

Accessing Nearby Copies of Replicated Objects in a Distributed Environment

C. Greg Plaxton¹ Rajmohan Rajaraman¹ Andréa W. Richa²

Abstract

Consider a set of shared objects in a distributed network, where several copies of each object may exist at any given time. To ensure both fast access to the objects as well as efficient utilization of network resources, it is desirable that each access request be satisfied by a copy “close” to the requesting node. Unfortunately, it is not clear how to efficiently achieve this goal in a dynamic, distributed environment in which large numbers of objects are continuously being created, replicated, and destroyed.

In this paper, we design a simple randomized algorithm for accessing shared objects that tends to satisfy each access request with a nearby copy. The algorithm is based on a novel mechanism to maintain and distribute information about object locations, and requires only a small amount of additional memory at each node. We analyze our access scheme for a class of cost functions that captures the hierarchical nature of wide-area networks. We show that under the particular cost model considered: (i) the expected cost of an individual access is asymptotically optimal, and (ii) if objects are sufficiently large, the memory used for objects dominates the additional memory used by our algorithm with high probability. We also address dynamic changes in both the network as well as the set of object copies.

1 Introduction

The advent of high-speed networks has made it feasible for a large number of geographically dispersed computers to cooperate and share objects (e.g, files, words of memory). This has resulted in the implementation of large distributed databases like the World Wide Web on wide-area networks. The large size of the databases and the rapidly growing demands of the users has in turn overloaded the underlying network resources. Hence, an important goal is to make efficient use of network resources when providing access to shared objects.

As one might expect, the task of designing efficient algorithms for supporting access to shared objects over wide-area networks is extremely challenging, both from a practical as well as a theoretical perspective. With respect to any interesting measure of performance (e.g., latency, throughput), the optimal bound achievable by a given network is a complex function of many parameters, including edge delays, edge capacities, buffer space, communication overhead, patterns of user communication, and so on. Ideally, we would like to take all of these factors into account when optimizing performance with respect to a given measure. However, such a task may not be feasible in general because the many network parameters interact in a fairly complex manner. For this reason, we adopt a simplified model in which the combined effect of the detailed network parameter values is assumed to be captured by a single function that specifies the cost of communicating a fixed-length message between any given pair of nodes. We anticipate that analyzing algorithms under this model will significantly aid in the design of practical algorithms for modern distributed networks.

Accessing shared objects. Consider a set \mathcal{A} of m objects being shared by a network G , where several copies of each object may exist. In this paper, we consider the basic problem of *reading* objects

¹Department of Computer Science, University of Texas at Austin, Austin, TX 78712. Supported by the National Science Foundation under Grant No. CCR-9504145. Email: {plaxton, rraj}@cs.utexas.edu.

²School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213. Supported by Bruce Maggs’s National Young Investigator Award under Grant No. CCR-94-57766. Email: aricha@cs.cmu.edu.

in \mathcal{A} . Motivated by the need for efficient network utilization, we seek algorithms that minimize the cost of the read operation. We do not address the *write* operation, which involves the additional consideration of maintaining consistency among the various object copies. The problem of consistency, although an important one, is separate from our main concern, namely, that of studying locality. Our results for the read apply for the write in scenarios where consistency either is not required or is enforced by an independent mechanism.

We differentiate between *shared* and *unshared* copies of objects. A copy is shared if any node can read this copy; it is unshared if only the node which holds the copy may read it. We say that a node u *inserts* (resp., *deletes*) a copy of object A (that u holds) if u declares the copy shared (resp., unshared).

We refer to the set of algorithms for read, insert, and delete operations as an *access scheme*. Any access scheme that efficiently supports these operations incurs an overhead in memory. It is desirable that this overhead be small, not only because of space considerations, but also because low overhead usually implies fast adaptability to changes in the network topology or in the set of object copies.

The main difficulty in designing an access scheme that is efficient with respect to both time and space is the competing considerations of these measures. For example, consider an access scheme in which each node stores the location of each copy of each object in the network. This allows very fast read operations since a node can easily determine the location of the closest copy of any desired object. However, such an access scheme is impractical because: (i) it incurs a prohibitively large memory overhead, and (ii) every node of the network has to be informed whenever a copy of an object is inserted or deleted. At the other extreme, one might consider an access scheme using no additional memory. In this case insert and delete operations are fast, but read operations are costly since it may be necessary to search the entire network in order to locate a copy of some desired object.

Our access scheme. We design a simple randomized access scheme that exploits locality and distributes control information to achieve low overhead in memory. The central part of our access scheme is a mechanism to maintain and locate the addresses of copies of objects. For a single object, say A , we can provide such a mechanism by the following approach. We embed an n -node “virtual” height-balanced tree T one-to-one into the network. Each node u of the network maintains information associated with the copies of A residing in the set of nodes that form the subtree of T rooted at u . Given the embedding of T , the read operation may be easily defined as follows. When a node u attempts to read A , u first checks its local memory for a copy of A or information about copies of A in the subtree of T rooted at u . If u is unable to locate any copy on the basis of local information, it forwards its request to its parent in T .

Naive extensions of the above approach to account for all objects require significant overhead in memory for control information at individual nodes. We overcome this problem by designing a novel method to embed the different trees associated with different objects. Our embedding enables us to define simple algorithms for read, insert, and delete operations, and to prove their efficiency for a class of cost functions that is appropriate for modeling wide-area networks.

The cost model. As indicated above, we assume that a given function determines the cost of communication between each pair of nodes in the network. Our analysis is geared towards a restrictive class of cost functions which we believe to be of practical interest. The precise set of assumptions that we make with respect to the cost function is stated in Section 2. Our primary assumption is that for all nodes x and costs r , the ratio of the number of nodes within cost $2r$ of node x to the number of nodes within cost r of node x is bounded from above and below by constants greater than 1 (unless the entire network is within cost $2r$ of node x , in which case the ratio may be as low as 1).

There are several important observations we can make concerning this primary assumption on the cost function. First, a number of commonly studied fixed-connection network families lead naturally to cost functions satisfying this assumption. For example, fixed-dimension meshes satisfy this assumption if the cost of communication between two nodes is defined as the minimum number of hops between them. As another example, fat-tree topologies can be made to satisfy our assumption if the cost of communication

between two nodes is determined by the total cost of a shortest path between them, where the cost assigned to individual edges grows at an appropriate geometric rate as we move higher in the tree. The latter example is of particular interest here, because of all the most commonly studied fixed-connection network families, the fat-tree may provide the most plausible approximation to the structure of current wide-area networks.

Even so, it is probably inappropriate to attempt to model the Internet, say, with any kind of uniform topology, including the fat-tree. Note that our assumption on the cost function is purely “local” in nature, and allows for the possibility of a network with a highly irregular global structure. This may be the most important characteristic of our cost model.

Performance bounds. We show that our access scheme achieves optimality or near-optimality in terms of several important complexity measures for the restricted class of cost functions discussed above. In particular, our scheme achieves the following bounds:

- The expected cost for any read request is asymptotically optimal.
- If the number of objects that can be stored at each node is q , then the additional memory required is $O(q \log^2 n)$ words whp¹, where a word is an $O(\log n)$ -bit string. Thus, if the objects are sufficiently large, i.e., $\Omega(\log^2 n)$ words, the memory for objects dominates the additional memory.
- The number of nodes that need to be updated upon the addition or removal of a node is $O(\log n)$ expected and $O(\log^2 n)$ whp.
- The expected cost of an insert (resp., delete) operation at node u is $O(C)$ (resp., $O(C \log n)$), where C is the maximum cost of communicating a single word message between any two nodes.

An obvious shortcoming of our analysis is that it only applies to the restricted class of cost functions discussed above. While we do not expect that all existing networks fall precisely within this restricted class, we stress that: (i) our access scheme is well-defined, and functions correctly, for arbitrary networks, and (ii) we expect that our access scheme would have good practical performance on any existing network. (Although we have not attempted to formalize any results along these lines, it seems clear that our performance bounds would only degrade significantly in the presence of a large number of non-trivial violations of our cost function assumptions.)

Related work. The basic problem of sharing memory in distributed systems has been studied extensively in different forms. Most of the earlier work in this area, e.g., emulations of PRAM on completely-connected distributed-memory machines (e.g., [10, 16]) or bounded-degree networks (e.g., [14]), and algorithms for providing concurrent access to a set of shared objects [13], assume that each of the nodes of the network has knowledge of a hash function that indicates the location of any copy of any object.

The basic problem of locating an object arises in every distributed system [11], and was formalized by Mullender and Vitányi [12] as an instance of the distributed matchmaking problem. Awerbuch and Peleg [3], and subsequently Bartal et al. [4] and Awerbuch et al. [1], give near-optimal solutions in terms of cost to a related problem by defining sparse-neighborhood covers of graphs. Their studies do not address the overhead due to control information and hence, natural extensions of their results to our problem may require an additional memory of m words at some node. However, we note that their schemes are designed for arbitrary cost functions, whereas we have focused on optimizing performance for a restricted class of cost functions.

In recent work, access schemes for certain Internet applications have been described in [8, 9, 17]. Some of the ideas in our scheme are similar to those in [17]; however, the two schemes differ considerably in the details. Moreover, the schemes of [8] and [17] have not been analyzed. As in our study, the results

¹We use the abbreviation “whp” throughout the paper to mean “with high probability” or, more precisely, “with probability $1 - n^{-c}$, where n is the number of nodes in the network and c is a constant that can be set arbitrarily large by appropriately adjusting other constants defined within the relevant context.”

of [9] concerning locality assume a restricted cost model. However, their cost model, which is based on the ultrametric, is different from ours. Also, their algorithms are primarily designed for problems associated with “hot spots” (i.e., popular objects).

A closely related problem is that of designing a dynamic routing scheme for networks [2, 6]. Such a scheme involves maintaining routing tables at different nodes of the network in much the same way as our additional memory. However, in routing schemes the size of additional memory is a function of network size, i.e., n , while in our problem the overhead is primarily a function of m . Straightforward generalizations of routing schemes result in access schemes that require an additional memory of m words at each node.

The remainder of this paper is organized as follows. Section 2 defines the model of computation. Section 3 formally describes our access scheme. Section 4 contains a formal statement of the main results. Section 5 analyzes the algorithm and establishes the main results. Section 6 discusses directions for future research.

2 Model of Computation

We consider a set V of n nodes, each with its own local memory, sharing a set \mathcal{A} of $m = \text{poly}(n)$ objects. We define our model of computation by characterizing the following aspects of the problem: (i) objects, (ii) communication, (iii) local memory, (iv) local computation, and (v) complexity measures.

Objects. Each object A has a unique $(\log m)$ -bit identification. For i in $[\log m]$, we denote the i th bit of the identification of A by A^i . (For any positive integer x , we use $[x]$ to denote the set $\{0, \dots, x - 1\}$.) Each object A consists of $\ell(A)$ words, where a word is an $O(\log n)$ -bit string.

Communication. Nodes communicate with one another by means of messages; each message consists of at least one word. We assume that the underlying network supports reliable communication.

We define the cost of communication by a function $c : V^2 \mapsto \mathbf{R}$. For any two nodes u and v in V , $c(u, v)$ is the cost of transmitting a single-word message from u to v . We assume that c is symmetric and satisfies the triangle inequality. We also assume for simplicity that for u, v , and w in V , $c(u, v)$ equals $c(u, w)$ iff v equals w .

The cost of transmitting a message of length ℓ from node u to node v is given by $f(\ell)c(u, v)$, where $f : \mathbf{N} \mapsto \mathbf{R}^+$ is any non-decreasing function such that $f(1)$ equals 1.

Given any u in V and any real r , let $M(u, r)$ denote the set $\{v \in V : c(u, v) \leq r\}$. We refer to $M(u, r)$ as the *ball* of *radius* r around u . We assume that there exist real constants $\delta > 8$ and Δ such that for any node u in V and any real $r \geq 1$, we have:

$$\min\{\delta|M(u, r)|, n\} \leq |M(u, 2r)| \leq \Delta|M(u, r)| \tag{1}$$

Local Memory. We partition the local memory of each node u into two parts. The first part, the *main memory*, stores objects. The second part, the *auxiliary memory*, is for storing possible control information.

Local Computation. There is no cost associated with local computation. (Although the model allows an arbitrary amount of local computation at zero cost, our algorithm does not perform any particularly complex local operations.)

Complexity measures. We evaluate any solution on the basis of four different complexity measures. The first measure is the cost of reading an object. The second measure is the size of the auxiliary memory at any node. The remaining two measures concern the dynamic nature of the problem, where we address the complexity of inserting or deleting a copy of an object and adding or removing a network node. The third measure is the cost of inserting or deleting a copy of an object. The fourth measure is *adaptability*, which is defined as the number of nodes whose auxiliary memory is updated upon the addition or removal of a node. (Our notion of adaptability is analogous to that of [6].)

3 The Access Scheme

In this section, we present our access scheme for shared objects. We assume that n is a power of 2^b , where b is a fixed positive integer to be specified later. For each node x in V , we assign a label independently and uniformly at random from $[n]$. For i in $[\log n]$, let x^i denote the i th bit of the label of x . Note that the label of a node x is independent of the $(\log n)$ -bit unique identification of the node. For all x in V (resp., A in \mathcal{A}), we define $x[i] = x^{(i+1)b-1} \dots x^{ib}$ (resp., $A[i] = A^{(i+1)b-1} \dots A^{ib}$), for i in $[(\log n)/b]$. We also assign a total order to the nodes in V , given by the bijection $\beta : V \rightarrow [n]$. We partition the auxiliary memory of each node into two parts, namely the *neighbor table* and the *pointer list* of the node.

Neighbor table. For each node x , the neighbor table of x consists of $(\log n)/b$ levels. The i th level of the table, i in $[(\log n)/b]$, consists of *primary*, *secondary*, and *reverse* (i, j) -neighbors, for all j in $[2^b]$. The *primary* (i, j) -neighbor y of x is such that $y[k] = x[k]$ for all k in $[i]$, and either: (i) $i < (\log n)/b - 1$ and y is the node of minimum $c(x, y)$ such that $y[i] = j$, if such a node exists, or (ii) y is the node with largest $\beta(y)$ among all nodes z such that $z[i]$ matches j in the largest number of rightmost bits. Let d be a fixed positive integer, to be specified later. Let y be the primary (i, j) -neighbor of x . If $y[i] = j$, then let $W_{i,j}$ denote the set of nodes w in $V \setminus \{y\}$ such that $w[k] = x[k]$, for k in $[i]$, $w[i] = j$, and $c(x, w)$ is at most $d \cdot c(x, y)$. Otherwise, let $W_{i,j}$ be the empty set. The set of *secondary* (i, j) -neighbors of x is the subset U of $\min\{d, |W_{i,j}|\}$ nodes u with minimum $c(x, u)$ in $W_{i,j}$; that is, $c(x, u)$ is at most $c(x, w)$, for all w in $W_{i,j}$, and for all u in U . A node w is a *reverse* (i, j) -neighbor of x iff x is a primary (i, j) -neighbor of w .

Pointer list. Each node x also maintains a pointer list $Ptr(x)$ with pointers to copies of some objects in the network. Formally, $Ptr(x)$ is a set of triples (A, y, k) , where A is in \mathcal{A} , y is a node that holds a copy of A , and k is an upper bound on the cost $c(x, y)$. We maintain the invariant that there is at most one triple associated with any object in $Ptr(x)$. The pointer list of x may only be updated as a result of insert and delete operations. All the pointer lists can be initialized by inserting each shared copy in the network at the start of the computation. We do not address the cost of initializing the auxiliary memories of the nodes.

Let r be the node with highest $\beta(r)$ such that there exists i in $[(\log n)/b]$ satisfying: (i) $r[k] = A[k]$ for all k in $[i]$, (ii) $r[i]$ matches $A[i]$ in the largest number of rightmost bits, and (iii) if $i < (\log n)/b - 1$, there is no node y with $y[k] = A[k]$ for all k in $[i + 1]$. We call r the *root* node for object A . The uniqueness of the root node for each A in \mathcal{A} is crucial to guarantee the success of every read operation.

In this section and throughout the paper, we use the notation $\langle \alpha \rangle_k$ to denote the sequence (of length $k + 1$) $\alpha_0, \alpha_1, \dots, \alpha_k$ (of length $k + 1$). When clear from the context, k will be omitted. In particular, a *primary neighbor sequence* for A is a maximal sequence $\langle u \rangle_k$ such that u_0 is in V , u_k is the root node for A , and u_{i+1} is the primary $(i, A[i])$ -neighbor of u_i , for all i . It is worth noting that the sequence $\langle u \rangle$ is such that the label of node u_i satisfies $(u_i[i - 1], \dots, u_i[0]) = (A[i - 1], \dots, A[0])$, for all i . We now give an overview of the read, insert, and delete operations.

Read. Consider a node x attempting to read an object A . The read operation proceeds by successively forwarding the *read request* for object A originating at node x along the primary neighbor sequence $\langle x \rangle$ for A with $x_0 = x$. When forwarding the read request, node x_{i-1} also informs x_i of the current best upper bound k on the cost of sending a copy of A to x . On receiving the read request with associated upper bound k , node x_i proceeds as follows. If x_i is the root node for A , then x_i requests that the copy of A associated with k be sent to x . Otherwise, x_i communicates with its primary and secondary $(i, A[i])$ -neighbors to check whether the pointer list of any of these neighbors has an entry (A, z, k_1) such that k_1 is at most k . Then, x_i updates k to be minimum of k and the smallest value of k_1 thus obtained (if any). If k is within a constant factor of the cost of following $\langle x \rangle$ up to x_i , that is, k is $O(\sum_{j=0}^{i-1} c(x_j, x_{j+1}))$, then x_i requests that the copy of A associated with the upper bound k be sent to x . Otherwise, x_i forwards the read request to x_{i+1} .

Insert. An *insert request* for object A generated by node y updates the pointer lists of some nodes

that form a prefix subsequence of the primary neighbor sequence $\langle y \rangle$ for A with $y_0 = y$. When such an update arrives at a node y_i by means of an insert message, y_i updates its pointer list if the upper bound $\sum_{j=0}^{i-1} c(y_j, y_{j+1})$ on the cost of getting object A from y , is smaller than the current upper bound associated with A in this list. In other words, y_i updates $Ptr(y_i)$ if (A, \cdot, \cdot) is not in this list, or if (A, \cdot, k) is in $Ptr(y_i)$ and k is greater than $\sum_{j=0}^{i-1} c(y_j, y_{j+1})$. Node y_i forwards the insert request to node y_{i+1} only if $Ptr(y_i)$ is updated.

Delete. A *delete request* for object A generated by node y eventually removes all triples of the form (A, y, \cdot) from the pointer lists $Ptr(y_i)$, where $\langle y \rangle$ is the primary neighbor sequence for A with $y_0 = y$, making the copy of A at y unavailable to other nodes in the network. Upon receiving such a request by means of a delete message, node y_i checks whether the entry associated with A in its pointer list is of the form (A, y, \cdot) . In case it is not, the delete procedure is completed and we need to proceed no further in updating the pointer lists in $\langle y \rangle$. Otherwise, y_i deletes this entry from its pointer list, and checks for entries associated with A in the pointer lists of its reverse $(i-1, A[i-1])$ -neighbors. If an entry is found, y_i updates $Ptr(y_i)$ by adding the entry $(A, w, k + c(w, y_i))$, where w is the reverse $(i-1, A[i-1])$ -neighbor of y_i with minimum upper bound k associated with A in its pointer list. A delete message is then forwarded to y_{i+1} .

The read, insert, and delete procedures for an object A are formally described in Figure 1. The messages and requests in the figure are all with respect to object A . A *read request* is generated by node x ($= x_0$) sends a message $Read(x, \infty, \cdot)$ to itself, if x does not hold a copy of A . A read message $Read(x, k, y)$ indicates a read request for object A generated at node x , and that the current best upper bound on the cost of reading A is k and such a copy resides at y . An *insert* (resp., *delete*) *request* is generated when node y ($= y_0$) sends a message $Insert(y, 0)$ (resp., $Delete(y)$) to itself. An insert message $Insert(y, k)$ indicates to its recipient node z that the best known upper bound on the cost incurred by z to read the copy of A located at y is k . We assume that y holds a copy of A and that this copy is unshared (resp., shared) when an insert (resp., delete) request for A is generated at y .

The correctness of our access scheme follows from the two points below:

- (1) The insert and delete procedures maintain the following invariants. For any A in \mathcal{A} and any y in V , there is at most one entry associated with A in the pointer list of y . If y holds a shared copy of A and $\langle y \rangle$ is the primary neighbor sequence for A with $y_0 = y$, then: (i) there is an entry associated with A in the pointer list of every node in $\langle y \rangle$, and (ii) the nodes that have a pointer list entry associated with the copy of A at y form a prefix subsequence of $\langle y \rangle$. The preceding claims follow directly from the insert and delete procedures as described.
- (2) Every read request for any object A by any node x is successful. That is, it locates and brings to x a shared copy of A , if such a copy is currently available. The read operation proceeds by following the primary neighbor sequence $\langle x \rangle$ for A with $x_0 = x$, until either a copy of A is located or the root for A is reached. By point (1), there exists a shared copy of A in the network if and only if the root for A has a pointer to it.

4 Results

In this section, we formally state the main results of our access scheme. In Theorems 1, 2, 3, and 4, we prove bounds on the cost of a read, the cost of an insert or delete, the size of auxiliary memory, and the adaptability of our access scheme. Let C denote $\max\{c(u, v) : u, v \in V\}$.

Theorem 1 *Let x be any node in V and let A be any object in \mathcal{A} . If y is the nearest node to x that holds a shared copy of A , then the expected cost of a read operation is $O(f(\ell(A))c(x, y))$.*

When a node x tries to read an object A which has currently no shared copy in the network, then the expected cost of the associated operation is $O(C)$.

<p><i>Action of x_i on receiving a message $Read(x, k, y)$:</i></p> <p>If $i > 0$ and $x_i[i-1] \neq A[i-1]$, or $i = (\log n)/b - 1$ (that is, x_i is the root for A) then:</p> <ul style="list-style-type: none"> • Node x_i sends a message <i>Satisfy</i>(x) to node v such that (A, v, \cdot) is in $Ptr(x_i)$, requesting it to send a copy of A to x. If $Ptr(x_i)$ has no such entry, then there are no shared copies of A. <p>Otherwise:</p> <ul style="list-style-type: none"> • Let U be the set of secondary $(i, A[i])$-neighbors of x_i. Node x_i requests a copy of A with associated upper bound at most k from each node in $U \cup \{x_{i+1}\}$. • Each node u in $U \cup \{x_{i+1}\}$ responds to the request message received from x_i as follows: if there exists an entry (A, v, q_v) in $Ptr(u)$ and if $q'_v = q_v + c(x_i, u) + \sum_{j=0}^{i-1} c(x_j, x_{j+1})$ is at most k, then u sends a success message <i>Success</i>(v, q'_v) to x_i. • Let U' be the set of nodes u from which x_i receives a response message <i>Success</i>(u, k_u). If U' is not empty, then x_i updates (k, y) to be (k_z, z), where z is a node with minimum k_u over all u in U'. • If $k = O(\sum_{j=0}^{i-1} c(x_j, x_{j+1}))$ then x_i sends a message <i>Satisfy</i>(x) to node y, requesting y to send a copy of A to x. Otherwise, x_i forwards a message <i>Read</i>(x, k, y) to x_{i+1}. 	<p><i>Action of y_i on receiving a message $Delete(y)$:</i></p> <p>If (A, y, \cdot) is in $Ptr(y_i)$, then:</p> <ul style="list-style-type: none"> • Let U be the set of reverse $(i-1, A[i-1])$-neighbors of y_i. Node y_i removes (A, y, \cdot) from $Ptr(y_i)$, and requests a copy of A from each u in U. • Each u in U responds to the request message from y_i by sending a message <i>Success</i>($v, q_v + c(y_i, u)$) to y_i iff (A, v, q_v) is in $Ptr(u)$. • Let U' be the set of nodes u such that y_i receives a message <i>Success</i>(u, k_u) in response to the request message it sent. If $U' > 0$ then y_i inserts (A, w, k_w) into $Ptr(y_i)$, where w is the node in U' such that $k_w \leq k_u$, for all u in U'. • If $y_i[i-1] = A[i-1]$ then y_i sends a message <i>Delete</i>(y) to y_{i+1}.
<p><i>Action of y_i on receiving a message $Insert(y, k)$:</i></p> <p>If (A, \cdot, \cdot) is not in $Ptr(y_i)$, or (A, \cdot, k') is in $Ptr(y_i)$ and $k' > k$, then:</p> <ul style="list-style-type: none"> • Node y_i accordingly creates or replaces the entry associated with A in $Ptr(y_i)$ by inserting (A, y, k) into this list. • If $y_i[i-1] = A[i-1]$ then y_i sends a message <i>Insert</i>($y, k + c(y_i, y_{i+1})$) to y_{i+1}. 	

Figure 1: Actions on receiving messages *Read*, *Insert*, and *Delete* for object A .

Theorem 2 *The expected cost of an insert operation is $O(C)$, and that of a delete operation is $O(C \log n)$.*

Theorem 3 *Let q be the number of objects that can be stored in the main memory of each node. The size of the auxiliary memory at each node is $O(q \log^2 n)$ words whp.*

Theorem 4 *The adaptability of our scheme is $O(\log n)$ expected and $O(\log^2 n)$ whp.*

5 Analysis

In this section, we analyze the access scheme described in Section 3, and establish the main results described in Section 4. Section 5.1 presents some useful properties of balls. Section 5.2 presents properties of primary and secondary neighbors. Section 5.3 presents the proofs of Theorems 1 and 2. Sections 5.4 and 5.5 present the proofs of Theorems 3 and 4, respectively.

Several constants appear in the model, the algorithms, and the analysis: δ and Δ appear in the model, b and d appear in the algorithms, γ and ε appear in the analysis. We choose b , d , γ , and ε such that the following inequalities hold:

$$\gamma \geq \max\{\Delta^2, 4\Delta\} \quad (2)$$

$$2^b \geq \max\{\Delta^2 \gamma^3, 4\gamma^2 \Delta, 2\Delta(\gamma + 1)\} \quad (3)$$

$$\varepsilon \geq \max\{6\Delta/\gamma, 4e^{-\gamma/4\Delta}, 6(d+1)/2^b, 6(2e/2^b)^{d/2}, 6(e\Delta\gamma^2/d)^d\}, \quad (4)$$

$$\varepsilon < (10 \cdot 2^{b \log_\delta 2})^{-1} \quad (5)$$

One assignment of values to the constants γ , b , d , and ε that satisfies the above inequalities is the following. Set γ equal to $2^{b/3}/\Delta^{2/3}$, d equal to $e2^{2b/3+1}/\Delta^{1/3}$ and ε equal to $6e\Delta^{5/3}/2^{b/3}$. The preceding assignment satisfies Equations 3 and 4 if b is set sufficiently large. Equations 2 and 5 can be satisfied by setting b large enough so that $2^b \geq \Delta^8$ and $2^{b(1/3 - \log_\delta 2)} > 60e\Delta^{5/3}$.

5.1 Properties of Balls

Given any u in V and any integer k in $[1, n]$, let $N(u, k)$ denote the unique set of k nodes such that for any v in $N(u, k)$ and w not in $N(u, k)$, $c(u, v)$ is less than $c(u, w)$. (For integers a and b , we let $[a, b]$ denote the set $\{k \in \mathbf{Z} : a \leq k \leq b\}$.) We refer to $N(u, k)$ as the *ball of size k* around u . For convenience, if k is greater than n , we let $N(u, k)$ be V .

Lemma 5.1 *Let u, v , and w be in V and let k_0 and k_1 be positive integers. If v is in $N(u, k_0)$ and w is in $N(v, k_1)$, then w is in $N(u, \Delta k_0 + \Delta^2 k_1)$.*

Proof: Let r_0 and r_1 denote $c(u, v)$ and $c(v, w)$ respectively. The node w is contained in the ball $M(u, c(u, w))$. If r_0 is at least r_1 , then $|M(u, c(u, w))|$ is at most $|M(u, r_0 + r_1)|$, which by Equation 1, is at most Δk_0 . Otherwise, $|M(u, r_0 + r_1)|$ is at most $|M(v, 4r_1)|$, which by Equation 1, is at most $\Delta^2 k_1$. Therefore, w belongs to $N(u, \Delta k_0 + \Delta^2 k_1)$. ■

Given any subset S of V and some node u in S , let $q(u, S)$ (resp., $r(u, S)$) denote the largest (resp., smallest) integer k such that $N(u, k)$ is a subset (resp., superset) of S . Let $Q(u, S)$ and $R(u, S)$ denote $N(u, q(u, S))$ and $N(u, r(u, S))$, respectively.

Lemma 5.2 *Let u be in V , let S be a subset of V , and let k be in $[1, n]$. Then $N(u, k)$ is a subset (resp., superset) of S iff $N(u, k)$ is a subset of $Q(u, S)$ (resp., superset of $R(u, S)$).*

Proof: If $N(u, k)$ is a subset of S then $q(u, S)$ is at least k ; hence, $N(u, k)$ is a subset of $Q(u, S)$. If $N(u, k)$ is a subset of $Q(u, S)$ then $N(u, k)$ is a subset of S because $Q(u, S)$ is a subset of S . If $N(u, k)$ is a superset of S then $r(u, S)$ is at most k ; hence, $N(u, k)$ is a superset of $R(u, S)$. If $N(u, k)$ is a superset of $R(u, S)$ then $N(u, k)$ is a superset of S because $R(u, S)$ is a superset of S . ■

Lemma 5.3 *Let u belong to V , and let k_0 and k_1 denote positive integers such that $k_1 \geq \Delta^2 k_0$. For any v in $N(u, k_0)$, $q(v, N(u, k_1))$ is at least k_1/Δ and $R(v, N(u, k_1))$ is a subset of $N(u, \Delta k_1)$.*

Proof: Let r_0 (resp., r_1) denote the radius of $N(u, k_0)$ (resp., $N(u, k_1)$). Since $k_1 \geq \Delta^2 k_0$, by Equation 1, we obtain $r_1 - r_0 \geq (r_1 + r_0)/2$.

Let w be the node in $Q(v, N(u, k_1))$ such that $c(v, w)$ is maximum. By the definitions of N and Q , we have: $r_1 - r_0 \leq c(v, w)$. It follows that $2c(v, w)$ is at least $r_1 + r_0$ and $M(v, 2c(v, w))$ is a superset of $N(u, k_1)$. We now obtain a lower bound on $q(v, N(u, k_1))$ as follows:

$$\begin{aligned} q(v, N(u, k_1)) &= |M(v, c(v, w))| \\ &\geq |M(v, 2c(v, w))|/\Delta \\ &\geq k_1/\Delta. \end{aligned}$$

We now place an upper bound on $r(v, N(u, k_1))$. Let w be the node in $R(v, N(u, k_1))$ such that $c(v, w)$ is maximum. By the definitions of N and R , we have: $r_1 - r_0 \leq c(v, w) \leq r_1 + r_0$. It follows that $2(r_1 - r_0)$ is at least $c(v, w)$ and $M(v, c(v, w))$ is a subset of $M(u, 2r_1)$. Therefore,

$$\begin{aligned} r(v, N(u, k_1)) &= |M(v, c(v, w))| \\ &\leq |M(u, 2r_1)| \\ &\leq \Delta k_1. \end{aligned}$$

We refer to any predicate on V that only depends on the label of v as a *label predicate*. Given any node u in V and a label predicate \mathcal{P} on V , let $p(u, \mathcal{P})$ denote the node v such that: (i) $\mathcal{P}(v)$ holds, and

(ii) for any node w such that $\mathcal{P}(w)$ holds, $c(u, v)$ is at most $c(u, w)$. (We let $p(u, \mathcal{P})$ be null if such a v is not defined.) Let $P(u, \mathcal{P})$ be $M(u, c(u, p(u, \mathcal{P})))$, if $p(u, \mathcal{P})$ is not null, and V , otherwise.

For u in V and i in $[(\log n)/b]$, let $\lambda_{\geq i}(u)$ denote the string of $(\log n - ib)$ bits given by $u[(\log n)/b - 1] \cdots u[i + 1]u[i]$. For convenience, we let $\lambda_{> i}(u)$ denote $\lambda_{\geq i+1}(u)$. For all i and all u in V , let $\mathcal{P}_i(u)$ hold iff $u[i] = A[i]$. For all i and all u in V , let $\mathcal{P}_{< i}(u)$ denote $\bigwedge_{j \in [i]} \mathcal{P}_j(u)$. Let $\mathcal{P}_{\leq i}(u)$, $\mathcal{P}_{> i}(u)$, and $\mathcal{P}_{\geq i}(u)$ be defined similarly. (We note that for u and v in V and nonnegative integers i and j , if $(u \neq v) \vee ((u = v) \wedge (i \neq j))$, then $\mathcal{P}_i(u)$ and $\mathcal{P}_j(v)$ are independent random variables. Also, each of the predicates defined above is a label predicate.)

Lemma 5.4 *Let S and S' be subsets of V and let u belong to S . Let \mathcal{P} be a label predicate on V and for each v in S' , let $\lambda_{\geq 0}(v)$ be chosen independently and uniformly at random.*

1. *Given that $P(u, \mathcal{P}) \subseteq S$, we have: (i) for each node v in $S' \setminus P(u, \mathcal{P})$, $\lambda_{\geq 0}(v)$ is independently and uniformly random, and (ii) for each node v in $P(u, \mathcal{P}) \setminus \{p(u, \mathcal{P})\}$, $\mathcal{P}(v)$ is false.*
2. *Given that $P(u, \mathcal{P}) \not\subseteq S$, we have: (i) for each node v in $S' \setminus Q(u, S)$, $\lambda_{\geq 0}(v)$ is independently and uniformly random, and (ii) for each node v in $Q(u, S)$, $\mathcal{P}(v)$ is false.*
3. *Given that $P(u, \mathcal{P}) \supseteq S$, we have: (i) for each node v in $S' \setminus R(u, S)$, $\lambda_{\geq 0}(v)$ is independently and uniformly random, and (ii) for each node v in $R(u, S) \setminus \{r(u, S)\}$, $\mathcal{P}(v)$ is false.*

Proof: We first consider Part 1 of the lemma. Part 1(i) follows from the independence of $\mathcal{P}(v)$ and $\mathcal{P}(w)$ for any two distinct nodes. By the definition of P , $\mathcal{P}(p(u, \mathcal{P}))$ holds and for each node v in $P(u, \mathcal{P})$, $\mathcal{P}(v)$ is false. This proves Part 1(ii). Parts 2 and 3 follow similarly. For Part 2, we note that the event $P(u, \mathcal{P}) \not\subseteq S$ is equivalent to the event that for each node v in $Q(u, S)$, $\mathcal{P}(v)$ is false. For Part 3, we note that the event $P(u, \mathcal{P}) \supseteq S$ is equivalent to the event that for each node in $R(u, S) \setminus \{r(u, S)\}$, $\mathcal{P}(v)$ is false. ■

The following claim follows from repeated application of Part 1 of Lemma 5.4.

Corollary 5.4.1 *Let S be an arbitrary subset of V , let i be in $[(\log n)/b - 1]$, and let S' be a subset of V such that $\lambda_{\geq 0}(u)$ is independently and uniformly random for each u in S' . Given a sequence of nodes u_0, u_1, \dots, u_i such that for all j in $[i]$, $u_{j+1} = p(u_j, \mathcal{P}_{\leq j})$ and $P(u_j, \mathcal{P}_{\leq j}) \subseteq S$, we have:*

1. *For each node u in $S' \setminus \bigcup_{j \in [i]} P(u, \mathcal{P}_{\leq j})$, $\lambda_{\geq 0}(u)$ is independently and uniformly random.*
2. *The random variable $\lambda_{> i}(u_i)$ is independently and uniformly random and for each node u in $\bigcup_{j \in [i]} P(u_j, \mathcal{P}_{\leq j}) \setminus \{u_i\}$, $\mathcal{P}_{\leq i}(u)$ is false;* ■

5.2 Properties of Neighbors

In this section, we establish certain claims concerning the different types of neighbors that are defined in Section 3. We differentiate between root and nonroot primary (i, j) -neighbors. A root primary (i, j) -neighbor w of v is a primary (i, j) -neighbor w of v such that $w[i] \neq j$ or $i = (\log n)/b - 1$. A primary neighbor that is not a root primary neighbor is a nonroot primary neighbor.

Lemma 5.5 *Let u and v be in V , and let k denote $|M(u, c(u, v))|$. For any j in $[2^b]$, we have: (i) for any i in $[(\log n)/b - 1]$, the probability that u is a primary (i, j) -neighbor of v is at most $e^{-((k/\Delta)-2)/2^{(i+1)b}}$, and (ii) for any i in $[(\log n)/b]$, the probability that u is a root primary (i, j) -neighbor of v is at most $e^{-n/2^{(i+1)b}}$.*

Proof: Consider the ball $M(v, c(v, u))$. By Equation 1, $|M(v, c(v, u))|$ is at least $|M(v, 2c(v, u))|/\Delta$. Since $M(v, 2c(v, u))$ is a superset of $N(u, k)$, we obtain that $|M(v, c(v, u))|$ is at least k/Δ . Since i is less than $(\log n)/b - 1$, the probability that u is a primary (i, j) neighbor of v is at most:

$$\begin{aligned} & (1 - 1/2^{(i+1)b})^{(k/\Delta)-2} \\ & \leq e^{-((k/\Delta)-2)/2^{(i+1)b}}. \end{aligned}$$

If u is a root primary (i, j) -neighbor of v , then $u[\ell]$ equals $v[\ell]$ for each ℓ in $[i]$ and there does not exist any node w in V such that $w[i]$ equals j and $w[\ell]$ equals $v[\ell]$ for each ℓ in $[i]$. Therefore, the probability that u is a root primary (i, j) -neighbor of v is at most:

$$\begin{aligned} & (1/2^{ib})(1 - 1/2^{(i+1)b})^{n-1}(1 - 1/2^b) \\ & \leq (1/2^{ib})(1 - 1/2^{(i+1)b})^n \\ & \leq (1/2^{ib})e^{-n/2^{(i+1)b}}. \end{aligned}$$

■

Corollary 5.5.1 *Let u and v be in V , let i be in $[(\log n)/b]$, and let j be in $[2^b]$. If u is a primary (i, j) -neighbor of v , then v is in $N(u, O(2^{ib} \log n))$ whp.*

■

For any u in V , let a_u denote the total number of triples (i, j, v) such that i belongs to $[(\log n)/b]$, j belongs to $[2^b]$, v belongs to V , and u is a primary or secondary (i, j) -neighbor of v . Lemma 5.6 is used in the proof of Theorem 4, while Corollary 5.6.1 is used in the proofs of Theorems 2 and 3.

Lemma 5.6 *Let u be in V and let i be in $[(\log n)/b]$. Then, the number of nodes of which u is an i th level primary neighbor is $O(\log n)$ whp. Also, $E[a_u] = O(\log n)$ and a_u is $O(\log^2 n)$ whp.*

Proof: Given a node v in V , i in $[(\log n)/b - 1]$, and j in $[2^b]$, it follows from Lemma 5.5 that the probability that u is a root primary (i, j) -neighbor of v is at most $(1/2^{ib})e^{-n/2^{(i+1)b}}$. Given a node v in V and j in $[2^b]$, the probability that u is a root $((\log n)/b, j)$ -primary neighbor of v is at most $1/n$.

Fix j in $[2^b]$. Let ℓ equal $(\log n - \log \log n)/b - \Omega(1)$, where the constant in the $\Omega(1)$ term is chosen appropriately. We consider two cases: i is less than ℓ , and otherwise. If i is at most ℓ , then the probability that there exists v in V such that u is a root primary (i, j) -neighbor of v is at most:

$$\begin{aligned} & n(1/2^{ib})e^{-n/2^{(i+1)b}} \\ & \leq ne^{-\Omega(\log n)} \\ & = O(1/\text{poly}(n)). \end{aligned}$$

If i is at least ℓ , then given v in V , the probability that u is a root primary (i, j) -neighbor of v is at most $1/2^{\ell b} = O((\log n)/n)$. It follows from Chernoff bounds [5] that the number of nodes for which u is a root primary (i, j) -neighbor is $O(\log n)$ whp.

We now consider secondary and nonroot primary neighbors. For any i in $[(\log n)/b]$, u is a secondary or nonroot primary (i, j) -neighbor of v only if j is $u[i]$ and u is one of the $(d+1)$ nodes w in V with minimum $c(v, w)$ whose lowest ib bits match those of v . We now fix u and i , set j to $u[i]$, and obtain an upper bound on the probability that u is one of the at most $(d+1)$ nodes w with minimum $c(v, w)$ and whose first ib bits match those of v .

Consider a node v in $N(u, \mu^{k+1}2^{(i+1)b}) \setminus N(u, \mu^k 2^{(i+1)b})$, where μ is a real constant that will be specified later. If k equals zero, then the probability that u is a primary or secondary (i, j) -neighbor of v is at most

$1/2^{ib}$. Otherwise, consider the ball $M(v, c(v, u))$. By the low-expansion condition, $|M(v, c(v, u))|$ is at least $|M(v, 2c(v, u))|/\Delta$. We are given that $M(u, c(u, v))$ is a superset of $N(u, \mu^k 2^{(i+1)b})$. Since $M(v, 2c(v, u))$ is a superset of $M(u, c(u, v))$, we obtain that $|M(v, c(v, u))|$ is at least $\mu^k 2^{(i+1)b}/\Delta$. The probability that u is a primary or secondary (i, j) -neighbor of v is at most:

$$\begin{aligned} & d \binom{\mu^k 2^{(i+1)b}/\Delta}{d} (1 - 1/2^{(i+1)b})^{(\mu^k 2^{(i+1)b}/\Delta) - d} / (2^{ib} 2^{(i+1)bd}) \\ & \leq d (e \mu^k 2^{(i+1)b} / (\Delta d))^d e^{-\mu^k/\Delta} (1 - 1/2^{(i+1)b})^{-d} / (2^{ib} 2^{(i+1)bd}) \\ & \leq 4d (e \mu^k / (\Delta d))^d (e^{-\mu^k/\Delta} / 2^{ib}) \\ & \leq 1 / ((2\mu)^k 2^{ib}). \end{aligned}$$

(The third inequality holds since $d \leq 2^b \leq 2^{ib}$ and $(1 - 1/2^{ib})^{-2^{ib}}$ is at most 4. The fourth inequality follows by choosing μ large enough with respect to Δ and d such that $e^{\mu^k/\Delta} \geq (2^k \Delta^k / d^{d-1}) (\mu^k / \Delta)^{d+1}$ for all $k \geq 1$.)

Thus, the expected number of nodes for which u is a secondary or nonroot primary neighbor is at most:

$$\begin{aligned} & \sum_{i \in [(\log n)/b], j=u[i]} \sum_{k \geq 0} \sum_{v \in N(u, \mu^{k+1} 2^{(i+1)b}) \setminus N(u, \mu^k 2^{(i+1)b})} 1 / ((2\mu)^k 2^{ib}) \\ & \leq \sum_{i \in [(\log n)/b], j=u[i]} 2^b \mu \\ & = O(\log n). \end{aligned}$$

To obtain a high probability bound on the number of nodes for which u is a secondary or nonroot primary neighbor, we proceed as follows. For any v not in $N(u, \Theta(2^{(i+1)b} \log n))$, it follows from Lemma 5.5 that the probability that u is a secondary or nonroot primary (i, j) -neighbor of v is $O(1/\text{poly}(n))$. For any v in $N(u, \Theta(2^{(i+1)b} \log n))$, the probability that u is a secondary or nonroot primary (i, j) -neighbor of v is at most $1/2^{(i+1)b}$. Therefore, the number of nodes for which u is a secondary or nonroot primary neighbor is $O(\log^2 n)$ whp.

The bounds on expectation and the high probability bounds together establish that $E[a_u]$ is $O(\log n)$ and a_u is $O(\log^2 n)$ whp. \blacksquare

Corollary 5.6.1 *For any u in V , the total number of reverse neighbors of u is $O(\log^2 n)$ whp, and expected $O(\log n)$.*

Proof: The desired claim follows directly from Lemma 5.6 since v is a reverse (i, j) -neighbor of u only if u is a primary (i, j) -neighbor of v . \blacksquare

For any u and v in V and i in $[(\log n)/b]$, v is said to be an i -leaf of u if there exists a sequence $v = v_0, v_1, \dots, v_{i-1}, v_i = u$, such that for all j in $[i]$, v_{j+1} is a primary $(j, v_{j+1}[j])$ -neighbor of v_j . Lemma 5.7 is used in the proof of Theorem 3.

Lemma 5.7 *Let u belong to V , and let i be in $[(\log n)/b]$. Then the number of i -leaves of u is $O(2^{ib} \log n)$ whp.*

Proof: We will establish the lemma by showing that if v is an i -leaf of u , then v is in $N(u, c_0 2^{ib} \log n)$ whp, where c_0 is a real constant that is specified later. By Corollary 5.5.1, we have that for all j in $[i]$, v_j is in $N(v_{j+1}, c_1 2^{(j+1)b} \log n)$ whp for some real constant c_1 . We will prove by induction on j in $[i+1]$ that $v = v_0$ is in $N(v_j, c_0 2^{jb} \log n)$ whp.

The induction base follows trivially. For the induction step, let us assume that v belongs to $N(v_j, c_0 2^{jb} \log n)$. By Corollary 5.5.1, v_j belongs to $N(v_{j+1}, c_1 2^{jb} \log n)$ whp. Applying Lemma 5.1 with the substitution (v_{j+1}, v_j, v) for (u, v, w) , we obtain that v is in $N(v_{j+1}, (\Delta c_1 + \Delta^2 c_0) 2^{jb} \log n)$. Since $\Delta^2 \leq 2^b$, we can choose c_0 large enough such that $c_0(2^b - \Delta^2)$ is at least Δc_1 . It thus follows that v is in $N(v_{j+1}, c_0 2^{(j+1)b})$.

Applying the above inductive claim with $j = i$, we obtain that v is in $N(u, O(2^{ib} \log n))$ whp. The desired claim follows. \blacksquare

5.3 Cost of operations

Consider a read request originating at node x for an object A . Let y denote a node that has a copy of A . In the following, we show that the expected cost of a read operation is $O(f(\ell(A))c(x, y))$. Letting y to be the node with minimum $c(x, y)$ among the set of nodes that have a copy of A , this bound implies that the expected cost is asymptotically optimal.

Let $\langle x \rangle$ and $\langle y \rangle$ be the primary neighbor sequences for A with $x_0 = x$ and $y_0 = y$, respectively. For any nonnegative integer i , let A_i (resp., D_i) denote the ball of smallest radius around x_i (resp., y_i) that contains x_{i+1} (resp., y_{i+1}). Let B_i (resp., E_i) denote the set $\cup_{0 \leq j \leq i} A_j$ (resp., $\cup_{0 \leq j \leq i} D_j$). Let C_i denote the ball of smallest radius around x_i that contains all of the secondary $(i, A[i])$ -neighbors of x_i . For convenience, we define $B_{-1} = E_{-1} = \emptyset$.

It is useful to consider an alternative view of x_i, y_i, A_i , and D_i . For any nonnegative i , if x_{i+1} (resp., y_{i+1}) is not the root node for A , then x_{i+1} (resp., y_{i+1}) is $p(x_i, \mathcal{P}_{\leq i})$ (resp., $p(y_i, \mathcal{P}_{\leq i})$) and A_i (resp., D_i) is $P(x_i, \mathcal{P}_{\leq i})$ (resp., $P(y_i, \mathcal{P}_{\leq i})$).

Let γ be an integer constant that is chosen later appropriately. For any nonnegative integer i and any integer j , let X_i^j (resp., Y_i^j) denote the ball $N(x, \gamma^j 2^{(i+1)b})$ (resp., $N(y, \gamma^j 2^{(i+1)b})$). Let i^* denote the least integer such that the radius of $X_{i^*}^1$ is at least $c(x, y)$. Let a_i (resp., b_i) denote the radius of X_i^1 (resp., Y_i^1).

Lemma 5.8 *For all i such that $i \geq i^*$, X_i^2 is a superset of Y_i^1 .*

Proof: By the definition of i^* , a_i is at least $c(x, y)$. Therefore, $M(y, 2a_i)$ is a superset of X_i^1 . Hence, $M(y, 2a_i)$ contains at least $\gamma 2^{(i+1)b}$ nodes and is a superset of Y_i^1 .

By Equation 1, $|M(x, 3a_i)|$ is at most $\Delta^2 |M(x, a_i)| \leq \Delta^2 \gamma 2^{(i+1)b} \leq \gamma^2 2^{(i+1)b}$. Thus, $M(x, 3a_i)$ is a subset of X_i^2 . Since $M(x, 3a_i)$ is a superset of $M(y, 2a_i)$, which is a superset of Y_i^1 , the claim holds. \blacksquare

Lemma 5.9 *For all i in $[(\log n)/b - 2]$, we have $2^{b \log_\Delta 2} a_i \leq a_{i+1} \leq 2^{b \log_\delta 2} a_i$ and $2^{b \log_\Delta 2} b_i \leq b_{i+1} \leq 2^{b \log_\delta 2} b_i$. For $i = (\log n)/b - 2$, we have $a_{i+1} \leq 2^{b \log_\delta 2} a_i$ and $b_{i+1} \leq 2^{b \log_\delta 2} b_i$. Also, a_{i^*} and b_{i^*} are both $O(c(x, y))$.*

Proof: Since $\gamma \leq 2^b$, for all i in $[(\log n)/b - 2]$, we have $|X_{i+1}^1| = 2^b |X_i^1|$ (resp., $|Y_{i+1}^1| = 2^b |Y_i^1|$). Therefore, for all i in $[(\log n)/b - 2]$, it follows from Equation 1 that $2^{b \log_\Delta 2} a_i \leq a_{i+1} \leq 2^{b \log_\delta 2} a_i$ (resp., $2^{b \log_\Delta 2} b_i \leq b_{i+1} \leq 2^{b \log_\delta 2} b_i$). For $i = (\log n)/b - 2$, $|X_{i+1}^1| \leq 2^b |X_i^1|$ (resp., $|Y_{i+1}^1| \leq 2^b |Y_i^1|$), and hence, $a_{i+1} \leq 2^{b \log_\delta 2} a_i$ (resp., $b_{i+1} \leq 2^{b \log_\delta 2} b_i$).

If $i^* > 0$, then a_{i^*} (resp., b_{i^*}) is at most $2^{b \log_\delta 2} c(x, y)$. Otherwise, a_{i^*} (resp., b_{i^*}) is $O(2^{\log_\delta \gamma}) = O(c(x, y))$, since δ and γ are constants. \blacksquare

We define two sequences $\langle s_i \rangle$ and $\langle t_i \rangle$ of nonnegative integers as follows:

$$s_i = \begin{cases} 0 & \text{if } B_i \subseteq X_i^1, A_i \supseteq X_i^{-1}, C_i \supseteq X_i^2, \\ 1 & \text{if } B_i \subseteq X_i^1, A_i \supseteq X_i^{-1}, C_i \not\supseteq X_i^2, \\ 2 & \text{if } B_i \subseteq X_i^1, A_i \not\supseteq X_i^{-1}, \text{ and} \\ 3 + j & \text{if } 0 \leq j \leq i, B_{i-j} \not\subseteq X_i^1, B_{i-j-1} \subseteq X_i^1. \end{cases}$$

$$t_i = \begin{cases} 0 & \text{if } E_i \subseteq Y_i^1, \text{ and} \\ 1+j & \text{if } 0 \leq j \leq i, E_{i-j} \not\subseteq Y_i^1, E_{i-j-1} \subseteq Y_i^1. \end{cases}$$

Lemma 5.10 *If s_i is in $\{0, 1, 2\}$, then $c(x_i, x_{i+1})$ is $O(a_i)$. If t_i is 0, then $c(y_i, y_{i+1})$ is $O(b_i)$.*

Proof: The proof of the first claim follows from the observation that if s_i is in $\{0, 1, 2\}$ then $A_i \subseteq B_i \subseteq X_i^1$. The proof of the second claim follows from the observation that if t_i is 0 then $D_i \subseteq E_i \subseteq Y_i^1$. ■

We now determine an upper bound on the cost of read for A as follows. Let τ be the smallest integer $i \geq i^*$ such that $(s_i, t_i) = (0, 0)$. By Lemma 5.8, C_τ is a superset of D_τ , implying that a copy of A is located within τ forwarding steps along $\langle x \rangle$. By the definition of the primary and secondary neighbors, the cost of any request (resp., forward) message sent by node x_i is at most $d \cdot c(x_i, x_{i+1})$ (resp., $c(x_i, x_{i+1})$). Since a copy of A is located within τ forwarding steps, by the definition of the algorithm, the cost of all messages needed in locating the particular copy of A that is read is at most $O(\sum_{0 \leq j < \tau} (d^2 c(x_j, x_{j+1}) + c(y_j, y_{j+1})))$. The cost of reading the copy is at most $f(\ell(A))$ times the preceding cost. Since d is a constant, the cost of reading A is at most:

$$\sum_{0 \leq j < \tau} O(f(\ell(A))(c(x_j, x_{j+1}) + c(y_j, y_{j+1}))) \quad (6)$$

The remainder of the proof concerns the task of showing that $E[\sum_{0 \leq j < \tau} (c(x_j, x_{j+1}) + c(y_j, y_{j+1}))]$ is $O(c(x, y))$. A key idea is to establish that the sequence $\langle s_i, t_i \rangle$ corresponds to a two-dimensional random walk that is biased towards $(0, 0)$. Lemmas 5.11 and 5.12 below provide the important first step towards formalizing this notion.

Lemma 5.11 *Let i be in $[(\log n)/b - 1]$. Given arbitrary well defined values for s_j and t_j for all j in $[i]$ such that s_{i-1} is at least 3, the probability that s_i is less than s_{i-1} is at least $1 - \varepsilon^2$. Given arbitrary values for s_j and t_j for all j in $[i]$ such that t_{i-1} is at least 1, the probability that t_i is less than t_{i-1} is at least $1 - \varepsilon^2$. ■*

Lemma 5.12 *Let i be in $[(\log n)/b - 1]$. Given arbitrary well defined values for s_j and t_j for all j in $[i]$ such that s_{i-1} is at most 3, the probability that s_i is 0 is at least $1 - \varepsilon$. Given arbitrary values for s_j and t_j for all j in $[i]$ such that t_{i-1} is at most 1, the probability that t_i is 0 is at least $1 - \varepsilon$.*

In order to establish the above lemmas, we introduce some additional notation. For each $i \geq -1$, we define S_i and T_i as follows. Let $S_{-1} = T_{-1} = \emptyset$. For nonnegative i , we have:

$$S_i = \begin{cases} S_{i-1} \cup B_i \cup (C_i \cap R(x_i, X_i^2)) & \text{if } s_i \in \{0, 1\}, \\ S_{i-1} \cup B_i & \text{if } s_i = 2, \\ S_{i-1} \cup B_{i-s_i+2} \cup Q(x_{i-s_i+3}, X_i^1) & \text{otherwise.} \end{cases}$$

$$T_i = \begin{cases} T_{i-1} \cup E_i & \text{if } t_i = 0, \\ T_{i-1} \cup E_{i-t_i} \cup Q(y_{i-t_i+1}, Y_i^1) & \text{otherwise.} \end{cases}$$

Lemmas 5.13, 5.14, and 5.15 are used in the proofs of Lemmas 5.11 and 5.12.

Lemma 5.13 *Let i be in $[(\log n)/b - 1]$. Given arbitrary values for s_j and t_j for all j in $[i]$, we have:*

1. *For each node u not in $S_i \cup T_i$, each bit of $\lambda_{\geq 0}(u)$ is independently and uniformly drawn from $\{0, 1\}$.*
2. *There exists a subset S'_i of S_i of size at most $d + 1$ such that: (i) for each node u in S'_i , each bit of $\lambda_{> i}(u)$ is independently and uniformly random, and (ii) for each node u in $S_i \setminus S'_i$, $\mathcal{P}_{\leq i}(u)$ is false.*

3. There exists at most one node v in T_i such that: (i) each bit of $\lambda_{>i}(v)$ is independently and uniformly random, and (ii) for each node u in $T_i \setminus \{v\}$, $\mathcal{P}_{<i}(u)$ is false.

Proof: We prove Parts 1, 2, and 3 for all $i \geq -1$. The proof is by induction. For the induction base we set $i = -1$. Part 1 follows directly from the random assignment of labels. For Part 2, we set S'_{-1} to \emptyset , and the desired claim holds since S_{-1} is \emptyset . The claim of Part 3 holds vacuously since T_{-1} is \emptyset .

For the induction hypothesis, we assume that Parts 1, 2, and 3 of the lemma hold for $i - 1$. We first consider different cases depending on the value of s_i .

- (a) $s_i = 3 + j$, $j \in [i]$: The event $s_i = 3 + j$ is equivalent to the event $(B_{i-j-1} \subseteq X_i^1) \wedge (A_{i-j} \not\subseteq X_i^1)$. We first condition on the event $B_{i-j-1} \subseteq X_i^1$ by invoking Corollary 5.4.1 with the substitution $(X_i^1, V \setminus (S_{i-1} \cup T_{i-1}), i - j)$ for (S, S', i) . We next condition on the event $A_{i-j} \not\subseteq X_i^1$ by invoking Part 2 of Lemma 5.4 with the substitution $(x_{i-j}, X_i^1, V \setminus (S_{i-1} \cup T_{i-1} \cup B_{i-j-1}), \mathcal{P}_{\leq i})$ for (u, S, S', \mathcal{P}) . By combining Part (i) of both invocations, we have: (a.i) for each node v not in $S_{i-1} \cup T_{i-1} \cup B_{i-j-1} \cup Q(x_{i-j}, X_i^1)$, $\lambda_{\geq 0}(v)$ is independently and uniformly random. By combining Part (ii) of both invocations, we have: (a.ii) for each node v in $B_{i-j-1} \cup Q(x_{i-j}, X_i^1)$, $\mathcal{P}_{\leq i}(v)$ is false.

We set S'_i to $S'_{i-1} \setminus (B_{i-j-1} \cup Q(x_{i-j}, X_i^1))$.

- (b) $s_i = 2$: The event $s_i = 2$ is equivalent to the event $(B_i \subseteq X_i^1) \wedge (A_i \not\subseteq X_i^{-1})$. We first condition on the event $B_i \subseteq X_i^1$ by invoking Corollary 5.4.1 with the substitution $(X_i^1, V \setminus (S_{i-1} \cup T_{i-1}), i)$ for (S, S', i) . It follows from the preceding invocation and the definition of B_i that: (b.i) for each node not in $S_{i-1} \cup T_{i-1} \cup B_i$, $\lambda_{\geq 0}(v)$ is independently and uniformly random, and (b.ii) for each node v in $B_i \setminus \{x_{i+1}\}$, $\mathcal{P}_{\leq i}(v)$ is false.

We set S'_i to $S'_{i-1} \setminus (B_i \setminus \{x_{i+1}\})$.

- (c) $s_i \in \{0, 1\}$: The event $s_i \in \{0, 1\}$ is equivalent to the event $(B_i \subseteq X_i^1) \wedge (A_i \supseteq X_i^{-1})$. We condition on the event $B_i \subseteq X_i^1$ by invoking Corollary 5.4 with the substitution $(X_i^1, V \setminus (S_{i-1} \cup T_{i-1}), i)$ for (S, S', i) . It follows from the preceding invocation and the definition of B_i that: (i) for each node v not in $S_{i-1} \cup T_{i-1} \cup B_i$, $\lambda_{\geq 0}(v)$ is independently and uniformly random, and (ii) for each node v in $B_i \setminus \{x_{i+1}\}$, $\mathcal{P}_{\leq i}(v)$ is false.

Let S'_i equal the set $\{v \in C_i \cap R(x_i, X_i^2) : \mathcal{P}_{\leq i}(v)\}$. By the definition of C_i , $|S'_i|$ is at most $d + 1$. If $C_i \not\subseteq X_i^2$, then $C_i \subseteq R(x_i, X_i^2)$ and it follows from the definition of C_i that: (c.i) for each node v not in $S_{i-1} \cup T_{i-1} \cup B_i \cup C_i$, $\lambda_{\geq 0}(v)$ is independently and uniformly random, and (c.ii) for each node v in S'_i , $\lambda_{>i}(v)$ is independently and uniformly random, and for each node v in $(B_i \cup C_i) \setminus S'_i$, $\mathcal{P}_{\leq i}(v)$ is false. If $C_i \supseteq X_i^2$ then $C_i \supseteq R(x_i, X_i^2)$ and it follows from Part 3 of Lemma 5.4 that: (c.i) for each node v not in $S_{i-1} \cup T_{i-1} \cup B_i \cup R(x_i, X_i^2)$, $\lambda_{\geq 0}(v)$ is independently and uniformly random, (c.ii) for each node v in S'_i , $\lambda_{>i}(v)$ is independently and uniformly random, and for each node v in $(B_i \cup R(x_i, X_i^2)) \setminus S'_i$, $\mathcal{P}_{\leq i}(v)$ is false.

We thus obtain from (a.i), (b.i), and (c.i) and the definition of S_i that: (i) for each node u not in $S_i \cup T_{i-1}$, $\lambda_{\geq 0}(u)$ is independently and uniformly random. We obtain from (a.ii), (b.ii), and (c.ii) and the definition of S'_i that: (ii) for each node u in S'_i , $\lambda_{>i}(u)$ is independently and uniformly random, and for each node in $S_i \setminus S'_i$, $\mathcal{P}_{\leq i}(u)$ is false. We next consider two cases depending on the value of t_i .

- (d) $t_i = 1 + j$, $j \in [i]$: This case is similar to Case (a). The event $t_i = 1 + j$ is equivalent to the event $(E_{i-j-1} \subseteq Y_i^1) \wedge (D_{i-j} \not\subseteq Y_i^1)$. We first condition on the event $E_{i-j-1} \subseteq Y_i^1$ by invoking Corollary 5.4.1 with the substitution $(Y_i^1, V \setminus (S_i \cup T_{i-1}), i - j)$ for (S, S', i) . We next condition on the event $D_{i-j} \not\subseteq Y_i^1$ by invoking Part 2 of Lemma 5.4 with the substitution $(y_{i-j}, Y_i^1, V \setminus (S_i \cup T_{i-1} \cup E_{i-j-1}), \mathcal{P}_{\leq i})$ for (u, S, S', \mathcal{P}) .

By combining Part (i) of both invocations, we have: (d.i) for each node v not in $S_i \cup T_{i-1} \cup E_{i-j-1} \cup Q(y_{i-j}, Y_i^1)$, $\lambda_{\geq 0}(v)$ is independently and uniformly random. By combining Part (ii) of both invocations, we have: (d.ii) for each node v in $E_{i-j-1} \cup Q(y_{i-j}, Y_i^1)$, $\mathcal{P}_{\leq i}(v)$ is false.

- (e) $t_i = 0$: This case is similar to Case (b). The event $t_i = 2$ is equivalent to the event $E_i \subseteq Y_i^1$. We invoke Corollary 5.4.1 with the substitution $(Y_i^1, V \setminus (S_i \cup T_{i-1}), i)$ for (S, S', i) to obtain that: (e.i) for each node not in $S_i \cup T_{i-1} \cup E_i$, $\lambda_{\geq 0}(v)$ is independently and uniformly random, and (e.ii) for each node v in $E_i \setminus \{y_{i+1}\}$, $\mathcal{P}_{\leq i}(v)$ is false.

To complete the induction step, we consider each part of the statement of the lemma separately:

1. By (i), (d.i), and (e.i) and the definition of T_i , it follows that given arbitrary values for s_j and t_j , $j \in [i]$, for each node u not in $S_i \cup T_i$, $\mathcal{P}_{\geq 0}(u)$ is independently and uniformly random.
2. This part follows directly from (ii) above.
3. By (d.ii) and (e.ii), it follows that given arbitrary values for s_j and t_j , $j \in [i]$: (i) $\lambda_{> i}((y_{i+1}))$ is independently and uniformly random, and (ii) for each node u in $T_i \setminus \{y_{i+1}\}$, $\mathcal{P}_{\leq i}(u)$ is false. ■

Lemma 5.14 *Let i be a nonnegative integer. If s_i is in $\{0, 1\}$, $|S_i|$ is at most $|X_i^3|$; otherwise, $|S_i|$ is at most $|X_i^1|$. The size of T_i is at most $|Y_i^1|$.*

Proof: The proof follows from the relevant definitions and the inequality $\Delta^2 \leq \gamma$. ■

Lemma 5.15 *Let i be in $[(\log n/b) - 1]$. Given arbitrary values for s_k and t_k for all k in $[i]$ such that s_{i-1} is $3 + j$ for some j in $[i + 1]$, the probability that B_{i-j-1} is a subset of X_{i-1}^2 is at least $1 - \varepsilon^2/2$. Given arbitrary values for s_k and t_k for all k in $[i]$ such that t_{i-1} is $1 + j$ for some j in $[i + 1]$, the probability that E_{i-j-1} is a subset of Y_{i-1}^2 is at least $1 - \varepsilon^2/2$.*

Proof: Let \mathcal{E} denote the event that the random variables $s_k, t_k, k \in [i]$, take the respective arbitrary values. Let us assume that \mathcal{E} holds. We begin with the proof of the first claim. Since s_i is $3 + j$, B_{i-j-1} is not a subset of X_{i-1}^1 , and B_{i-j-2} is a subset of X_{i-1}^1 .

By Part 1 of Lemma 5.13, it follows that given \mathcal{E} , for each node u not in $S_{i-1} \cup T_{i-1}$, $\lambda_{\geq 0}(u)$ is independently and uniformly random. By Lemma 5.14, $|S_{i-1} \cup T_{i-1}|$ is at most $\gamma 2^{ib+1}$. By Lemma 5.3, since $\gamma \geq \Delta^2$, $q(x_{i-j-1}, X_{i-1}^2)$ is at least $\gamma^2 2^{ib}/\Delta$. Therefore, the probability that A_{i-j-1} is not a subset of $Q(x_{i-j-1}, X_{i-1}^2)$ is at most

$$\begin{aligned} (1 - 1/2^{(i-j)b})^{(\gamma^2/\Delta - 2\gamma)2^{ib}} &\leq e^{-(\gamma^2/\Delta - 2\gamma)2^{jb}} \\ &\leq \varepsilon^2/2. \end{aligned}$$

(The last inequality follows from the inequalities $\gamma \geq 4\Delta$ and $e^{-\gamma/\Delta} \leq \varepsilon^2/2$.)

The proof of the second claim is analogous to the above proof and is obtained by substituting (t, D, E, y, Y) for (s, A, B, x, X) . ■

Proof of Lemma 5.11: Let \mathcal{E} denote the event that the random variables $s_j, t_j, j \in [i]$, take the respective arbitrary values. Let us assume that \mathcal{E} holds. We begin with the proof of the first claim. Let s_{i-1} be $3 + j$ for some j in $[i]$. Thus, B_{i-j-1} is not a subset of X_{i-1}^1 , and B_{i-j-2} is a subset of X_{i-1}^1 . We show that given \mathcal{E} , with probability at least $1 - \varepsilon^2$, B_{i-j} is a subset of X_i^1 .

It follows from Lemma 5.15 that: (a) B_{i-j-1} is a subset of X_{i-1}^2 with probability at least $1 - \varepsilon^2/2$. Let us now assume that \mathcal{E} and the event that B_{i-j-1} is a subset of X_{i-1}^2 hold. We now show: (b) the

probability that B_{i-j} is a subset of X_i^1 is at least $1 - \varepsilon/2$. By Lemma 5.14, $|X_{i-1}^2 \cup T_{i-1}|$ is at most $\gamma(\gamma + 1)2^{ib}$. By Lemma 5.3, since $\gamma \geq \Delta^2$, $q(x_{i-j}, X_i^1)$ is at least $\gamma 2^{(i+1)b}/\Delta$. Therefore, the probability that A_{i-j-1} is not a subset of $Q(x_{i-j-2}, X_{i-1}^2)$ is at most

$$\begin{aligned} (1 - 1/2^{(i-j+1)b})^{(\gamma/\Delta - \gamma^2/2^b - \gamma/2^b)2^{(i+1)b}} &\leq e^{-(\gamma/\Delta - \gamma^2/2^b - \gamma/2^b)2^{2b}} \\ &\leq \varepsilon^2/2. \end{aligned}$$

(The last inequality follows from the inequalities $2\Delta(\gamma + 1) \leq 2^b$ and $e^{-\gamma/2\Delta} \leq \varepsilon^2/2$.)

It follows from (a) and (b) above that with probability at least $(1 - \varepsilon^2)$, s_i is less than s_{i-1} , thus establishing the first claim of the lemma. The proof of the second claim is analogous to the above proof and is obtained by substituting (t, D, E, y, Y) for (s, A, B, x, X) . \blacksquare

Proof of Lemma 5.12: Let \mathcal{E} denote the event that the random variables $s_j, t_j, j \in [i]$, take the respective arbitrary values. Let us assume that \mathcal{E} holds. We begin with the proof of the first claim. If s_{i-1} is in $\{0, 1, 2\}$, B_{i-1} is a subset of X_{i-1}^1 . If s_{i-1} is 3, then by Lemma 5.15, B_{i-1} is a subset of X_{i-1}^2 with probability at least $1 - \varepsilon^2/2$. We now assume that B_{i-1} is a subset of X_{i-1}^2 .

We first show: (a) the probability that B_i is a subset of X_i^1 is at least $1 - \varepsilon/3 + \varepsilon^2/2$. By Part 1 of Lemma 5.13, it follows that given \mathcal{E} , for each node u not in $S_{i-1} \cup T_{i-1}$, $\lambda_{\geq 0}(u)$ is independently and uniformly random. By Lemma 5.14, $|S_{i-1} \cup T_{i-1}|$ is at most $\gamma^3 2^{ib+1}$. By Lemma 5.3, since x_i is in X_{i-1}^2 and $2^b \geq \Delta^2 \gamma$, $q(x_i, X_i^1)$ is at least $\gamma 2^{(i+1)b}/\Delta$. Therefore, the probability that A_i is not a subset of $Q(x_i, X_i^1)$ is at most

$$\begin{aligned} (1 - 1/2^{(i+1)b})^{2^{(i+1)b}(\gamma/\Delta - 2\gamma^3/2^b)} &\leq e^{-(\gamma/\Delta - 2\gamma^3/2^b)} \\ &\leq e^{-\gamma/2\Delta} \\ &\leq \varepsilon/3 - \varepsilon^2/2. \end{aligned}$$

(The second inequality follows from the inequality $4\gamma^2\Delta \leq 2^b$ and the last inequality holds since $e^{-\gamma/2\Delta} \leq \varepsilon/4 \leq \varepsilon/3 - \varepsilon^2/2$.)

We next show: (b) the probability that A_i is a superset of X_i^{-1} is at least $1 - \varepsilon/3$. By Lemma 5.3, since $\Delta^2 \gamma^3 \leq 2^b$, $r(x_i, X_i^{-1})$ is at most $\Delta 2^{(i+1)b}/\gamma$. By Lemma 5.13: (i) for each node u not in $S_{i-1} \cup T_{i-1}$, $\lambda_{\geq 0}(u)$ is uniformly random, and (ii) there are at most $d + 1$ nodes in $S_{i-1} \cup T_{i-1}$ for which the predicate $\mathcal{P}_{< i}$ holds. Therefore, the probability that A_i is a subset of $R(x_i, X_i^{-1})$ is at most $d/2^b + \Delta/\gamma$, which is at most $\varepsilon/3$. It follows that with probability at least $1 - \varepsilon/3$, A_i is a superset of X_i^{-1} .

We finally show: (c) given that B_i is a subset of X_i^1 and A_i is a superset of X_i^{-1} , the probability that C_i is a superset of X_i^2 is at least $1 - \varepsilon/3$. Let r_0 (resp., r_1) denote the radius of $R(x_i, X_i^{-1})$ (resp., $R(x_i, X_i^2)$). By definition, $r(x_i, X_i^{-1})$ is at least $2^{(i+1)b}/\gamma$. By Lemma 5.3, $r(x_i, X_i^2)$ is at most $\Delta \gamma^2 2^{(i+1)b}$. By Lemma 5.13: (i) for each node u not in $S_{i-1} \cup T_{i-1}$, $\lambda_{\geq 0}(u)$ is independently and uniformly random, and (ii) there are at most $d + 1$ nodes in $S_{i-1} \cup T_{i-1}$ for which the predicate $\mathcal{P}_{< i}$ holds.

Before calculating the probability that C_i is not a superset of X_i^2 , we need to show that the nodes in $R(x_i, X_i^2)$ are within a cost of $d \cdot c(x_i, x_{i+1})$. We first note that $c(x_i, x_{i+1})$ is at least the difference of the radii of X_i^{-1} and X_{i-1}^2 . Moreover, since $R(x_i, X_i^2)$ is a subset of X_i^3 , the radius of $R(x_i, X_i^2)$ is at most the sum of the radii of X_i^3 and X_{i-1}^2 . Since $(4\gamma^4)^{\log_2 2} \leq \gamma^2 \leq d$, the nodes in $R(x_i, X_i^2)$ are within a cost of $d \cdot c(x_i, x_{i+1})$ from x_i .

It now follows that the probability that C_i is not a superset of X_i^2 is at most

$$\begin{aligned} \binom{d+1}{d/2} (1/2^b)^{d/2} + \binom{\Delta \gamma^2 2^{(i+1)b}}{d} (1/2^{(i+1)b})^d &\leq (2e/2^b)^{d/2} + (e\Delta \gamma^2/d)^d \\ &\leq \varepsilon/3. \end{aligned}$$

(The last inequality follows from the inequalities: $(2e/2^b)^{d/2} \leq \varepsilon/6$ and $(e\Delta \gamma^2/d)^d \leq \varepsilon/6$.)

It follows from (a), (b), and (c) above that with probability at least $1 - \varepsilon$, s_i is 0, thus establishing the first claim of the lemma. The proof of the second claim is analogous to the proof of (a) and is obtained by substituting (t, D, E, y, Y) for (s, A, B, x, X) . \blacksquare

By the definitions of s_i and t_i , it follows that $0 \leq s_{i+1} \leq 3$ if $s_i \leq 2$, and $0 \leq s_{i+1} \leq s_i + 1$ otherwise. In addition, $0 \leq t_{i+1} \leq t_i + 1$, for all i . Let s'_i equal 0 if $s_i = 0$, equal 1 if $s_i \in \{1, 2, 3\}$, and equal $s_i - 2$ otherwise. Hence $0 \leq \max\{s'_{i+1}, t_{i+1}\} \leq \max\{s'_i, t_i\} + 1$, for all i . We now analyze the random walk corresponding to the sequence $\langle \max\{s', t\} \rangle$.

Random Walks. Let $W(U, F)$ be a directed graph in which U is the set of nodes and F is the set of edges. For all u in U , let \mathcal{D}_u be a probability distribution over the set $\{(u, v) \in F\}$ (let $\Pr_{\mathcal{D}_u}[(u, v) : (u, v) \notin F] = 0$, for convenience). A *random walk* on W starting at v_0 and according to $\{\mathcal{D}_u : u \in U\}$ is a random sequence $\langle v \rangle$ such that: (i) v_i is in U and (v_i, v_{i+1}) is in F , for all i , and (ii) given any fixed (not necessarily simple) path u_0, \dots, u_i in W and any fixed u_{i+1} in U , $\Pr[v_{i+1} = u_{i+1} \mid (v_0, \dots, v_i) = (u_0, \dots, u_i)] = \Pr[v_{i+1} = u_{i+1} \mid v_i = u_i] = \Pr_{\mathcal{D}_{u_i}}[(u_i, u_{i+1})]$.

Let H be the directed graph with node set \mathbf{N} and edge set $\{(i, j) : i \in \mathbf{N}, 0 \leq j \leq i + 1\}$. Let H' be the subgraph of H induced by the edges $\{(i + 1, i), (i, i + 1) : i \in \mathbf{N}\} \cup \{(0, 0), (1, 1)\}$.

Let p and q be reals in $(0, 1]$. We now define two random walks, $\omega_{p,q}$ and $\omega'_{p,q}$, on graphs H and H' , respectively. The walk $\omega_{p,q} = \langle w \rangle$ is characterized by: (i) $\Pr[w_{i+1} \leq j - 1 \mid w_i = j] \geq p$, for any integer $j > 1$, (ii) $\Pr[w_{i+1} = 0 \mid w_i = j] \geq q$, for j equal 0 or 1, and (iii) $\Pr[w_{i+1} = 2 \mid w_i = 1] \leq 1 - p$. The walk $\omega'_{p,q} = \langle w' \rangle$ is characterized by: (i) $\Pr[w'_{i+1} = j - 1 \mid w'_i = j] = p$, for all integer $j > 1$, (ii) $\Pr[w'_{i+1} = 0 \mid w'_i = j] = q$, for j equal 0 or 1, and (iii) $\Pr[w'_{i+1} = 2 \mid w'_i = 1] = 1 - p$. We note that the sequence $\langle \max\{s', t\} \rangle$ represents the random walk $\omega_{p,q}$ with appropriate values for p and q , as determined by Lemmas 5.11 and 5.12. We analyze random walk $\omega_{p,q}$ by first showing that $\omega_{p,q}$ “dominates” $\omega'_{p,q}$ with respect to the properties of interest. The random walk $\omega'_{p,q}$ is easier to analyze as it is exactly characterized by p and q . Lemmas 5.16 and 5.18 show that the bias of $\omega_{p,q}$ towards 0 is more than that of $\omega'_{p,q}$.

Lemma 5.16 *For all i and k in \mathbf{N} , for random walks $\omega_{p,q}$ and $\omega'_{p,q}$, we have $\Pr[w_i \leq k] \geq \Pr[w'_i \leq k]$.*

Proof: We prove the claim by induction on i . The base case $i = 0$ is trivial. Assume the claim holds for i and any k .

Let $k \geq 2$. Then

$$\begin{aligned} \Pr[w'_{i+1} \leq k] &= \Pr[w'_i \leq k - 1] + p \Pr[k \leq w'_i \leq k + 1] \\ &= (1 - p) \Pr[w'_i \leq k - 1] + p \Pr[w'_i \leq k + 1] \\ \Pr[w_{i+1} \leq k] &\geq \Pr[w_i \leq k - 1] + p \Pr[k \leq w_i \leq k + 1] \\ &= (1 - p) \Pr[w_i \leq k - 1] + p \Pr[w_i \leq k + 1]. \end{aligned}$$

If $k = 1$ or 0 , respectively, then

$$\begin{aligned} \Pr[w'_{i+1} \leq 1] &= \Pr[w'_i \leq 0] + p \Pr[1 \leq w'_i \leq 2] \\ &= (1 - p) \Pr[w'_i \leq 0] + p \Pr[w'_i \leq 2] \\ \Pr[w_{i+1} \leq 1] &\geq \Pr[w_i \leq 0] + p \Pr[1 \leq w_i \leq 2] \\ &= (1 - p) \Pr[w_i \leq 0] + p \Pr[w_i \leq 2]. \end{aligned}$$

$$\begin{aligned} \Pr[w'_{i+1} \leq 0] &= q \Pr[w'_i \leq 0] + q \Pr[w'_i = 1] \\ &= q \Pr[w'_i \leq 1] \\ \Pr[w_{i+1} \leq 0] &\geq q \Pr[w_i \leq 0] + q \Pr[w_i = 1] \\ &= q \Pr[w_i \leq 1]. \end{aligned}$$

The lemma follows by induction. ■

Let $z_i(\omega)$ be the random variable denoting the number of steps taken to reach node 0 starting from node i , for a random walk ω . An *excursion* of length ℓ in a graph W with node set \mathbf{N} is a walk that starts at node 0 and first returns to the start node at time ℓ , for all ℓ in \mathbf{N} . For all i such that $w_i = 0$, let $\ell_i(\omega)$ be the random variable that gives the length of the excursion in ω starting at time i . We note that for all i , $\ell_i(\omega)$ equals $z_0(\omega)$.

Lemma 5.17 $\Pr[z_i(\omega'_{p,q}) \leq \ell] \leq \Pr[z_{i-1}(\omega'_{p,q}) \leq \ell]$, for all ℓ and all $i > 0$.

Proof: We omit the subscript p, q for convenience. We use induction on ℓ . The base case $\ell = 0$ is trivial. Let $\ell \geq 1$. If $i > 2$ then:

$$\begin{aligned} \Pr[z_{i-1}(\omega') \leq \ell] &= p \Pr[z_{i-2}(\omega') \leq \ell - 1] + (1 - p) \Pr[z_i(\omega') \leq \ell - 1] \\ &\geq p \Pr[z_{i-1}(\omega') \leq \ell - 1] + (1 - p) \Pr[z_{i+1}(\omega') \leq \ell - 1] = \Pr[z_i(\omega') \leq \ell] \end{aligned}$$

where the inequality follows by induction. If $i = 2$ or 1 , then we have, respectively:

$$\begin{aligned} \Pr[z_1(\omega') \leq \ell] &= q + (1 - q - (1 - p)) \Pr[z_1(\omega') \leq \ell - 1] + (1 - p) \Pr[z_2(\omega') \leq \ell - 1] \\ &\geq p \Pr[z_1(\omega') \leq \ell - 1] + (1 - p) \Pr[z_2(\omega') \leq \ell - 1] \\ &\geq p \Pr[z_1(\omega') \leq \ell - 1] + (1 - p) \Pr[z_3(\omega') \leq \ell - 1] = \Pr[z_2(\omega') \leq \ell] \\ \Pr[z_0(\omega') \leq \ell] &= q + (1 - q) \Pr[z_1(\omega') \leq \ell - 1] \\ &\geq q \Pr[z_0(\omega') \leq \ell - 1] + (p - q) \Pr[z_1(\omega') \leq \ell - 1] + (1 - p) \Pr[z_2(\omega') \leq \ell - 1] \\ &= \Pr[z_1(\omega') \leq \ell] \end{aligned}$$

where the last inequalities in both calculations also follow by induction. ■

Lemma 5.18 For all i and ℓ in \mathbf{N} , we have $\Pr[z_i(\omega_{p,q}) \leq \ell] \geq \Pr[z_i(\omega'_{p,q}) \leq \ell]$.

Proof: We omit the subscript p, q for convenience. Let $p_j = \Pr[w_{i+1} \leq j - 1 \mid w_i = j]$, for $j > 1$, and $h_j = \Pr[w_{i+1} = j \mid w_i = j]$, for all j in \mathbf{N} . Note that $p \leq p_j$, for all $j > 1$, $q \leq \min\{p_1, h_0\}$, and $p \leq q$.

If $i \geq 2$ then:

$$\begin{aligned} \Pr[z_i(\omega') \leq \ell] &= p \Pr[z_{i-1}(\omega') \leq \ell - 1] + (1 - p) \Pr[z_{i+1}(\omega') \leq \ell - 1] \\ &\leq p_i \Pr[z_{i-1}(\omega') \leq \ell - 1] + (1 - p_i) \Pr[z_{i+1}(\omega') \leq \ell - 1] \\ &\leq p_i \Pr[z_{i-1}(\omega') \leq \ell - 1] + h_i \Pr[z_i(\omega') \leq \ell - 1] + (1 - p_i - h_i) \Pr[z_{i+1}(\omega') \leq \ell - 1] \end{aligned}$$

The inequalities above follow from Lemma 5.17 and the fact that $p \leq p_i$. If $i = 1$ or 0 then

$$\begin{aligned} \Pr[z_1(\omega') \leq \ell] &= q + (1 - q - (1 - p)) \Pr[z_1(\omega') \leq \ell - 1] + (1 - p) \Pr[z_2(\omega') \leq \ell - 1] \\ &\leq p_1 + h_1 \Pr[z_1(\omega') \leq \ell - 1] + (1 - p_1 - h_1) \Pr[z_2(\omega') \leq \ell - 1] \\ \Pr[z_0(\omega') \leq \ell] &= q + (1 - q) \Pr[z_1(\omega') \leq \ell - 1] \\ &\leq h_0 + (1 - h_0) \Pr[z_1(\omega') \leq \ell - 1] \end{aligned}$$

where we use Lemma 5.17, and the inequalities $q \leq \min\{p_1, h_0\}$ and $p \geq q$.

By induction on ℓ , we have that $\Pr[z_i(\omega') \leq \ell] \leq \Pr[z_i(\omega) \leq \ell]$, for all i . ■

We now formalize the notion of the domination of $\omega_{p,q}$ over $\omega'_{p,q}$. For any i , let τ_i (resp., τ'_i) denote the smallest $j \geq 0$ such that $w_{i+j} = 0$ (resp., $w'_{i+j} = 0$). We note that by letting $\langle w \rangle$ represent $\langle \max\{s', t\} \rangle$, the terminating step τ is given by $i^* + \tau_{i^*}$.

Lemma 5.19 For any i and $j \geq i$, we have $\Pr[\tau_i \leq j] \geq \Pr[\tau'_i \leq j]$.

Proof: In the following, we omit the subscript p, q for convenience. By the definitions of $\tau_i, z_i(\omega), \tau'_i$, and $z_i(\omega')$ we have:

$$\begin{aligned}\Pr[\tau_i \leq j] &= \sum_{0 \leq k \leq i} \Pr[w_i = k] \Pr[z_k(\omega) \leq j], \text{ and} \\ \Pr[\tau'_i \leq j] &= \sum_{0 \leq k \leq i} \Pr[w'_i = k] \Pr[z_k(\omega') \leq j].\end{aligned}$$

We now prove the desired claim as follows:

$$\begin{aligned}\Pr[\tau_i \leq j] &= \sum_{0 \leq k \leq i} \Pr[w_i = k] \Pr[z_k(\omega) \leq j] \\ &\geq \sum_{0 \leq k \leq i} \Pr[w_i = k] \Pr[z_k(\omega') \leq j] \\ &\geq \sum_{0 \leq k \leq i} \Pr[w'_i = k] \Pr[z_k(\omega') \leq j].\end{aligned}$$

(In the second step, we use Lemma 5.18. For the third step we first invoke Lemma 5.16 and then invoke Lemma A.1 with the substitution $(i, k, \Pr[w_i = k], \Pr[w'_i = k], \Pr[z_k(\omega') \leq j])$ for (m, i, p_i, q_i, n_i) . We note that one of the conditions for the latter invocation, namely, $\Pr[z_k(\omega') \leq j]$ is non-increasing with k , follows from Lemma 5.17.) \blacksquare

The following claim is proved using Raney's lemma [7, 15].

Lemma 5.20 For all i and ℓ in \mathbb{N} , we have $\Pr[\ell_i(\omega'_{p,q}) = \ell + 1 \mid w'_i = 0] \leq \max\{1 - q, 5(p - q)\} \Pr[\ell_i(\omega'_{p,q}) = \ell \mid w'_i = 0]$.

Proof: Since ω' is a random walk, $\Pr[\ell_i(\omega') = \ell \mid (w'_0, \dots, w'_{i-1}, w'_i) = (u_0, \dots, u_{i-1}, 0)] = \Pr[\ell_0(\omega') = \ell \mid w'_0 = 0]$, for any u_0, \dots, u_{i-1} in \mathbb{N} . For the remainder of the proof, we assume that w'_0 is 0.

For $\ell = 1$, the desired claim holds since $\Pr[\ell_0(\omega') = 2] / \Pr[\ell_0(\omega') = 1] = (1 - q)$. We now consider $\ell \geq 2$. Let \mathcal{E}_j denote the event that the random walk does not reach node 0 in the first j steps. That is, \mathcal{E}_j is the event that w'_k is non-zero for all k in $[1, j]$. For all j , let α_j denote the probability that w'_{j+1} is 1 and \mathcal{E}_{j+1} holds, given that w'_1 is 1. For convenience, we assume that α_{-1} equals $1/(p - q)$. We obtain that:

$$\Pr[\ell_0(\omega') = \ell] = (1 - q) \cdot \alpha_{\ell-2} \cdot q \tag{7}$$

It thus follows that the ratio of $\Pr[\ell_0(\omega') = \ell + 1]$ and $\Pr[\ell_0(\omega') = \ell]$ equals $\alpha_{\ell-1} / \alpha_{\ell-2}$. The remainder of the proof is devoted to obtaining an upper bound on α_{j+1} / α_j for all $j \geq 0$.

For $m \geq 0$, let β_m denote the probability that \mathcal{E}_{2m+1} holds and $w'_{2m+1} = 1$ and the edge $(1, 1)$ is not traversed in any of the first $2m + 1$ steps, given that w'_1 is 1. Using Raney's lemma [7, 15], we obtain that β_m equals $\frac{1}{2m+1} \binom{2m+1}{m} (p(1-p))^m$. By the definitions of α_j and β_m , it follows that:

$$\begin{aligned}\alpha_j &= \sum_{0 \leq m \leq \lfloor j/2 \rfloor} \beta_m \cdot (p - q) \cdot \alpha_{j-2m-1} \\ &= \sum_{0 \leq m \leq \lfloor j/2 \rfloor} \frac{1}{2m+1} \binom{2m+1}{m} (p(1-p))^m \cdot (p - q) \cdot \alpha_{j-2m-1}.\end{aligned}$$

We now prove by induction on $j \geq 2$ that α_{j+1} / α_j is at most $(p - q)$. The induction base holds since α_0 is 1 and α_1 is $5(p - q)$. For the induction hypothesis, we assume that α_{j+1} / α_j is at most $5(p - q)$ for

all $j \leq k - 1$. If k is even, then we have:

$$\begin{aligned}\alpha_{k+1}/\alpha_k &\leq \max_{0 \leq m \leq k/2} \alpha_{k-2m}/\alpha_{k-2m-1} \\ &\leq 5(p-q),\end{aligned}$$

where the last inequality follows from the induction hypothesis. If k is odd, then we have:

$$\begin{aligned}\alpha_{k+1}/\alpha_k &\leq \max \left\{ \left(\frac{1}{k} \binom{k}{(k-1)/2} (p-q)^2 + \frac{1}{k+2} \binom{k+2}{(k+1)/2} p \right) / \left(\frac{1}{k} \binom{k}{(k-1)/2} (p-q) \right), \right. \\ &\quad \left. \max_{0 \leq m \leq (k-3)/2} \alpha_{k-2m}/\alpha_{k-2m-1} \right\} \\ &\leq \max\{5(p-q), 5(p-q)\} \\ &= 5(p-q),\end{aligned}$$

where the second inequality follows from the induction hypothesis and the inequalities $1-p \leq (p-q)^2$, and $\binom{k+2}{(k+1)/2} \leq 4 \binom{k}{(k-1)/2}$. The desired claim of the lemma follows from the upper bound on α_{k+1}/α_k and Equation 7. \blacksquare

We now let ω and ω' denote the random walks $\omega_{p,q}$ and $\omega'_{p,q}$, respectively, where $p = 1 - 2\varepsilon^2$ and $q = 1 - 2\varepsilon$. Lemmas 5.12 and 5.11 imply that ω characterizes the random walk corresponding to the sequence $\langle \max\{s', t\} \rangle$. Consider the random walk ω' . Assume that at each step we only reveal whether $w'_i = 0$ or not. We can define a sequence $\langle v \rangle$ associated with $\langle w' \rangle$ as follows: $v_j = G$ iff $w'_j = 0$, and $v_j = B$ otherwise.

Lemma 5.21 *Let i be in $[(\log n/b) - 1]$. Given any fixed sequence $\langle v \rangle_{i-1}$ of B, G values, the probability that w'_i is 0 is at least $1 - 10\varepsilon$.* \blacksquare

Proof: Assume that $v_j = G$. What is the probability that $v_i = G$, $i > j$, if we know that $v_k = B$, for all integer k in the interval $[j+1, i]$? From Lemma 5.20, it follows that this probability is at least $1 - 10\varepsilon$, since this is a lower bound on the probability that there is an excursion of length $i - j$ starting at j in H' , given that there is an excursion of length at least $i - j$ starting at j in H' . Given any fixed B, G sequence $\langle u \rangle_{j-1}$, $\Pr[v_i = G \mid (v_0, \dots, v_{i-1}) = (u_0, \dots, u_{j-1}, G, B, \dots, B)] = \Pr[v_i = G \mid (v_j, \dots, v_{i-1}) = (G, B, \dots, B)]$. Since this holds for any $j > 0$ and since $w'_i = 0$ iff $v_i = G$, we have $\Pr[w'_i \mid (v_0, \dots, v_{i-1}) = (u_0, \dots, u_{i-1})] \geq 1 - 10\varepsilon$. \blacksquare

Our main claim about the random walk ω follows from Lemmas 5.19 and 5.21.

Lemma 5.22 *For any i in $[(\log n)/b - 1]$ and any nonnegative integer j , the probability that τ_i is at least j is at most $(10\varepsilon)^j$.*

Proof: By Lemma 5.21, the probability that τ'_i is at least j is at most $(10\varepsilon)^j$. The desired claim follows from Lemma 5.19. \blacksquare

Using Lemma 5.22, we derive an upper bound on $E[c(x_i, x_{i+1})]$ and $E[c(y_i, y_{i+1})]$ for all i .

Lemma 5.23 *For any i in $[(\log n)/b - 1]$, $E[c(x_i, x_{i+1})]$ and $E[c(y_i, y_{i+1})]$ are both $O(a_i)$.*

Proof: We first observe that $c(x_i, x_{i+1})$ (resp., $c(y_i, y_{i+1})$) is at most a_k (resp., b_k), where k is the least $j \geq i$ such that s_j (resp., t_j) belongs to $\{0, 1, 2\}$ (resp., $\{0\}$); if such a j does not exist, then k is $(\log n)/b - 1$. Thus, k is at most $i + \tau_i$. By Lemma 5.22, it follows that for any $j \geq i$, the probability that $k \geq j$ is at most $(10\varepsilon)^{j-i}$. By Lemma 5.9, we thus have:

$$\begin{aligned}E[c(x_i, x_{i+1})] &\leq \sum_{j \geq i} a_i (10\varepsilon)^{j-i} 2^{b \log_2 2^{j-i}} = O(a_i) \text{ and} \\ E[c(y_i, y_{i+1})] &\leq \sum_{j \geq i} b_i (10\varepsilon)^{j-i} 2^{b \log_2 2^{j-i}} = O(b_i),\end{aligned}$$

since $10\varepsilon 2^{b \log_\delta 2} < 1$. ■

We now use Lemmas 5.9, 5.22, and 5.23 to establish Theorem 1.

Proof of Theorem 1: By Equation 6, the expected cost of the read operation is bounded by the expected value of $f(\ell(A)) \sum_{0 \leq i < \tau} O(c(x_i, x_{i+1}) + c(y_i, y_{i+1}))$. We separately place bounds on $E[\sum_{0 \leq i < i^*} (c(x_i, x_{i+1}) + c(y_i, y_{i+1}))]$ and $E[\sum_{i^* \leq i < \tau} (c(x_i, x_{i+1}) + c(y_i, y_{i+1}))]$. By Lemmas 5.9 and 5.23, the first term is $O(a_{i^*} + b_{i^*})$.

We place a bound on $E[\sum_{i^* \leq i < \tau} (c(x_i, x_{i+1}) + c(y_i, y_{i+1}))]$ as follows. Since τ is $i^* + \tau_{i^*}$, by Lemma 5.22, we obtain that for any $j \geq 0$, the probability that $\tau \geq i^* + j$ is at most $(10\varepsilon)^j$. Therefore, $E[\sum_{i^* \leq i < \tau} (c(x_i, x_{i+1}) + c(y_i, y_{i+1}))]$ is at most:

$$\begin{aligned} & \sum_{j \geq 0} j (10\varepsilon)^j (a_{i^*+j} + b_{i^*+j}) \\ & \leq \sum_{j \geq 0} j (10\varepsilon)^j 2^{jb \log_\delta 2} (a_{i^*} + b_{i^*}) \\ & = O(a_{i^*} + b_{i^*}), \end{aligned}$$

since $10\varepsilon 2^{b \log_\delta 2} < 1$. By Lemma 5.9, the claim of the theorem follows. ■

Proof of Theorem 2: Consider an insert operation executed by x for any object. The expected cost of the operation is bounded by $E[\sum_{0 \leq i < \log n/b} c(x_i, x_{i+1})]$, which by Lemmas 5.9 and 5.23 is $O(a_{(\log n)/b-1}) = O(C)$.

We now consider the cost of the delete operation. By Lemma 5.6, for each i , the number of reverse (i, j) -neighbors of x_i for any j is $O(\log n)$ whp, where x_i is the i th node in the primary neighbor sequence of x . Therefore, the expected cost of the delete operation executed by x is bounded by the product of $E[\sum_{0 \leq i < \log n/b} c(x_i, x_{i+1})]$ and $O(\log n)$. By Lemma 5.23, it follows that the expected cost of a delete operation is $O(C \log n)$. ■

5.4 Auxiliary Memory

Proof of Theorem 3: We first place an upper bound on the size of the neighbor table of any u in V . By definition, the number of primary and secondary neighbors of u is at most $(d+1)2^b(\log n)/b$, which is $O(\log n)$. By Corollary 5.6.1, the number of reverse neighbors of u is $O(\log^2 n)$ whp.

We next place an upper bound on the size of the pointer list of any u in V . The size of $Ptr(u)$ is at most the number of triples of the form (A, v, \cdot) , where A is in \mathcal{A} and v is in V such that: (i) there exists i in $[(\log n)/b]$ such that v is an i -leaf of u , (ii) $A[j] = u[j]$ for all j in $[i]$, and (iii) A is in the main memory of v .

By Lemma 5.7, the number of i -leaves of u is $O(2^{ib} \log n)$ whp. The probability that $A[j] = u[j]$, for all j in $[i]$, is at most $1/2^{ib}$. Since the number of objects in the main memory of any node is at most ℓ , it follows that whp, $|Ptr(u)|$ is at most $\sum_{i \in [(\log n)/b]} O(\ell \log n)$ which is $O(\ell \log^2 n)$.

Combining the bounds on the sizes of the neighbor table and pointer list, we obtain that the size of the auxiliary memory of u is $O(\ell \log^2 n)$ whp. ■

5.5 Adaptability

Proof of Theorem 4: By Lemma 5.6, for any node u , the number of nodes of which u is a primary or secondary neighbor is $O(\log n)$ expected and $O(\log^2 n)$ whp. Moreover, u is a reverse neighbor of $O(\log n)$ nodes since u has $O(\log n)$ primary neighbors. Therefore, the adaptability of our scheme is $O(\log n)$ expected and $O(\log^2 n)$ whp. ■

6 Future Work

We would like to extend our study to more general classes of cost functions and determine tradeoffs among the various complexity measures. It would also be interesting to consider models that allow faults in the

network. We believe that our access scheme can be extended to perform well in the presence of faults, as the distribution of control information in our scheme is balanced among the nodes of the network.

Acknowledgments

The authors would like to thank Madhukar Korupolu and Satish Rao for several helpful discussions.

References

- [1] B. Awerbuch, Y. Bartal, and A. Fiat. Distributed paging for general networks. In *Proceedings of the 7th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 574–583, January 1996.
- [2] B. Awerbuch and D. Peleg. Routing with polynomial communication space tradeoff. *SIAM Journal of Discrete Mathematics*, 5:151–162, 1990.
- [3] B. Awerbuch and D. Peleg. Sparse partitions. In *Proceedings of the 31st Annual IEEE Symposium on Foundations of Computer Science*, pages 503–513, 1990.
- [4] Y. Bartal, A. Fiat, and Y. Rabani. Competitive algorithms for distributed data management. In *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*, pages 39–47, May 1992.
- [5] H. Chernoff. A measure of the asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Annals of Mathematical Statistics*, 23:493–509, 1952.
- [6] S. Dolev, E. Kranakis, D. Krizanc, and D. Peleg. Bubbles: Adaptive routing scheme for high-speed dynamic networks. In *Proceedings of the 27th Annual ACM Symposium on the Theory of Computing*, pages 528–537, 1995.
- [7] R. L. Graham, D. E. Knuth, and O. Patashnik. *Concrete Mathematics*. Addison-Wesley, Reading, MA, 1989.
- [8] J. D. Guyton and M. F. Schwartz. Locating nearby copies of replicated Internet servers. In *Proceedings of ACM SIGCOMM*, pages 288–298, 1995.
- [9] D. Karger, E. Lehman, T. Leighton, M. Levine, D. Lewin, and R. Panigrahy. Relieving hot spots on the World Wide Web. In *Proceedings of the 29th Annual ACM Symposium on the Theory of Computing*, May 1997.
- [10] R. Karp, M. Luby, and F. Meyer auf der Heide. Efficient PRAM simulation on a distributed memory machine. In *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*, pages 318–326, May 1992.
- [11] S. J. Mullender, editor. *Distributed Systems*. Addison-Wesley, 1993.
- [12] S. J. Mullender and P. M. B. Vitányi. Distributed match-making. *Algorithmica*, 3:367–391, 1988.
- [13] C. G. Plaxton and R. Rajaraman. Fast fault-tolerant concurrent access to shared objects. In *Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science*, pages 570–579, October 1996.
- [14] A. G. Ranade. How to emulate shared memory. *Journal of Computer and System Sciences*, 42:307–326, 1991.
- [15] G. N. Raney. Functional composition patterns and power series reversion. *Transactions American Mathematical Society*, 94:441–451, 1960.

- [16] E. Upfal and A. Wigderson. How to share memory in a distributed system. *JACM*, 34:116–127, 1987.
- [17] M. Van Steen, F. J. Hauck, and A. S. Tanenbaum. A model for worldwide tracking of distributed objects. In *Proceedings of TINA '96*, pages 203–212, September 1996.

A A technical lemma

Lemma A.1 *Let m be a nonnegative integer and let $\langle n \rangle_m$ be a sequence of non-increasing reals. Let $\langle p \rangle_m$ and $\langle q \rangle_m$ be two sequences of reals such that: (i) for all j in $[m]$, $\sum_{0 \leq i \leq j} p_i \geq \sum_{0 \leq i \leq j} q_i$ and (ii) $\sum_{0 \leq i \leq m} p_i = \sum_{0 \leq i \leq m} q_i$. Then, we have:*

$$\sum_{0 \leq i \leq m} p_i n_i \geq \sum_{0 \leq i \leq m} q_i n_i.$$

Proof: The proof is by induction on m . The induction basis is trivial. For the induction hypothesis, we assume that the statement of the lemma holds for m . We now establish the claim for $m + 1$.

$$\begin{aligned} \sum_{0 \leq i \leq m+1} p_i n_i &= q_0 n_0 + (p_0 - q_0) n_0 + \sum_{1 \leq i \leq m+1} p_i n_i \\ &\geq q_0 n_0 + (p_0 - q_0 + p_1) n_1 + \sum_{2 \leq i \leq m+1} p_i n_i \\ &\geq q_0 n_0 + \sum_{1 \leq i \leq m+1} q_i n_i \\ &= \sum_{0 \leq i \leq m} q_i n_i. \end{aligned}$$

(The third step follows from the inequalities $n_0 \geq n_1$ and $p_0 \geq q_0$, and the induction hypothesis. We note that the induction hypothesis can be invoked since $p_0 - q_0 + p_1 + \sum_{2 \leq i \leq j} p_i \leq \sum_{1 \leq i \leq j} q_i$ and $p_0 - q_0 + p_1 + \sum_{2 \leq i \leq m+1} p_i = \sum_{1 \leq i \leq m+1} q_i$.) ■