overhead to make localized cache access practical and beneficial.

[0020] The base coherence unit in the DMG is a page (a logical block of storage), but the DMG allows for operations at both the sub-page and the multi-page levels. The directory is a collection of directory entries, each encoding distributed sharing knowledge for a specific page. When concurrent cache operations are active on a page, the directory entry locks and synchronizes access to the distributed resource. Directory information is kept current through point-to-point messages sent between the affected nodes. The DMG cache coherence messaging dialog allows it to share pages from remote caches (e.g., when read requests miss in the local cache) and invalidate remote cached copies (e.g., when write requests supercede previous copies).

[0021] Embodiments of a directory placement scheme and ways to take advantage of data access locality are described below in section 1.1. Section 1.2 introduces a messaging dialog and goes on, in Section 1.3, to demonstrate how locking is utilized at the directory to manage concurrent page accesses according to one embodiment. Lastly, Section 1.4 describes how location-aware page sharing is used to improve performance for geographically distributed cache coherence according to one embodiment.

## 1.1 Directory Placement

[0022] According to one aspect, in order to coordinate cache coherence, the DMG maintains a directory with entries tracking information about every active page. Active pages are those that can be found in at least one cache in the associated share group. The directory entry tracks which nodes have copies of the associated page, maintains a distributed lock to protect concurrent accesses to the page, and maintains a queue to serialize operations while the nodes wait for the lock.

### 1.1.1 Dedicated Directory Server

[0023] According to one embodiment, a very simple directory placement scheme is used in which the entire directory is situated on the first node that joins a share group. Such a directory scheme, however, may not scale well with the share group size. In general, there are two competing interests involved in determining the performance of a directory placement scheme. DMG operations on a node containing the directory entry for a given page can be optimized to avoid unnecessary message sends, thereby preserving inter-node bandwidth and improving operation latency. However, the directory node for a page has to process all the relevant message traffic from other nodes in the share group.

### 1.1.2 Striped Directory Placement

[0024] According to another embodiment, directory entries are striped in a round-robin fashion across the nodes in each share group. Because the DMG often has to deal with multi-page operations, the stripe size should be made sufficiently large to avoid frequent splitting of operations. Such striping is easy to implement and has minimal performance overhead, but two problems may make it desirable to seek a better solution. Those two problems are as follows:

1. Striping only performs well when all of the nodes in a share group are actively sharing pages distributed throughout the data volume. If any nodes are inactive, or if nodes are only sharing pages whose directory entries are on remote nodes, the performance may suffer. There is no way for a striped directory placement to adapt to a given workload.

2. Because striping is implemented, in one aspect, using a modulus function across nodes, if a node enters or leaves the share group, there is no quick and easy way to redistribute directory entries. Invalidating all caches and starting the directories from scratch is one option.

### 1.1.3 Locality Conscious Directory Migration

[0025] According to one embodiment of the present invention, the DMG implements a Locality Conscious Directory Migration, or LCDM, adaptive directory placement scheme.

As in the striped directory placement embodiment, in the LCDM embodiment, directory entries are split up into chunks of multiple pages. However, the placement, or ownership, of those chunks is not preset or permanent. Instead, ownership of the directory chunks migrate to the node that is using those pages the most frequently. In addition to taking advantage of data access locality for improved I/O performance, LCDM advantageously helps in quickly recovering from node failure (or nodes being added to the share group), since there is no need to redistribute directory chunks. In one embodiment, directory chunk ownership is granted to the first node that accesses a page in the chunk.

[0026] In one embodiment, ownership changes hands only when a directory chunk is empty, after all of the nodes have evicted every page in the chunk from their caches. One node in the share group is designated as the global directory chunk coordinator ("global coordinator"). This node grants ownership of a directory entry, and stores a look-up table, or

other similar data structure, identifying ownership of each active page in the share group. The other nodes in the share group maintain local partial mirrors of the look-up table, but occasionally need to defer to the global coordinator. A global coordinator may itself also be or become an owner of a directory entry. FIG. 2 summarizes the scenarios relevant to

5 directory ownership changes, which are explored in more detail below.

A. When the local node already knows who the owner of a directory chunk is, a simple table lookup to identify or confirm the owner before an I/O operation is performed.

10 B. If the local node is accessing a directory chunk for the first time, it sends a DIRQUERY\_REQ message to the global coordinator node. If another node has already claimed ownership of the chunk, the coordinator responds to the requester with a DIRQUERY\_REPLY message containing the rank of the chunk's owner. The requester then communicates with the access node owning the chunk to complete the I/O operation(s).

15 C. If, in the previous scenario, the global coordinator's look-up indicates that no node owns the directory chunk in question, it responds to the requester with a DIRQUERY\_REPLY message containing the requester's rank and granting ownership. The rank is an index identifying the order in which the node joined the share group. An entry in the global coordinator's look-up table is made, and

20 synchronized with other nodes.

D. When a chunk owner no longer has any directory entries, it sends a relinquish ownership, RELINQ\_REQ, message to the global coordinator. If new requests arrive at the original chunk owner after a DMG node has relinquished chunk ownership, the original chunk owner responds with a DIOWNER\_NACK message to the requester.

25 This will cause the requester to re-query the global coordinator. The first node to query the global coordinator after the relinquishment will be granted ownership of the chunk.

[0027] In another embodiment, LCDM is enhanced by allowing ownership of directory chunks to migrate between nodes even when they are not empty. This allows for more

30 closely following I/O access patterns and improves the performance of the DMG by minimizing messaging requirements. However, this optimization generally has certain requirements as follows:

1. Directory chunk ownership can only be moved safely when the entire chunk is quiescent. If there is any active coherence traffic or locked pages, the move is not feasible.

2. Moving non-empty chunks involves shipping the sharing information to the new chunk owner. This must be done semi-atomically, so that new coherence requests are not started in the middle of the move.

3. Because the cost of transferring chunk ownership is not nil, the heuristic used to determine when to effect a transfer is very important, especially where there may be limited prior knowledge of cache workload.

4. Coordinating a directory chunk transfer can involve up to three nodes: the current chunk owner, the global coordinator and the new chunk owner. This may make the task of semi-atomically transferring more difficult.

## 1.2 Standard I/O Paths in the DMG

[0028] According to one embodiment, a message dialog is provided in the DMG for data access requests, e.g., read, write and update requests. Sequence dialogs are used (see FIGS. 3-5) to detail the message ordering. In each Figure, a variety of scenarios are explored to explain the range of conditions that must be dealt with at the DMG. These Figures differentiate between the node containing the client cache, the node owning the appropriate directory entry (directory node), and all other nodes that have a copy of the page or pages in question. It is important to note that in many cases, these entities coexist on the same node. Each role-sharing combination introduces new optimizations, and is handled as follows:

**Single Member Share Group:** In a single member share group, the directory exists only implicitly, as there is no ownership information for it to track. Requests from the client cache are trivial to handle, and can be answered immediately.

**Client Cache and Directory Coexist:** When the directory is local to the node making a DMG request, it removes the need for two message sends: the request at the start and the reply at the end.

**Directory and Page Owner Coexist:** When the directory node is also a page owner, it can handle share or invalidate requests locally, avoiding the need for extra messages.

### 1.2.1 DMG Reads

[0029] FIG. 3 details an example of the variety of message sequences that are observed when a read misses in the cache and comes down to the DMG in an attempt to share a page from another node. The following points provide a detailed description of each of the scenarios in Figure 2.

A. When the directory does not contain an entry for the requested page, the directory node creates a new entry, marks the requester as the only page owner, and replies with a SHARE\_NACK message. This causes the requester to go to an AMF ("Array Management Function") to get a copy of the page data from disk. In general, and as described in more detail in US Patent No 6,148,414, an AMF provides access to storage resources, e.g., disks, for associated host systems. For example, when a host requests an AMF to perform an operation on a resource, the AMF synchronizes with any other AMFs sharing control of the redundancy group that includes the resource to be operated on, so as to obtain a lock on the resource. The AMF is typically resident at the same node as the cache for which a miss occurred, however, a remote AMF may be used.

B. If the directory look-up finds a remote copy of the page, the DMG sends a SHARE\_REQ message to that node storing a cached copy ("sharer"). The sharer sends the requested page in a SHARE\_ACK message directly to the client node, bypassing the directory node. When the sharer receives confirmation from the communication layer that the data has been sent, it sends a READ\_DONE message to the directory node, allowing any waiting operations to commence.

C. If, in the previous scenario, the SHARE\_REQ message arrives at the sharer after it has asynchronously evicted the page (see Section 1.2.4 below), the sharer replies to the directory node with a SHARE\_NACK message. If there is another potential sharer, the directory node sends out a subsequent SHARE\_REQ message and continues normally. Otherwise, if no other node has the page, the directory node replies to the requester with a SHARE\_NACK message.

D. If the sharer dies after sending its SHARE\_ACK message to the requester but before being able to send the READ\_DONE message to the directory node, the directory node makes the assumption that the requester did not get the data. The directory node then sends a READ\_FAILED message to the requester, indicating that

the read may have failed. If the read was successful, the requester responds to the directory node with a READ\_FAILED\_NACK message. In the case of failure, the directory node needs to know whether the requester received the data or not, so that it can correctly handle further read requests as it finishes bringing the system to a quiescent state before commencing recovery.

E. In the previous example, if the sharer dies before it is able to send the SHARE\_ACK message to the requester, the requester responds to the DMG's eventual READ\_FAILED message with a READ\_FAILED\_ACK message.

### 1.2.2 DMG Writes

[0030] The message sequences when a write request comes down to the DMG from a requesting node are shown in FIG. 4 and described below.

A. If the directory does not contain an entry for the requested page, or finds an entry that indicates that the requester is the sole owner of the page, the directory node responds to the requester immediately with a WRITE\_REPLY message. Upon receipt of the reply, the requester's cache is allowed to commit the newly written data to cache as soon as it is ready.

B. If the directory entry for the page in question shows that other nodes have copies of the page in their caches, the directory node sends out INVALID\_REQ messages to each of them. Once all of the INVALID\_REPLY messages have come back to the directory node, it replies as above to the requester.

### 1.2.3 DMG Updates

[0031] As with reads, a number of messaging scenarios are possible in the DMG in response to an update (*i.e.*, a sub-page write) request. FIG. 5 shows examples of the message sequences for update requests, which are also described in more detail below.

A. When an UPDATE\_REQ message arrives at the directory node, an UPDATE\_NACK message is returned immediately if either there is no directory entry for the page, or the requester is the sole owner of the page. The UPDATE\_NACK message is always the directory's response when the requester already has a copy of the page or there is no copy of the page on any node. In the case where no nodes have a cached copy of the page, the requester reads the page data

from an AMF. Once the requester's cache has completed its sub-page write, the requester responds to the DMG with an UPDATE\_DONE message.

5 B. If the directory entry indicates that in addition to the requester owning a copy of the page, there is at least one other node with the page in cache, the directory node sends out UPINV\_REQ messages to each of the other page owners. Once each of these nodes has invalidated their copy of the page and responded to the directory node with a UPINV\_REPLY message, the directory node will continue by sending a UPDATE\_NACK message to the requester as above.

10 C. If the requester does not have a copy of the page in its cache and the directory entry indicates that at least one other node has a copy of the page, the directory node sends a SHINV\_REQ message to the node that last wrote the page and an UPINV\_REQ message to the other nodes that hold a copy. The node to which the SHINV\_REQ message is sent shares its copy of the page by forwarding it to the directory node in a SHINV\_REPLY message, and then invalidates the local copy  
15 once the send is successfully complete. Once all of the replies have been received by the directory, the directory node sends an UPDATE\_ACK message to the requester with the data from the share-invalidate target node.

20 D. If, in the previous scenario, the share-invalidate target node has synchronously evicted its copy of the page before receiving the SHINV\_REQ message from the directory node, the target node responds with a SHINV\_NACK message. Because the directory node will have already sent out invalidation requests to all other nodes with a copy of the page, at this point there will be no cached copies of the page. The directory node has no option but to respond to the requester with an UPDATE\_NACK message (after waiting for all of the invalidation replies to come in), causing the  
25 requester to read the page from an AMF. This is fine because the node that evicted its copy of the page would have first needed to flush its copy out to an AMF - the directory always keeps track of the last writer to ensure data integrity.

30 E. If, after sending a SHINV\_REQ message, the directory node notices that the share-invalidate target of that message has died, the directory node waits for the invalidation replies from other nodes and then fails the request by sending an UPDATE\_FAILED message to the requester.

#### 1.2.4 DMG Evictions

[0032] According to one aspect, to avoid unnecessary message latency, when a node cache evicts a page it simply sends out an asynchronous EVICT\_REQ message to the DMG to notify it of the eviction. Because the notification is asynchronous, the cache may receive requests for the page after evicting it, but they are handled by the I/O message scenarios detailed above.

### 1.3 DMG Locking

[0033] One of the DMG's responsibilities is to prevent multiple concurrent changes to page data. The DMG allows multiple readers to access a page at the same time. Multiple writers, however, are serialized in the order in which the write requests arrived at the directory node, e.g., placed in a FIFO queue or buffer. According to one aspect, the caches on the nodes of the system may also require some serialization or locking to reserve and manipulate page frames. Because the nodes may need to hold both cache and DMG locks for the same pages at the same time, the system preferably includes a deadlock handling mechanism. According to one embodiment, the cache locks are made subordinate to the DMG locks, which means a locally held cache lock for a given page must be released if a DMG operation on that page is received.

[0034] When a read request arrives at the directory node, if the directory entry indicates that there are no operations waiting for the page and zero or more readers currently active, the read is allowed to proceed immediately. In contrast, when a write or update request arrives at the directory node, it can only proceed if there is no activity whatsoever on the page. In any other situation, new requests to the directory are queued up in the directory entry until they are activated by the completion of a preceding operation. For multi-page operations, the DMG gathers all of the locks for an operation in increasing order before replying to the client node and allowing the associated operation to proceed. If the DMG were to ignore these constraints and respond out-of-order to the client node for segments of a multi-page operation, a deadlock could occur.

[0035] According to one aspect, one optimization in the DMG to avoid wasted effort in the face of multiple concurrent writes to the same page involves checking the lock wait queue before allowing a write to proceed. If the write at the head of the queue is immediately followed by one or more writes from any node, the DMG can preemptively invalidate the earlier writes and allow the last write to proceed immediately. Since concurrent writes are not expected to be a common case, the implementation of this optimization is not crucial.

#### 1.4 Location Awareness

[0036] According to one embodiment, the DMG optimizes its page sharing strategy based on the physical or logical proximity of nodes within a share group. In one aspect, the site id of each node is recorded by all of the nodes in each share group. When the directory node  
5 receives a read request, it looks first to see if any node within the same site (or close physical proximity) as the reader has the page. If so, it sends a share request to the nearest node. In another aspect, as a second option, sharing is done by the directory node itself (assuming it has the page in cache), to save the share request message send. If the directory node doesn't have the requested page cached, then any other node can be selected. In this case, no  
10 consideration is given to the relative distance between sites (*i.e.*, nodes are either near or far in relation to each other). This same algorithm also applies to page sharing in an update situation.

[0037] It should be appreciated that code including instructions for implementing aspects of the DMG, including LCDM, can be stored on a computer readable medium such as a CD,  
15 DVD, ROM, RAM or the like, or can be transmitted over a network connection to and from data access node devices.

[0038] While the invention has been described by way of example and in terms of the specific embodiments, it is to be understood that the invention is not limited to the disclosed embodiments. To the contrary, it is intended to cover various modifications and similar  
20 arrangements as would be apparent to those skilled in the art. Therefore, the scope of the appended claims should be accorded the broadest interpretation so as to encompass all such modifications and similar arrangements.

WHAT IS CLAIMED IS:

1           1.       A method of reducing a number messages sent between data access  
2 nodes sharing access to a data storage network so as to maintain traffic scalability, by  
3 maintaining a directory of page ownership entries, wherein ownership of an entry is initially  
4 granted to the first access node requesting access to a page in the entry, and wherein  
5 ownership of an entry automatically transfers to the node that is accessing pages in the entry  
6 more often so as to reduce the number of synchronization messages sent between nodes.

1           2.       A method of reducing bandwidth between geographically separated  
2 access nodes sharing access to data in a data storage network, comprising caching data  
3 locally to an access node to provide localized cache access to that data for that node, and  
4 maintaining data coherency for cached data between the access nodes using a directory based  
5 ownership scheme.

1           3.       A method of providing cache coherence between caches in a  
2 distributed set of data access nodes in a data storage network, the method comprising:  
3               maintaining a directory in at least one of a plurality of access nodes sharing  
4 access to the data storage network, said directory storing information about data accessed by  
5 the plurality of access nodes;

6               receiving, at a first data access node, a data access request from a host system,  
7 said data access request identifying data to be processed,

8               determining whether the first access node has the identified data stored in  
9 cache, and if not:

10              determining, using the directory, whether another node in the plurality of  
11 access nodes has a copy of the identified data stored in cache; and

12              if a node has a copy of the identified data stored in cache, sending one of a  
13 share request to that node to share the identified data so that the requesting node does not  
14 have to access the identified data from storage or an invalidate request to invalidate the copy  
15 of the data stored in that node's cache.

1           4.       The method of claim 3, wherein the data access request is a data read  
2 request, and wherein if no node has a copy of the identified data stored in cache, the method  
3 further comprises:

4               retrieving the identified data from storage; and

5 creating a new entry in the directory identifying the requesting node as an  
6 owner of the identified data.

1 5. The method of claim 3, wherein the data access request is a write  
2 request, and wherein if no node has a copy of the identified data stored in cache, the method  
3 further comprises:

4 committing written data to the requesting node's cache.

1 6. The method of claim 3, wherein the data access request is a write  
2 request, and wherein if one or more other nodes has a copy of the identified data stored in  
3 cache, the method further comprises:

4 sending an invalidate request to the other nodes that have a copy of the  
5 identified data stored in cache.

1 7. The method of claim 3, wherein the data access request is an update  
2 request, and if the requesting node has a copy of the identified data stored in cache, and  
3 wherein if one or more other nodes has a copy of the identified data stored in cache, the  
4 method further comprises:

5 sending an invalidate request to the other nodes that have a copy of the  
6 identified data stored in cache.

1 8. The method of claim 3, wherein the data access request is an update  
2 request, and if the requesting node has no copy of the identified data stored in cache, and  
3 wherein if one or more other nodes has a copy of the identified data stored in cache, the  
4 method further comprises:

5 sending a share-invalidate request to the node that last wrote the data; and  
6 sending an invalidate request to the other nodes that have a copy of the  
7 identified data stored in cache.

1 9. The method of claim 3, further comprising:  
2 determining whether another node has a directory lock on the data identified  
3 by the data access request;

4 if no node has a directory lock on the data, granting a directory lock on the  
5 identified data to the first access node; and

6 if a node has a directory lock on the data, waiting for the lock to be released  
7 before granting it to the first access node.

1                   10.     The method of claim 3, wherein an unowned directory entry is placed  
2 local to the first access node to access the corresponding data page, and wherein the first  
3 access node is given ownership of the directory entry.

1                   11.     The method of claim 3, wherein the data access request identifies one  
2 of a page of data, a sub-page of data, or multiple pages of data.

1                   12.     The method of claim 3, wherein the data access request is one of a data  
2 read request, a data write request and a data update request.

1                   13.     A method of reducing a number messages sent between data access  
2 nodes sharing access to a data storage network so as to maintain traffic scalability, the  
3 method comprising:

4                   maintaining a directory for storing information about data accessed by a  
5 plurality of data access nodes, said directory including entries representing one or more pages  
6 of data in the data storage network;

7                   receiving, at a first data access node, a data access request from a host system,  
8 said data access request identifying data to be processed;

9                   determining, using a global directory coordinator, whether a node has  
10 ownership of the directory entry for the identified data; and

11                   if no node has ownership of the directory entry, granting to the first access  
12 node ownership of the directory entry for the identified data; and

13                   if a node has ownership of the entry, identifying that node to the first node.

1                   14.     The method of claim 13, wherein one of said access nodes is  
2 designated as the global directory coordinator, and wherein the global coordinator grants  
3 ownership of directory entries.

1                   15.     The method of claim 13, wherein the data access request is one of a  
2 data read request, a data write request and a data update request, and wherein the data access  
3 request identifies one of a page of data, a sub-page of data, or multiple pages of data. .

1                   16.     The method of claim 13, wherein the first access node maintains the  
2 directory as a table in memory, and wherein determining includes performing a lookup in the  
3 directory table..

1           17.     The method of claim 13, further including transferring ownership of  
2 the directory entry from the current owner to the first access node if the entry is currently  
3 unused.

1           18.     A system for maintaining cache coherency between a plurality of data  
2 access nodes sharing access to a data storage network, the system comprising:  
3           a storage system for storing data; and  
4           a plurality of access nodes configured to access data in the storage system,  
5 wherein one of said nodes is configured to maintain a directory (directory node) for storing  
6 information about data accessed by the plurality of data access nodes, said directory including  
7 entries representing one or more pages of data in the data storage network;  
8           wherein upon receiving a data access request identifying data to be processed  
9 from a host, a first access node sends a request to the directory node, wherein the directory  
10 node determines whether a node has ownership of the directory entry for the identified data;  
11 and if no node has ownership of the directory entry, the directory node grants to the first  
12 access node ownership of the directory entry for the identified data, and if a node has  
13 ownership of the entry, the directory node identifies that node to the first access node.

1           19.     The method of claim 18, wherein ownership of the directory entry for  
2 the identified data changes to a second access node if the second access node requests  
3 ownership and the directory entry is not being used.

1           20.     The method of claim 18, wherein one or more of the access nodes are  
2 geographically remote from the other nodes.

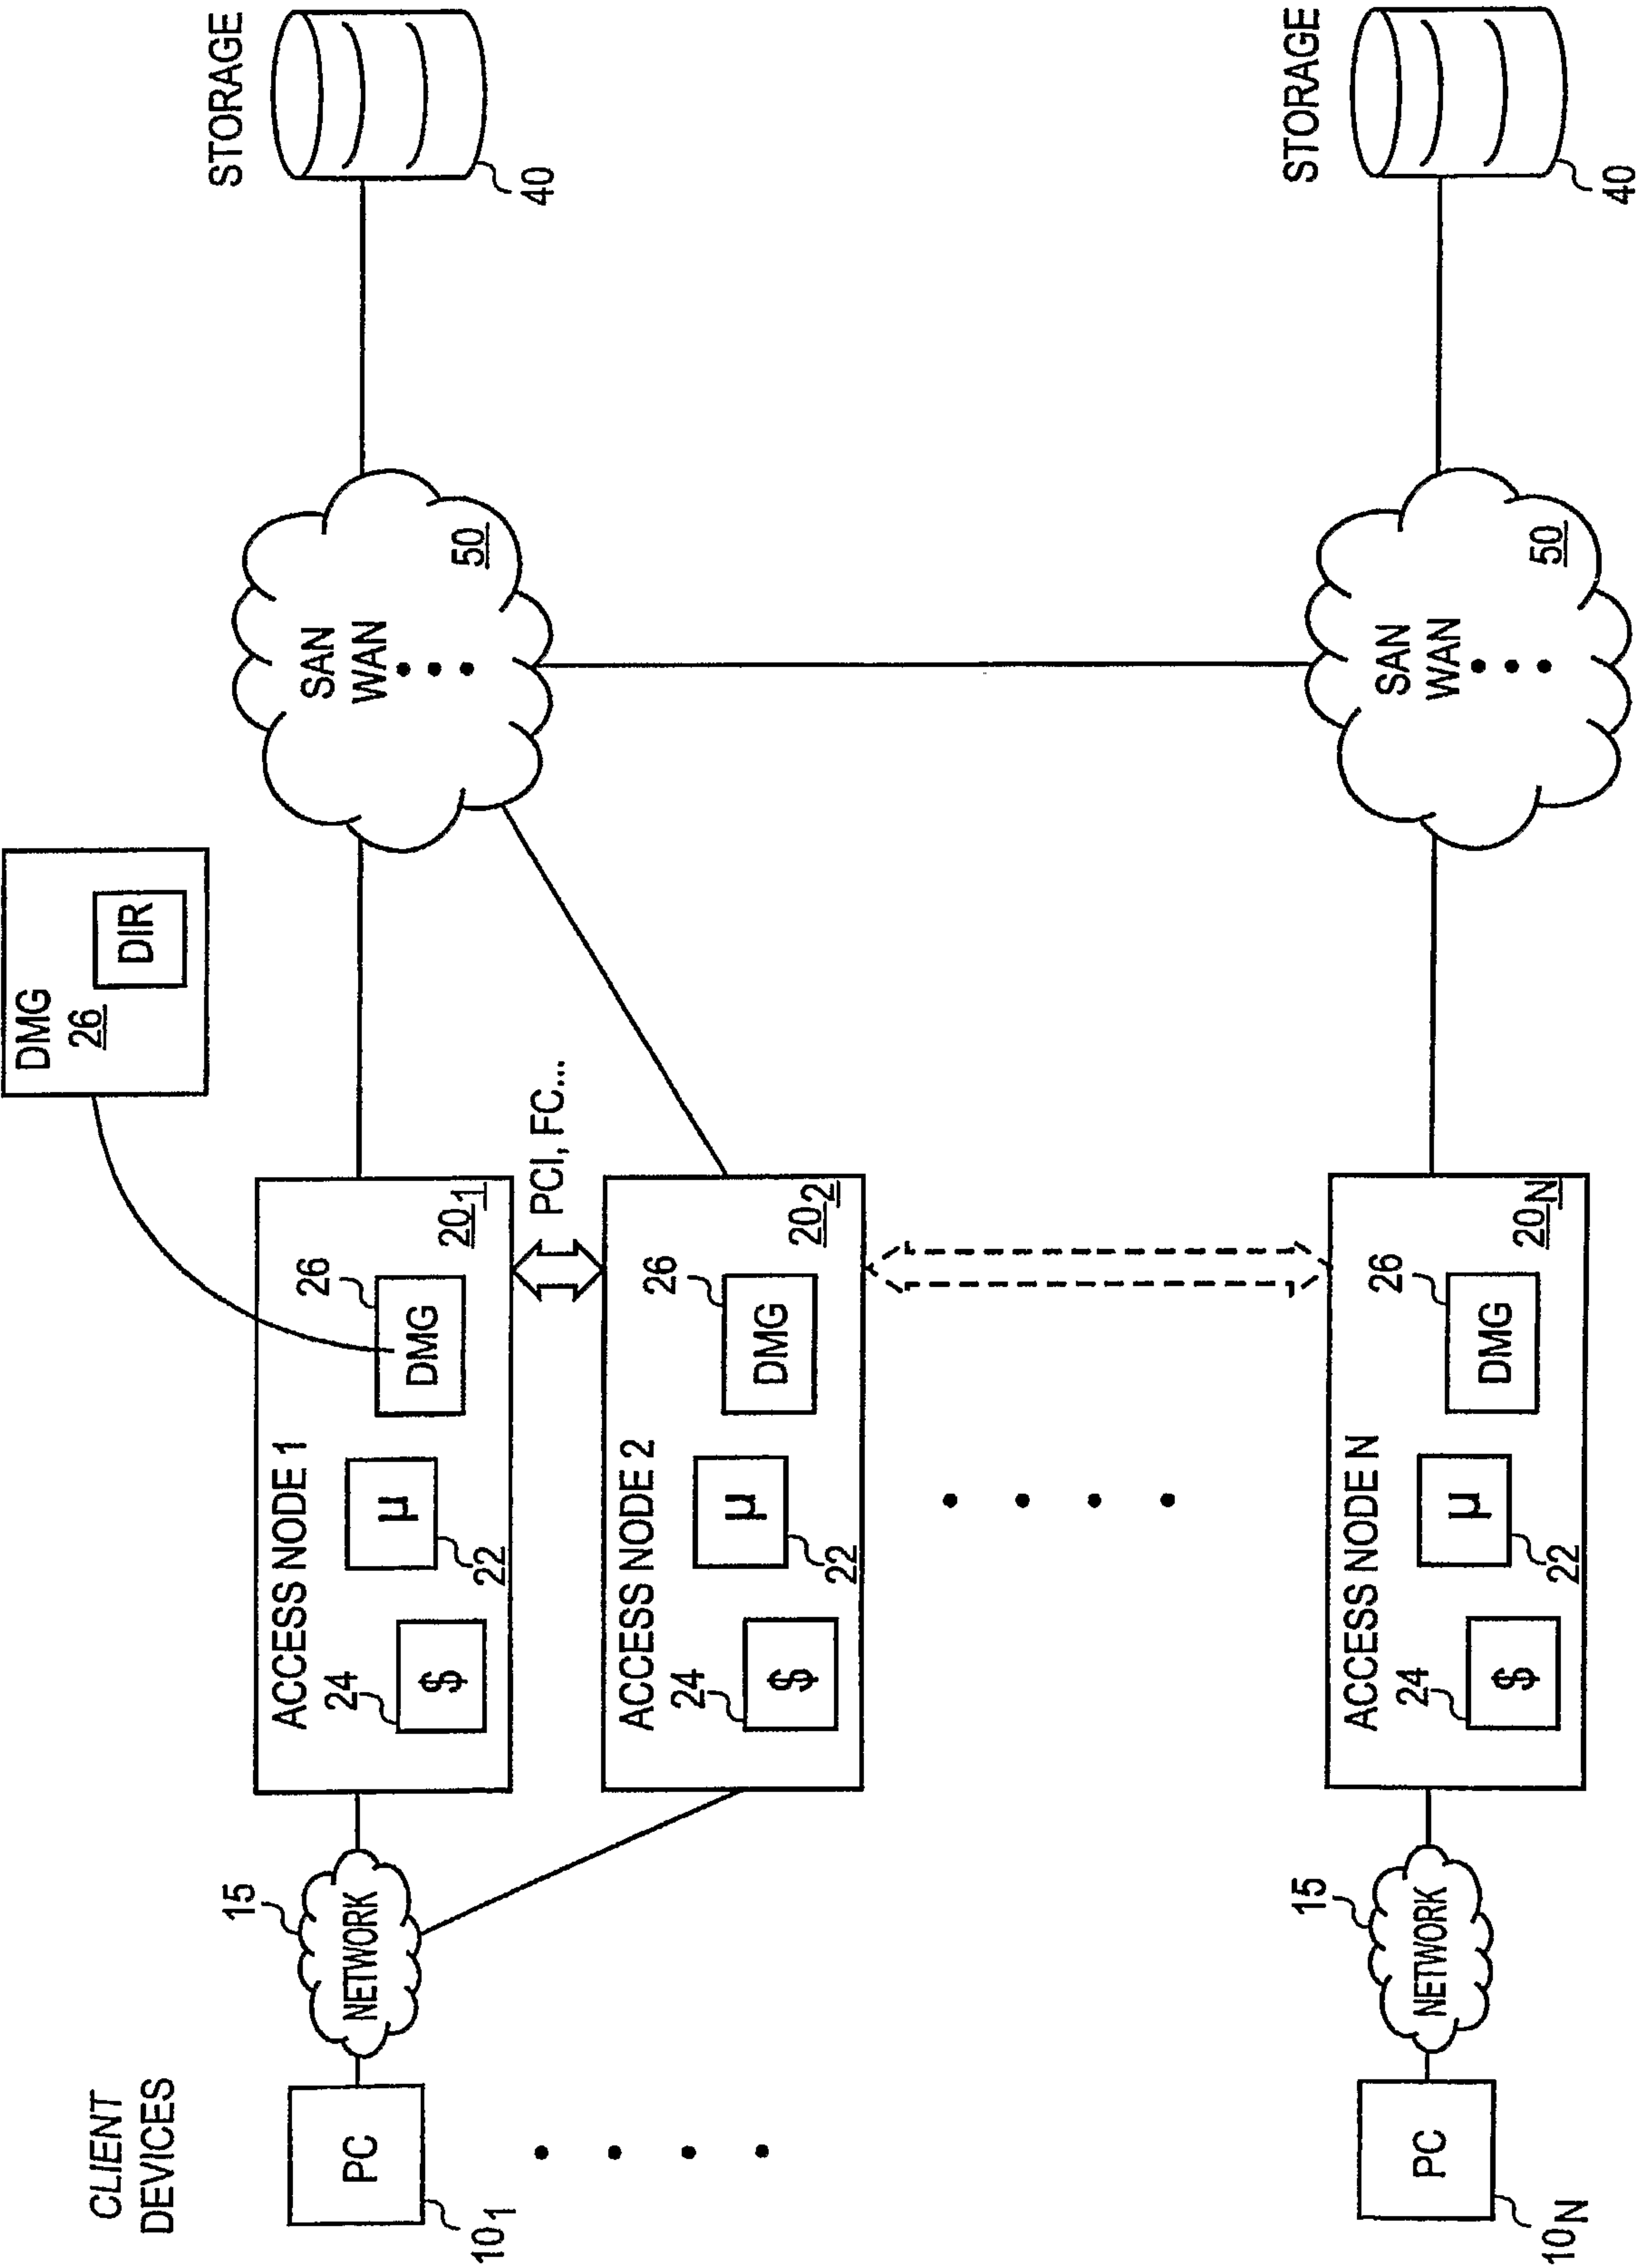
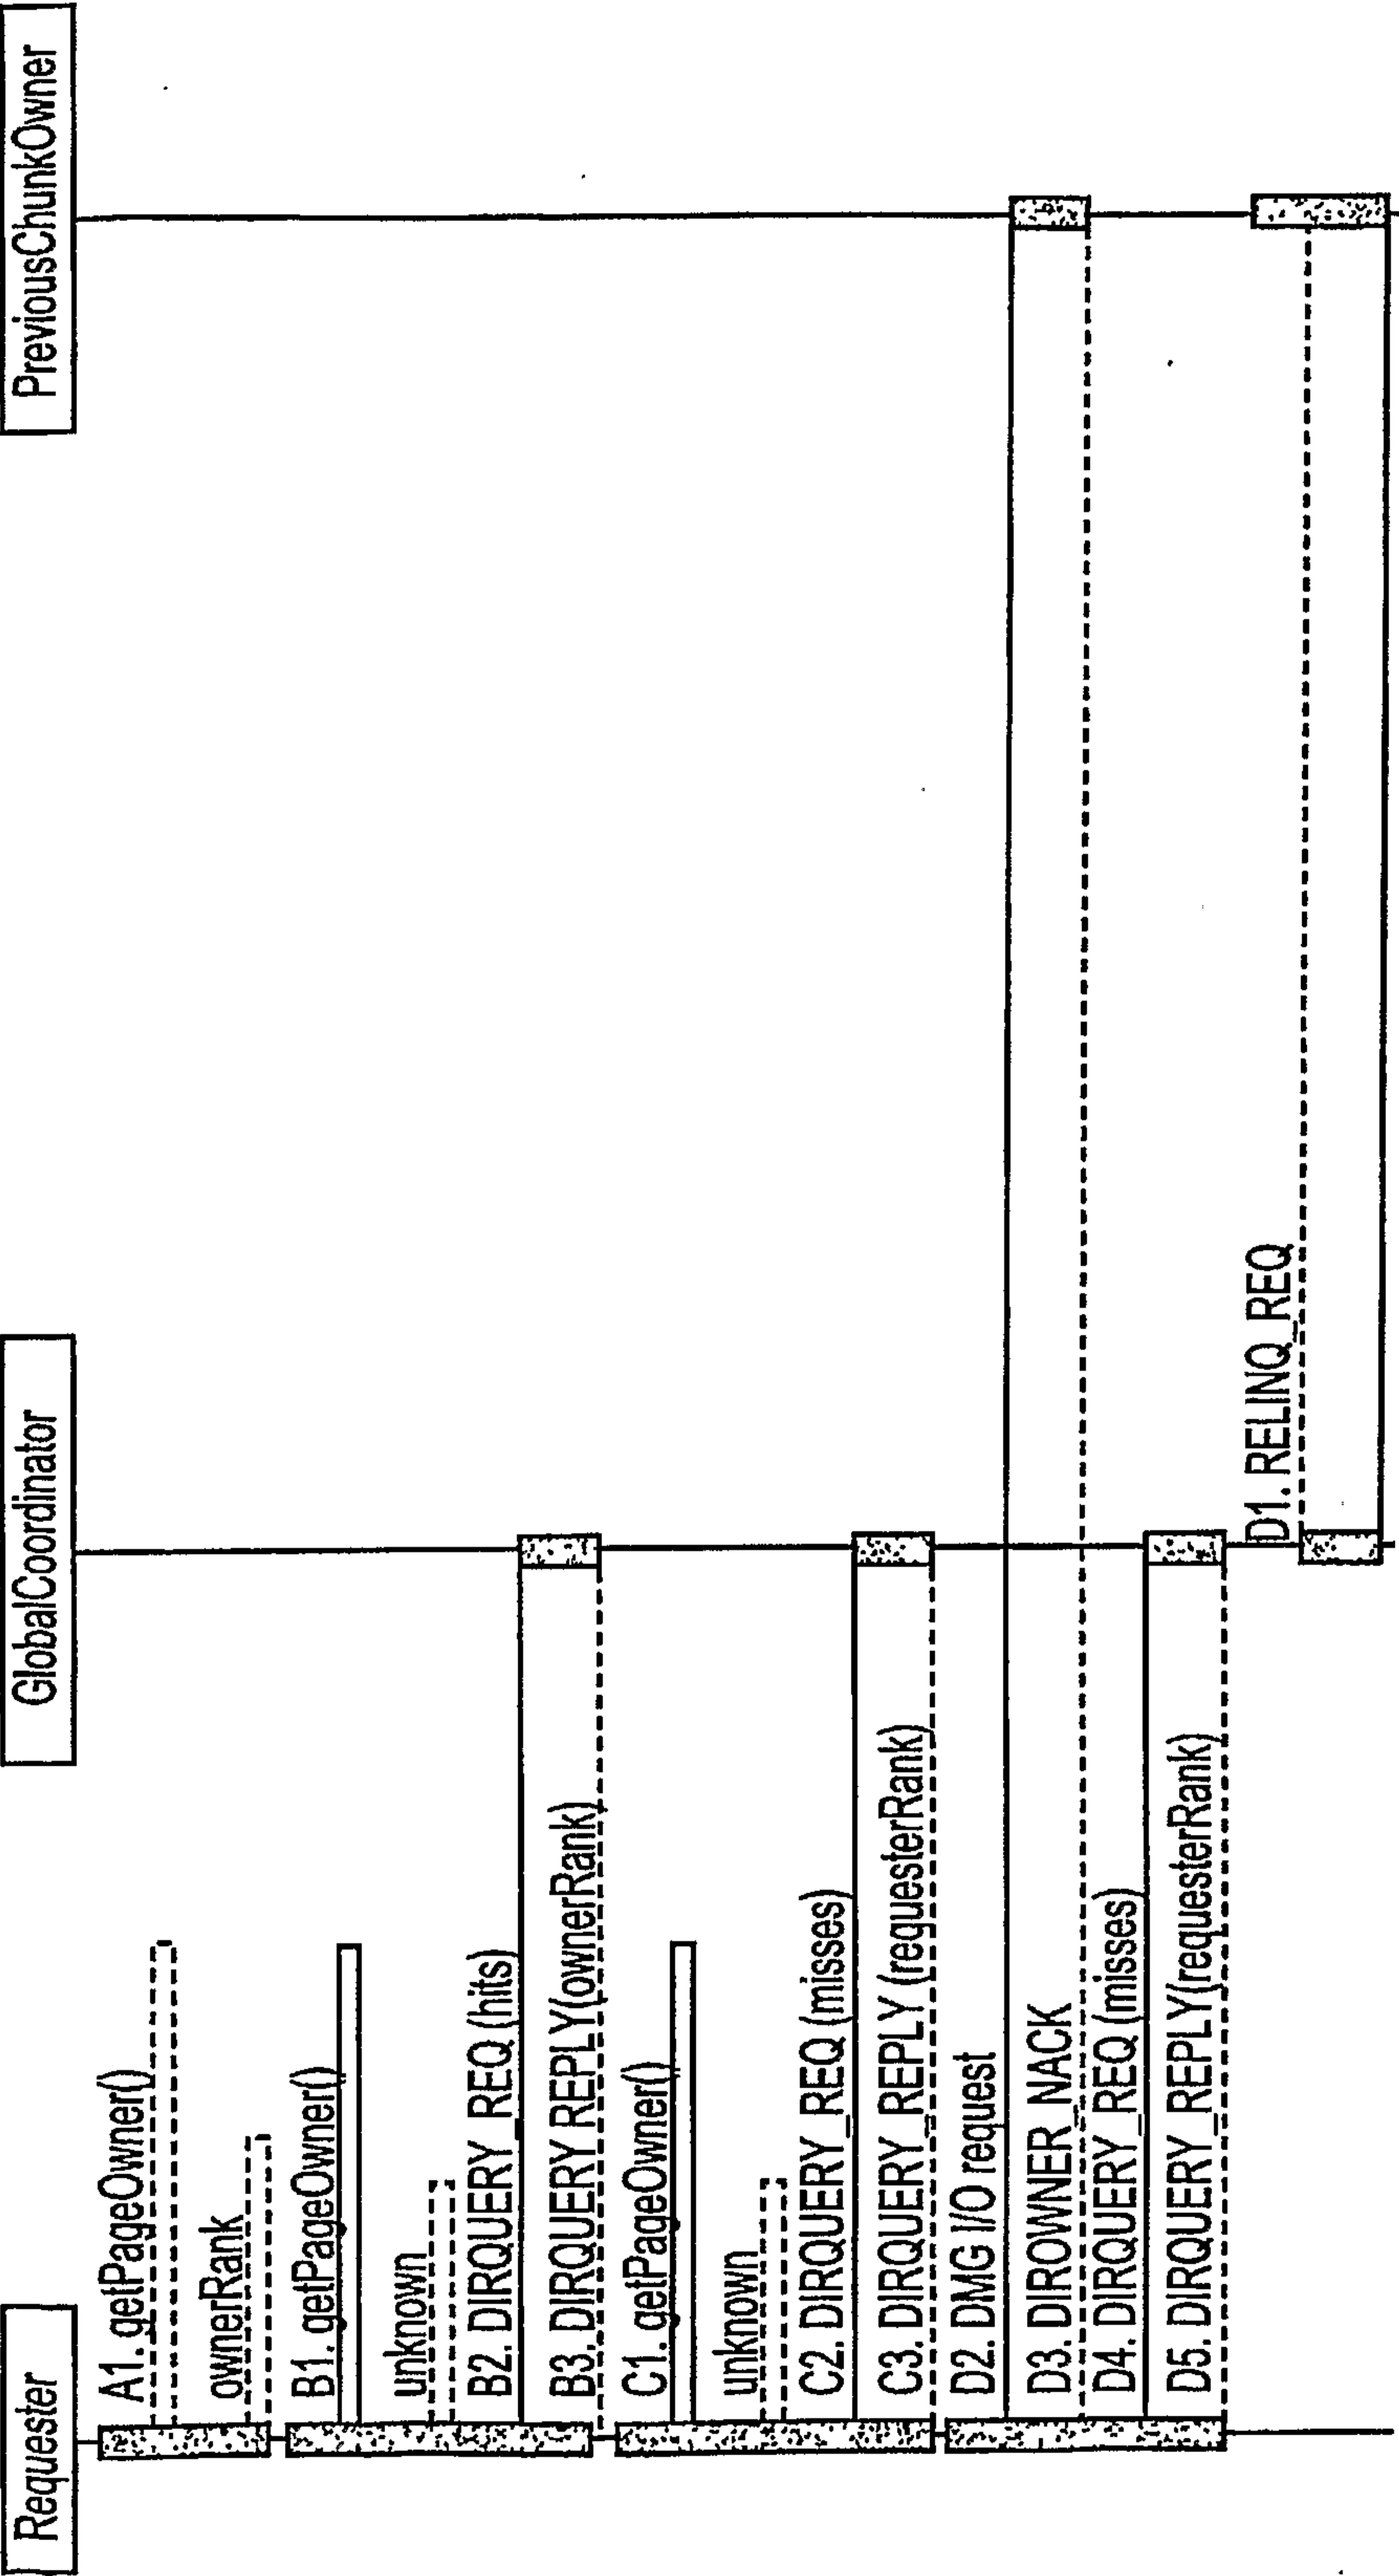


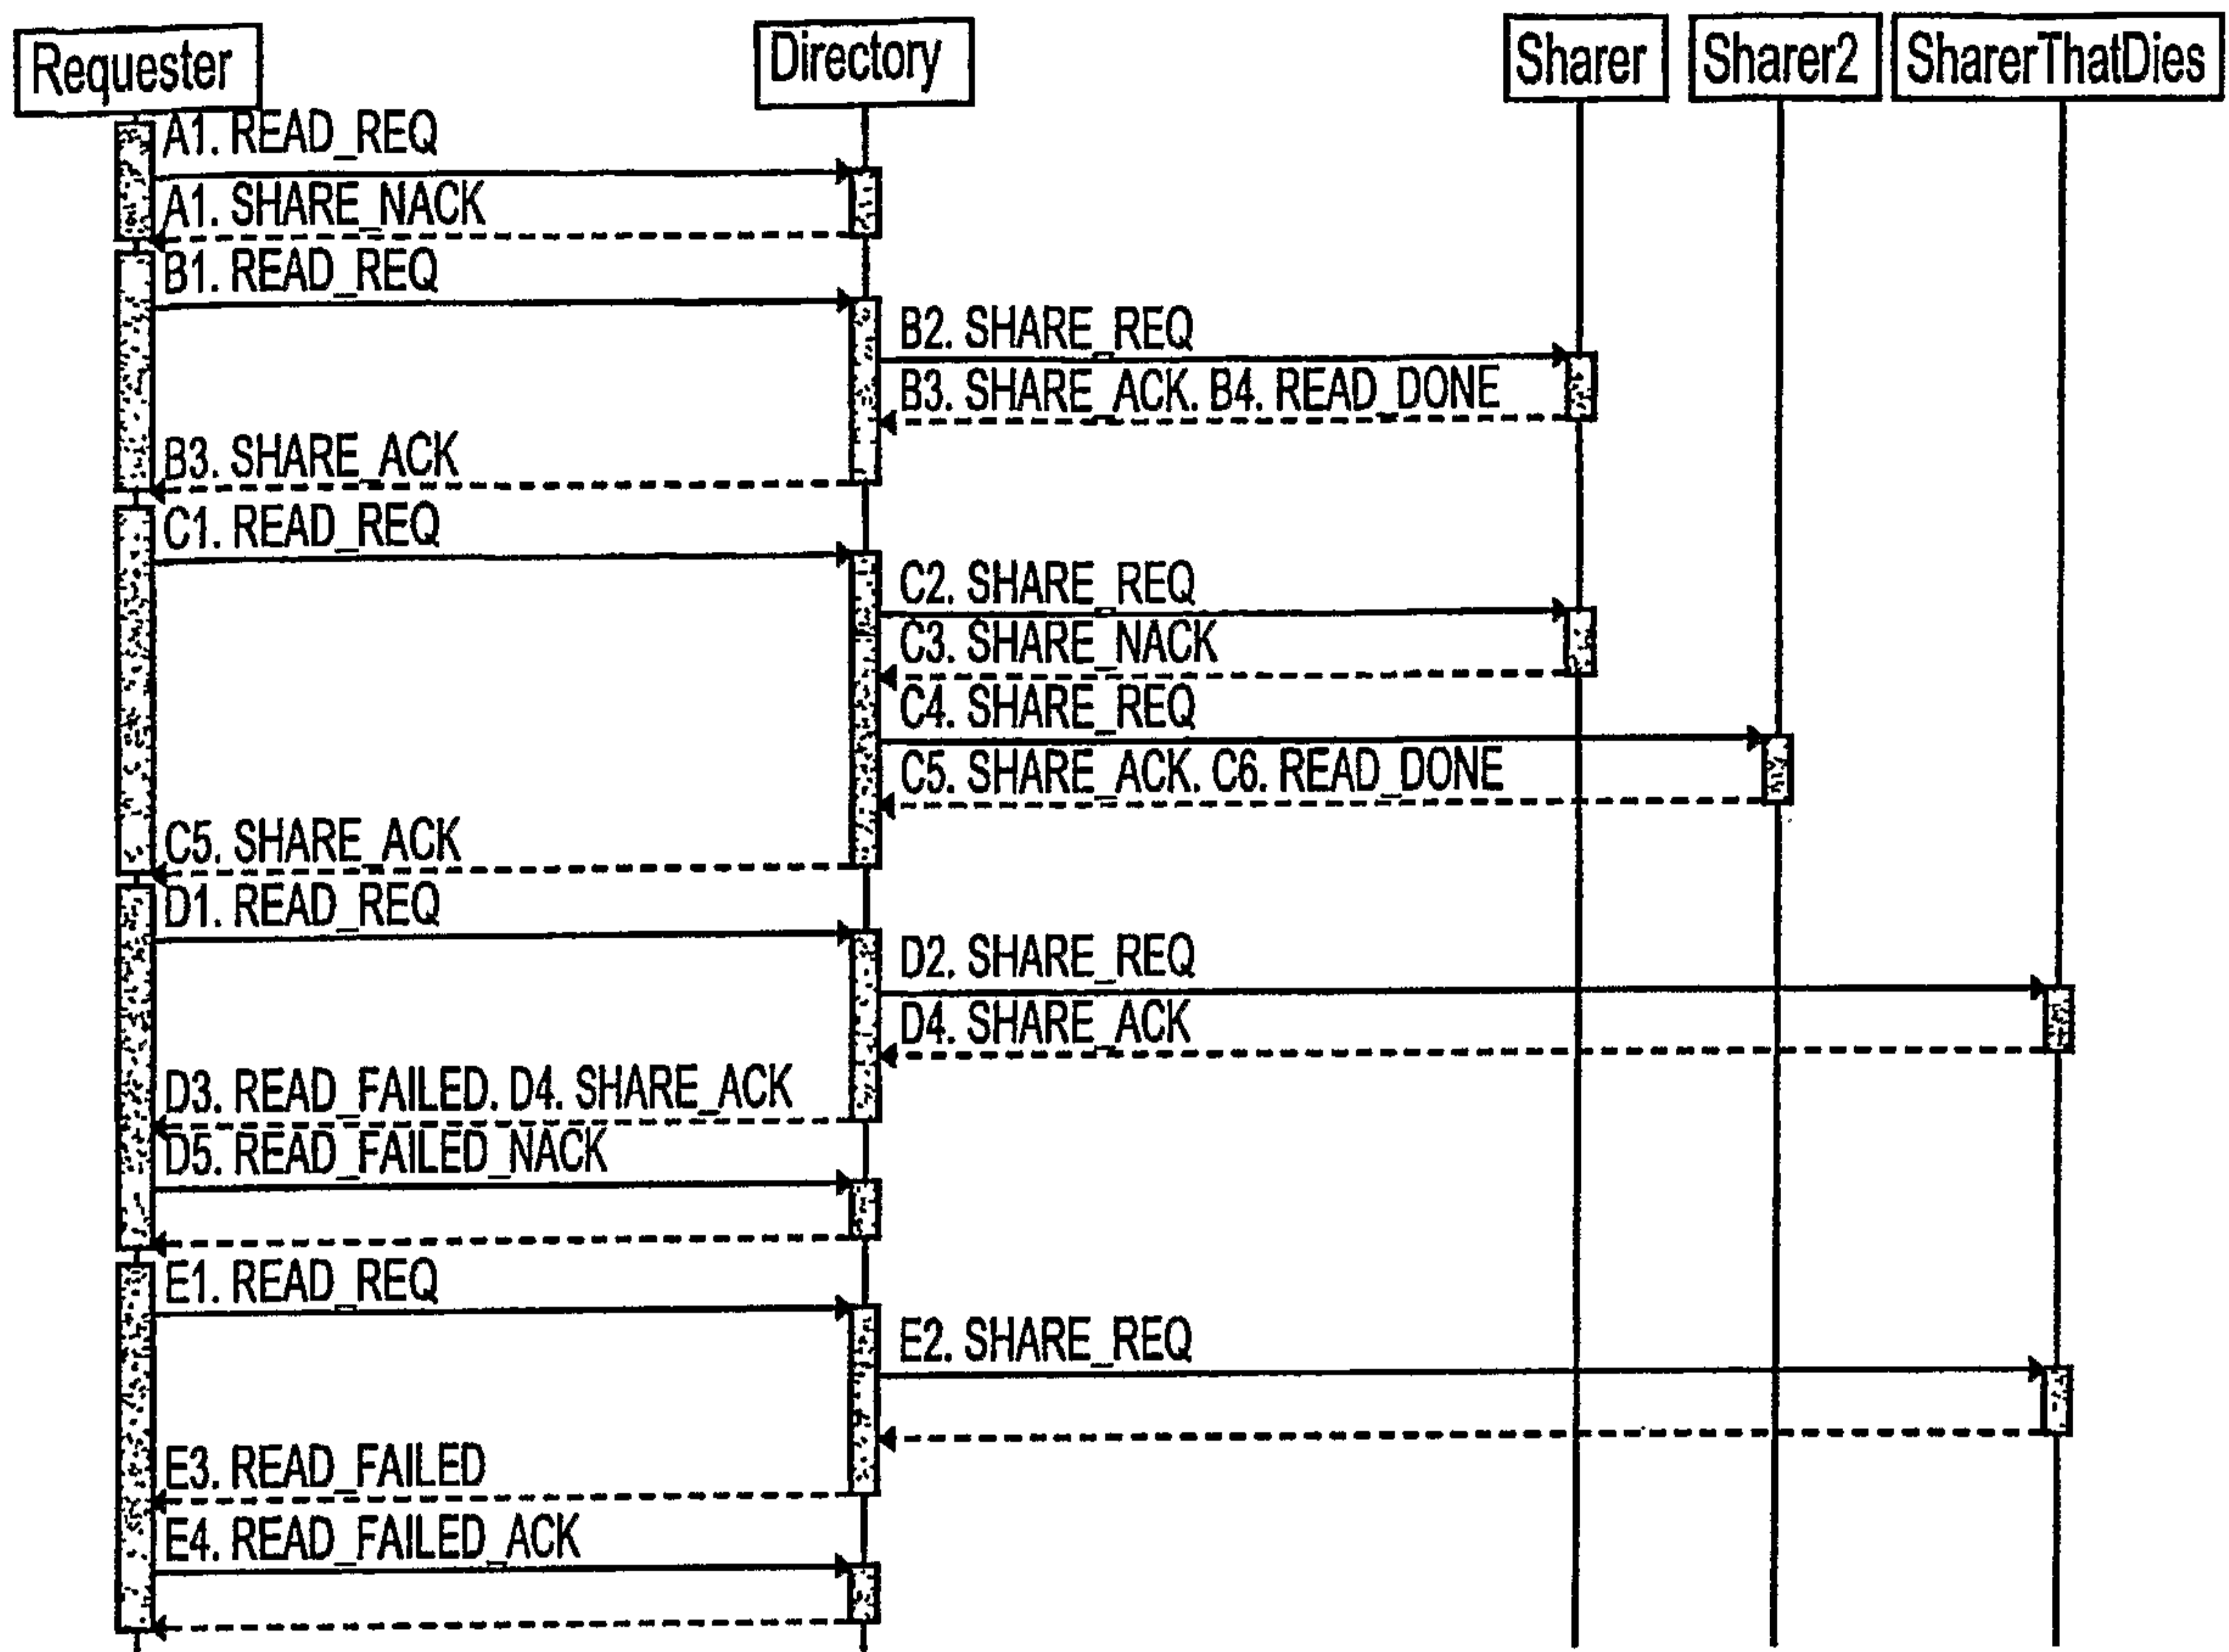
FIG. 1



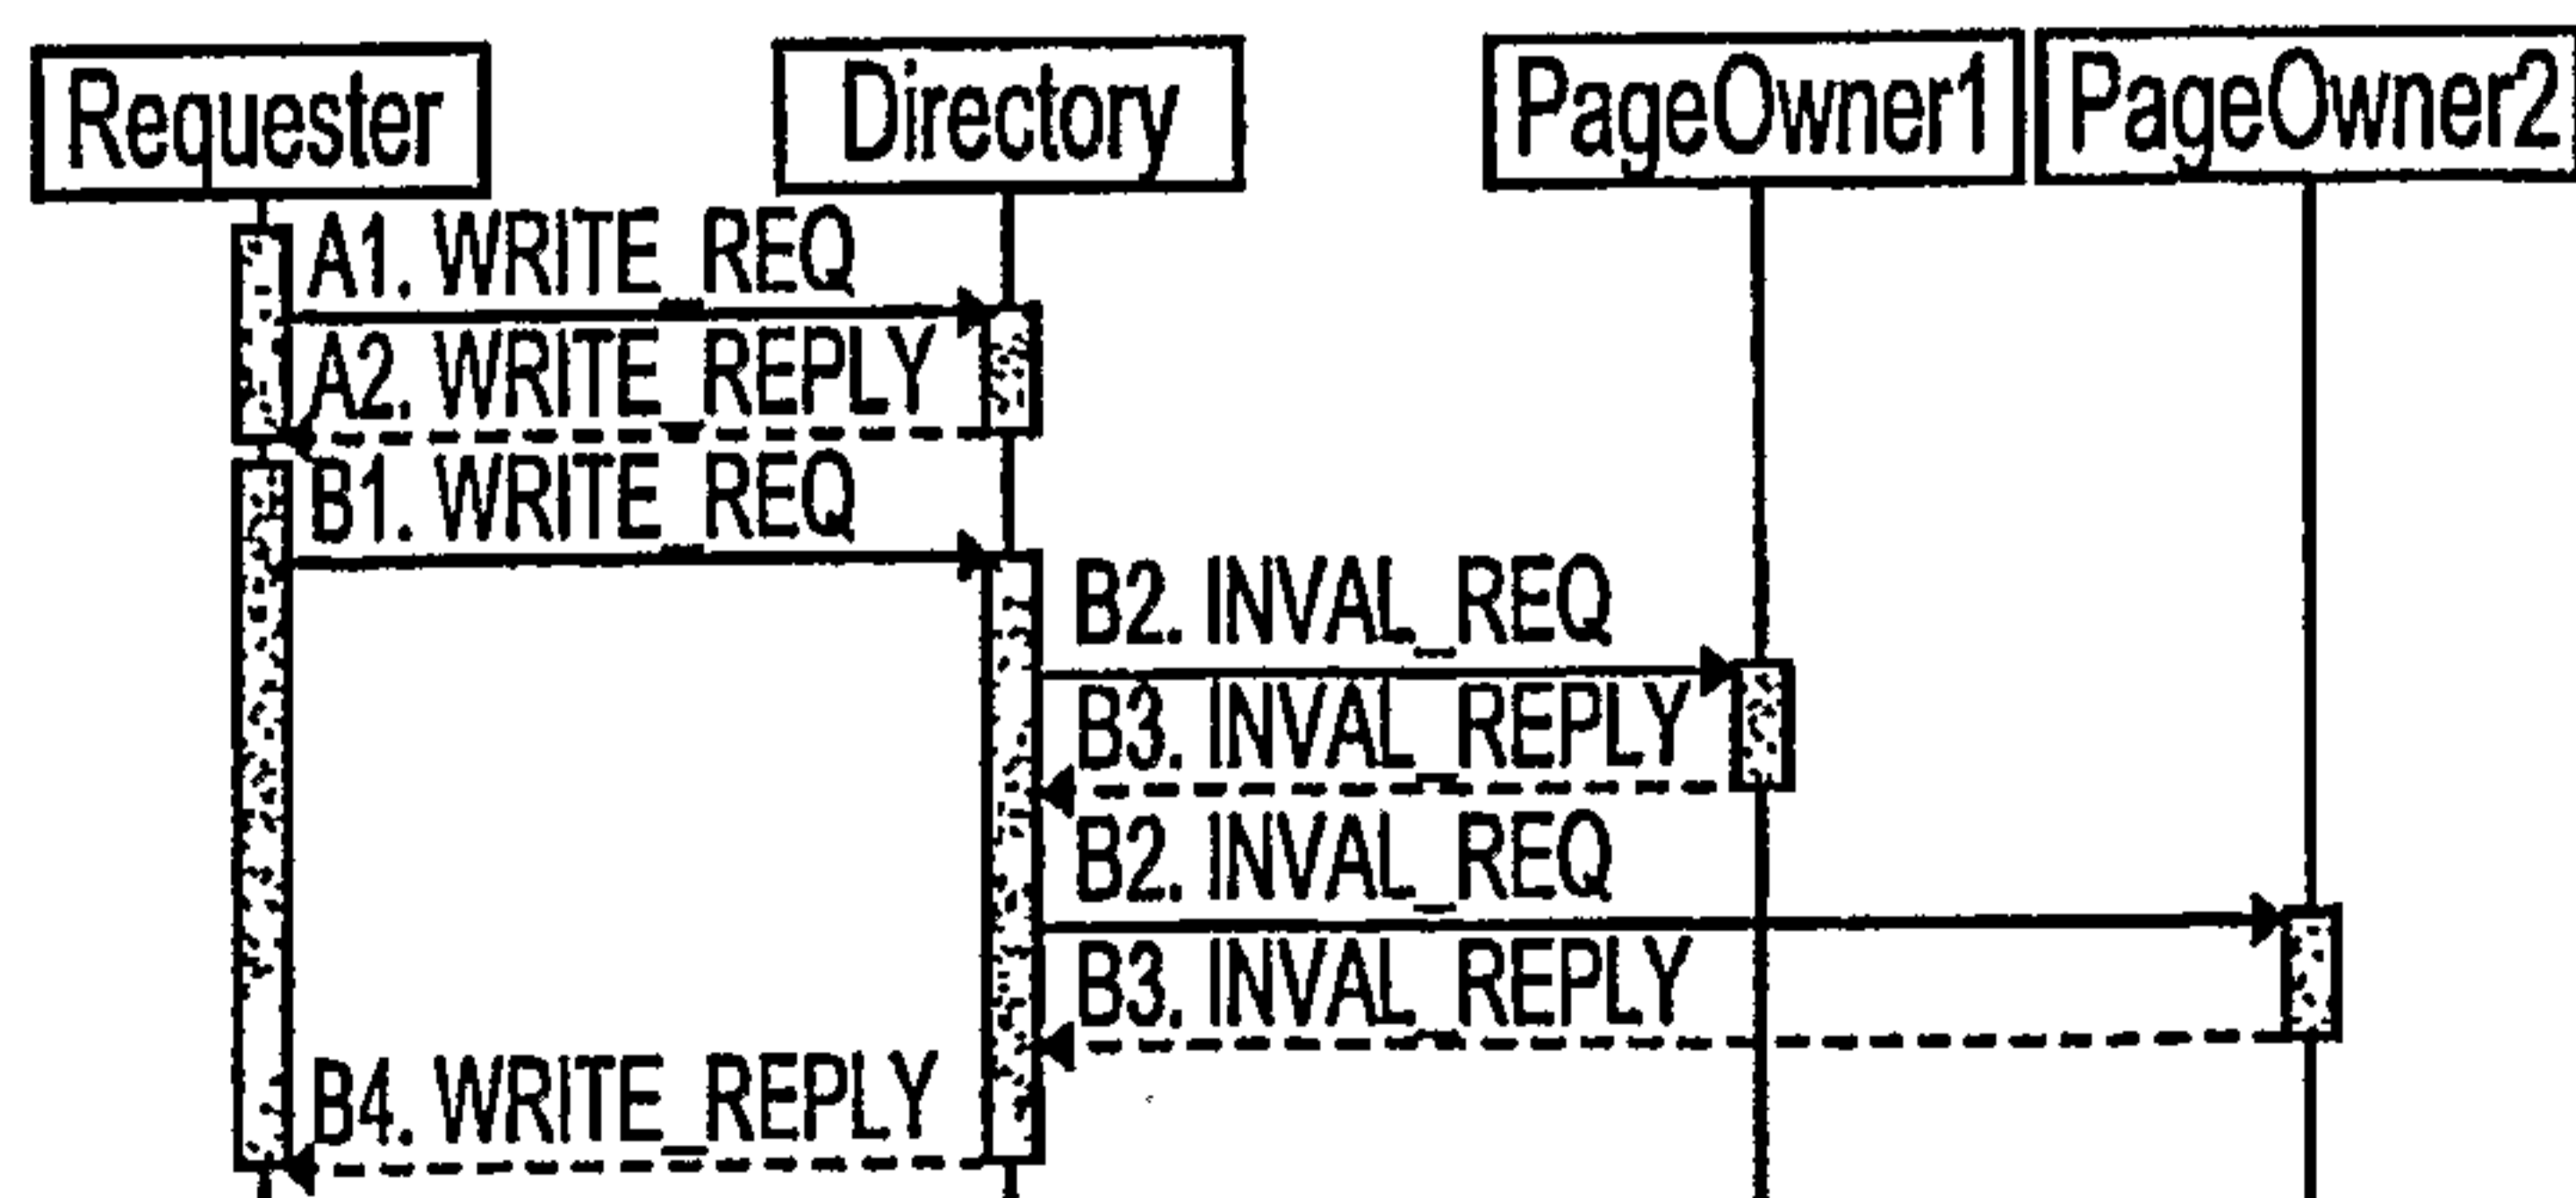
DMG LCDM Directory Ownership Changes

FIG. 2

3/4



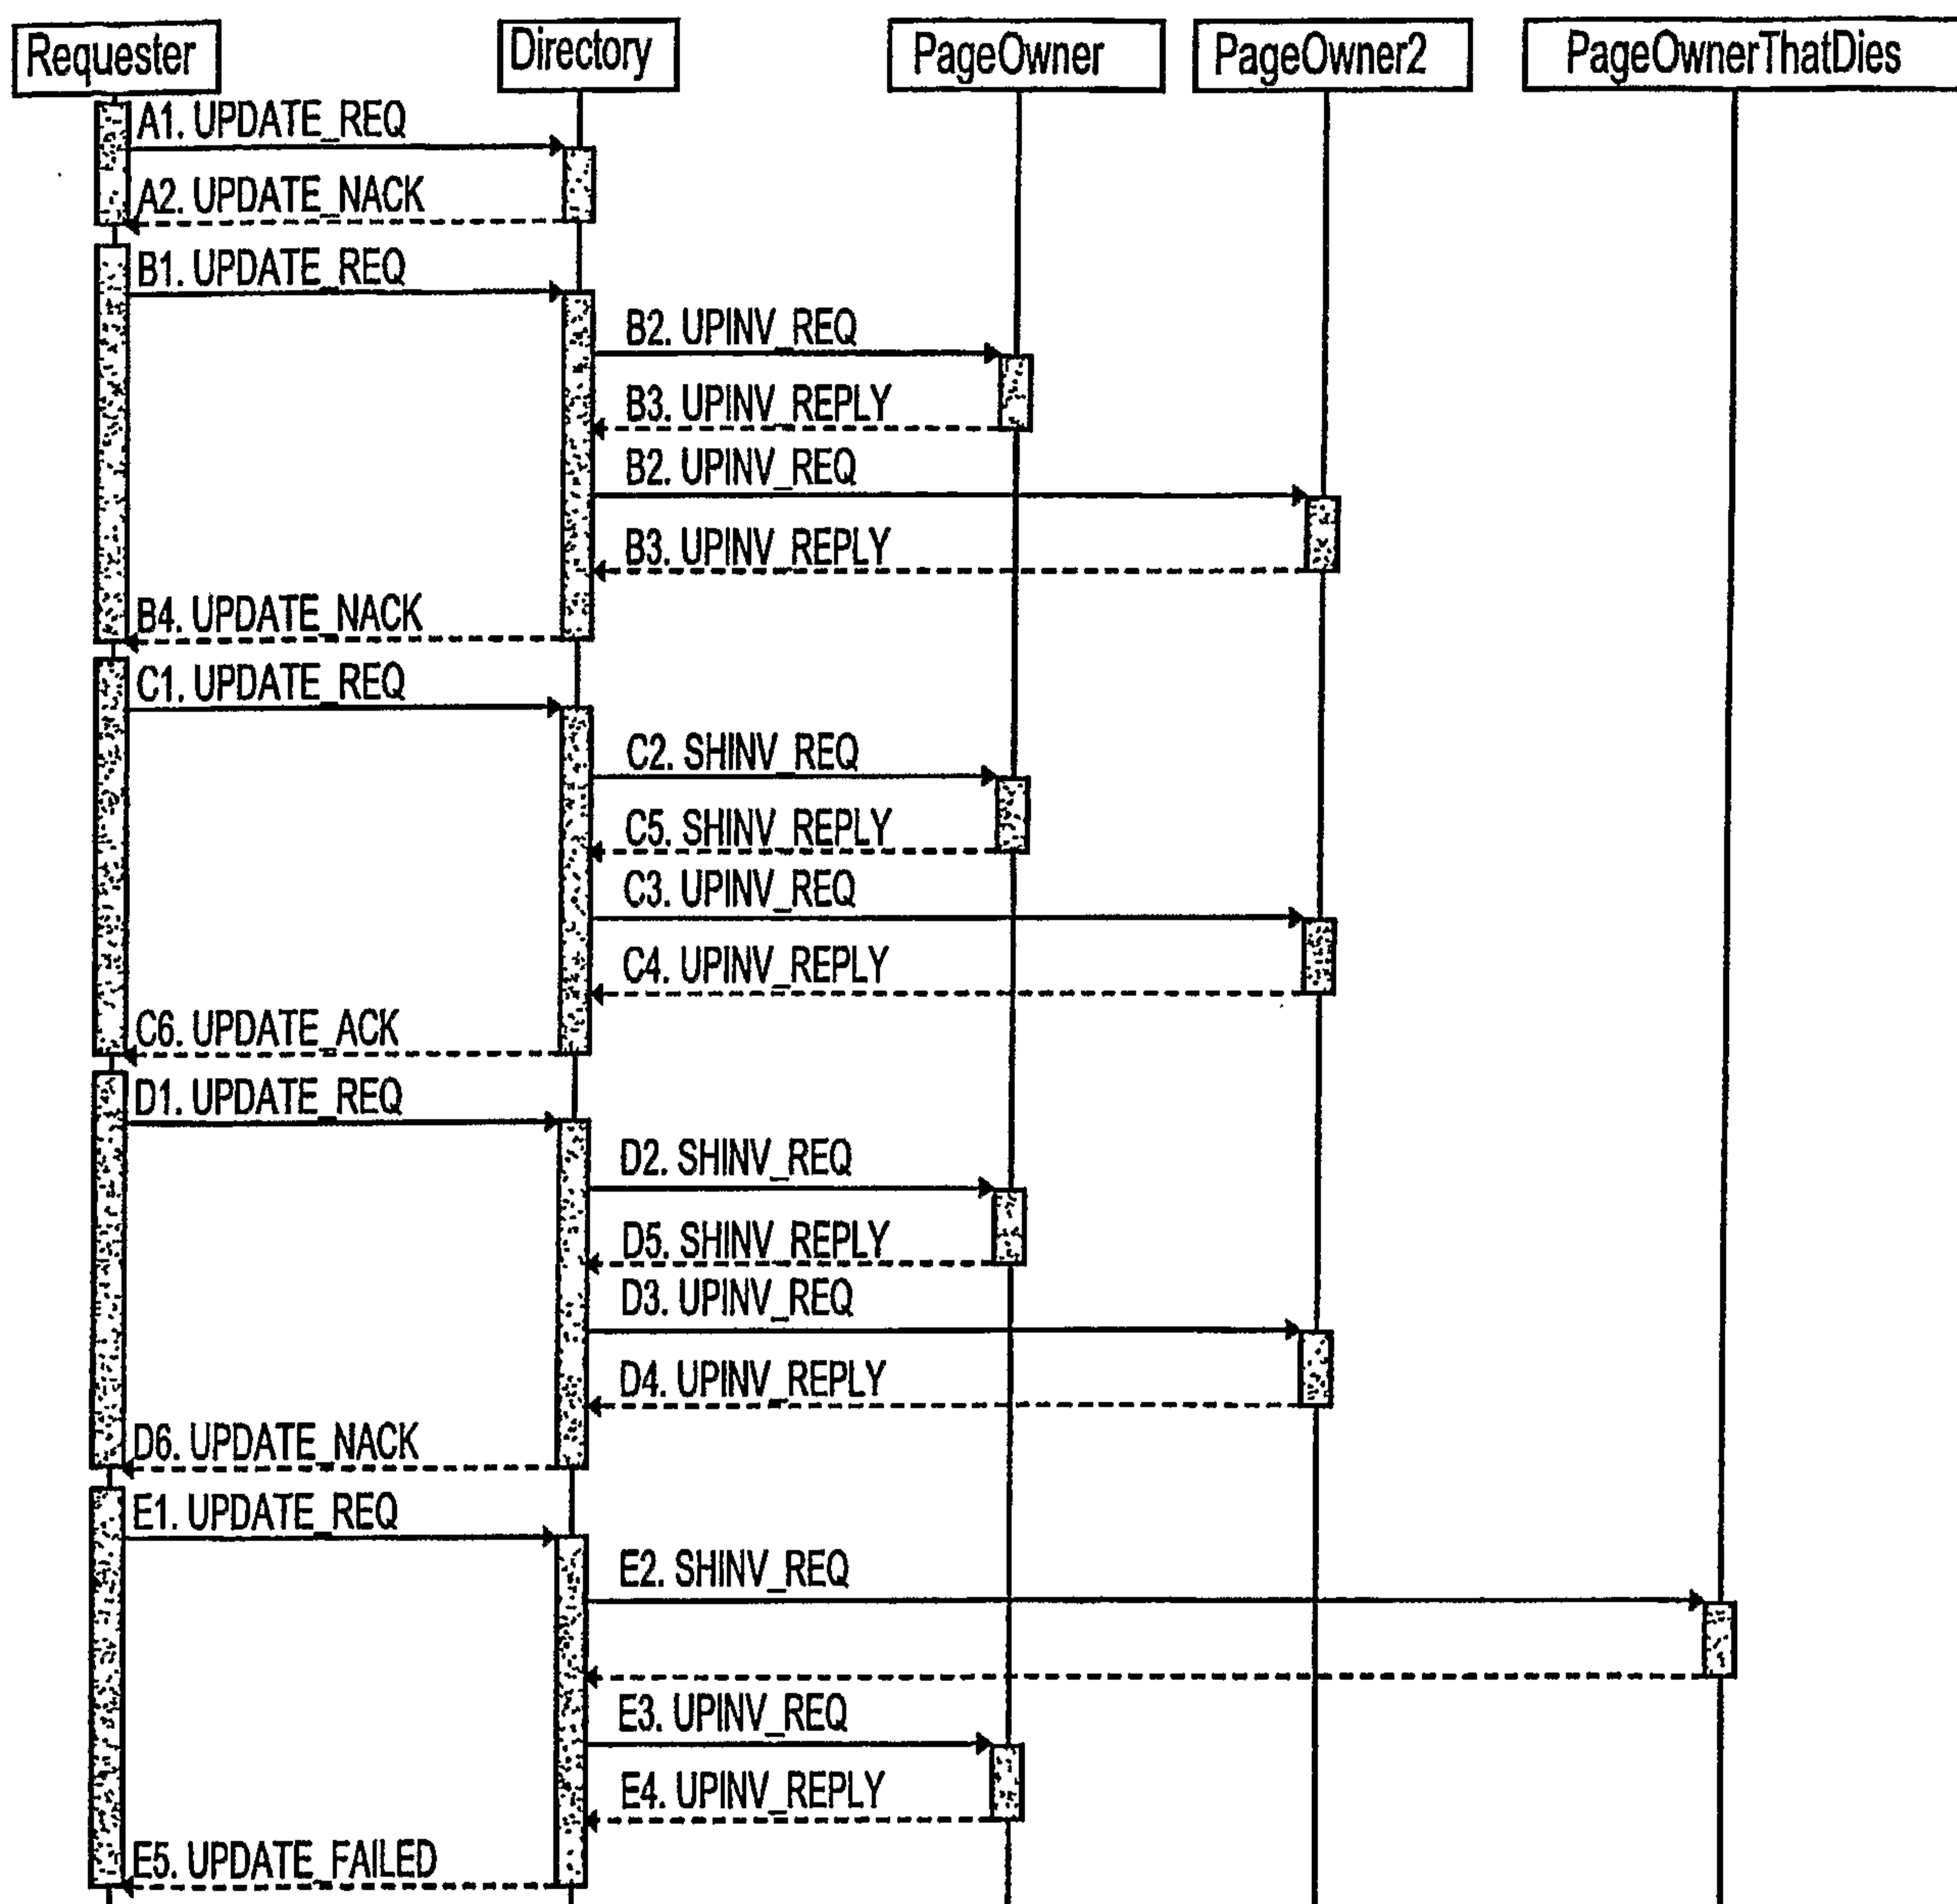
DMG Read Sequence Diagrams

**FIG. 3**

DMG Write Sequence Diagrams

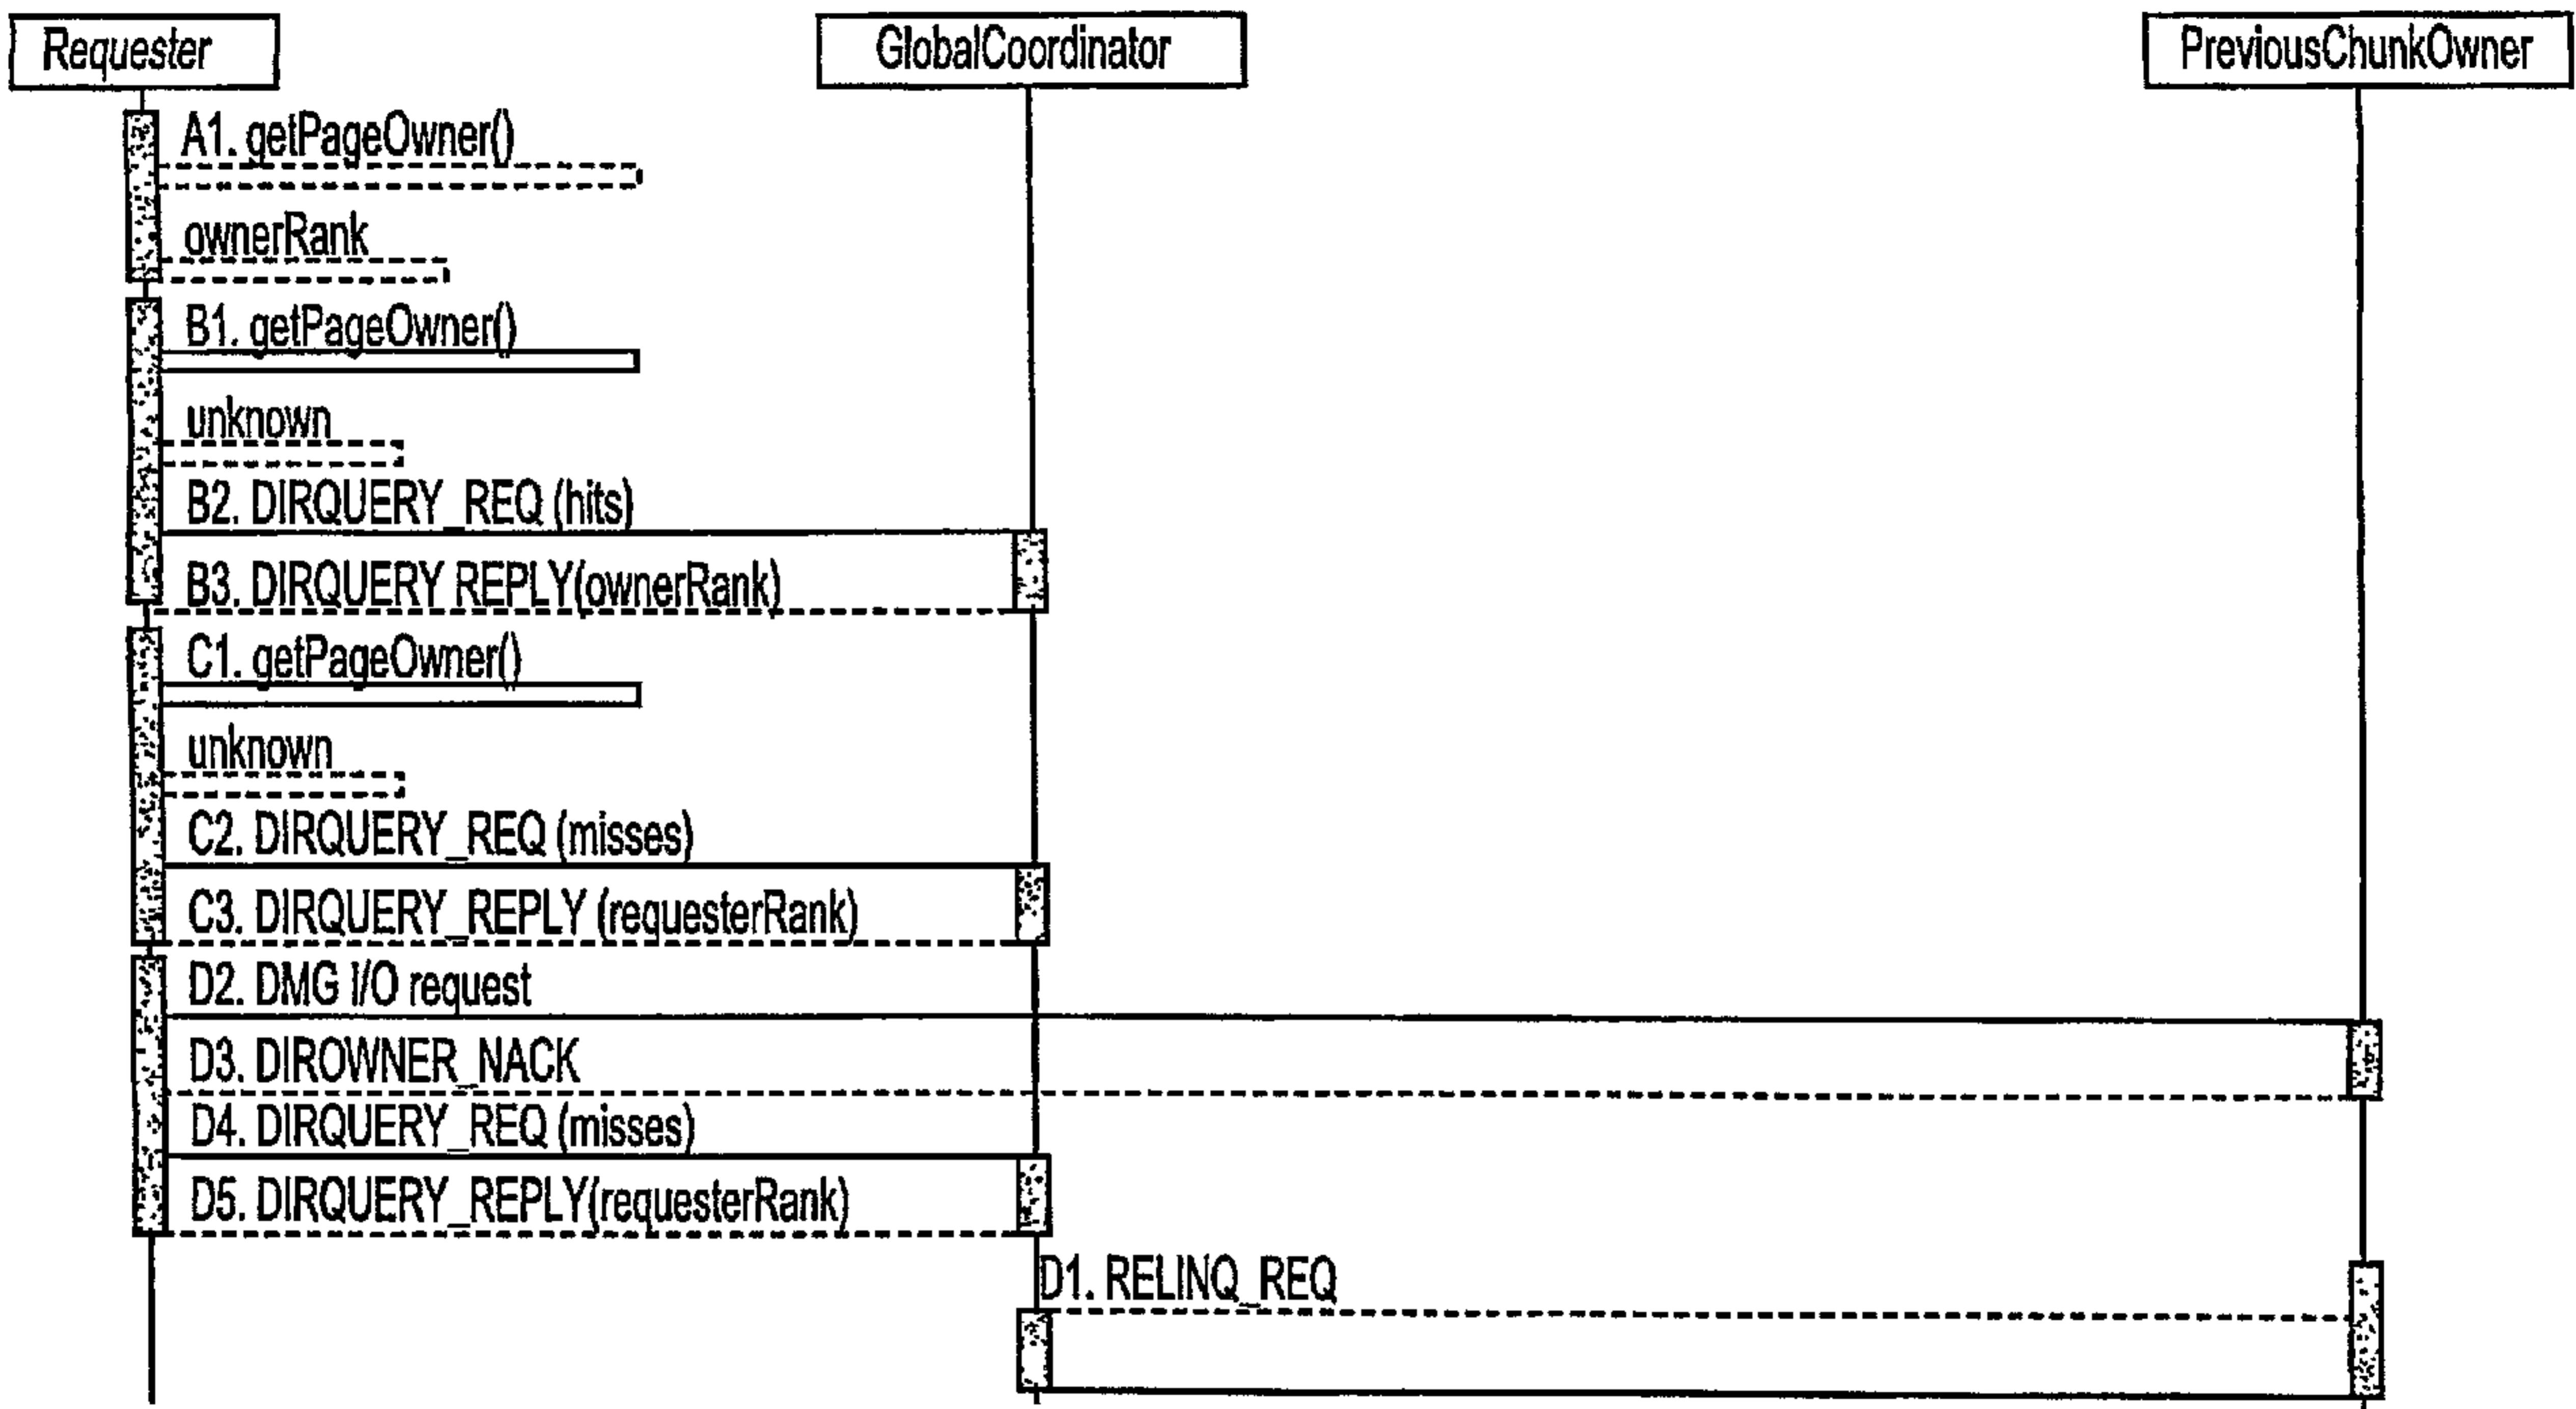
**FIG. 4**

4/4



DMG Update Sequence Diagrams

**FIG. 5**



DMG LCDM Directory Ownership Changes