

Distributed Object Location in a Dynamic Network*

Kirsten Hildrum, John D. Kubiatowicz, Satish Rao and Ben Y. Zhao
Computer Science Division, University of California at Berkeley
{hildrum, kubitron, satishr, ravenben}@cs.berkeley.edu

ABSTRACT

Modern networking applications replicate data and services widely, leading to a need for *location-independent routing* – the ability to route queries directly to objects using names independent of the objects’ physical locations. Two important properties of a routing infrastructure are *routing locality* and *rapid adaptation* to arriving and departing nodes. We show how these two properties can be efficiently achieved for certain network topologies. To do this, we present a new distributed algorithm that can solve the nearest-neighbor problem for these networks. We describe our solution in the context of Tapestry, an overlay network infrastructure that employs techniques proposed by Plaxton, Rajaraman, and Richa [20].

1 Introduction

In today’s chaotic network, data and services are mobile and replicated widely for availability, durability, and locality.¹ This has led to a renewed interest in techniques for routing queries to objects using names that are independent of their locations. The notion of *routing* is that queries are forwarded from node to node until they reach their destinations. The *location-independent routing* problem has spawned a host of proposals, many of them in the context of data sharing infrastructures such as OceanStore [14], FarSite [3], CFS [12] and PAST [24]. To permit locality optimizations, it is important that the routing process use as few network hops as possible and that these hops be as short as possible.

Properties that we would like from a location-independent routing infrastructure include:

1. *Deterministic Location*: Objects should be located if they exist anywhere in the network.
2. *Routing Locality*: Routes should have low *stretch*,² not just a small number of application-level hops. Sending queries to the nearest copy across the shortest path possible is the ideal.
3. *Minimality and Load Balance*: The infrastructure must not place undue stress on any of its components; this implies minimal storage and balanced computational load.
4. *Dynamic Membership*: The system must adapt to arriving and departing nodes while maintaining the above properties.

Although clearly desirable, the first property is not guaranteed by existing peer-to-peer systems such as Gnutella [18] and FreeNet [7].

A simple object location and routing scheme would employ a centralized directory of object locations. Servers would *publish* the

existence of objects by inserting entries into the directory. Clients would send *queries* to the directory, which forwards them to their destinations. This solution, while simple, induces a heavy load on the directory server. Moreover, when a nearby server happens to contain the object, the client must still interact with the potentially distant directory server. The average routing latency of this technique is proportional to the average diameter of the network – independent of the actual distance to the object. Worse, it is neither fault tolerant nor scalable, since the directory becomes a single point of both failure and contention.

Several recent proposals, Chord [26], CAN [22] and Pastry [23], address the load aspect of this problem by distributing the directory information over a large number of nodes. In particular, they can find an object with a polylogarithmic number of application-level network hops while ensuring that no node contains much more than its share of the directory entries. Moreover, they can support the introduction and removal of new participants in the peer-to-peer network. Unfortunately, while these approaches use a number of overlay hops that is polylogarithmic, the actual network latencies incurred by queries can be significantly more than those incurred by finding the object in the centralized directory.

An alternative solution is to broadcast an object’s location to every node in the network. This allows clients to easily find the nearest copy of the object, but requires a large amount of resources to publish and maintain location information, including both network bandwidth and storage. Furthermore, it requires full knowledge of the participants of the network. In a dynamic network, maintaining a list of participants is a significant problem in its own right.

We describe our results in the context of the Tapestry overlay routing and location infrastructure [32, 30]. Tapestry uses as a starting point the distributed data structure of Plaxton, Rajaraman, and Richa [20], which we will refer to as the PRR scheme. Their proposal yields routing locality with balanced storage and computational load. However, it does not provide dynamic maintenance of membership. The original statement of the algorithm required a static set of participating nodes as well as significant work to preprocess this set to generate a routing infrastructure. Additionally, the PRR scheme was unable to adapt to changes such as node failures. This paper extends their algorithms to a dynamic network.

1.1 Related Work

Several existing object location schemes exhibit routing locality, including Plaxton, Rajaraman, and Richa (PRR) [20] and Awerbuch and Peleg [1], and Rajaraman, Richa, Vöcking and Vuppuluri [21]. All of these provide the publication and deletion of objects with only a logarithmic number of messages and guarantee a low stretch, where stretch is defined as the ratio between the actual latency or distance to an object and the shortest distance. The PRR scheme finds objects with constant stretch for a specific class of network topologies while ensuring that no node has too many

*This research supported by NSF career award #ANI-9985250, NFS ITR award #CCR-0085899, NSF Graduate Research Fellowship, UC MICRO award #00-049, and NSF grant #CCR-0105533.

¹By locality we mean the ability to exploit local resources over remote ones whenever possible [31].

²Stretch is the ratio between the distance traveled by a query to an object and the minimal distance from the query origin to the object.

Scheme	Insert Cost	Space	Stretch, Metric	Hops	Balanced?
CHORD [26]	$O(\log^2 n)$	$O(n \log n)$	-	$O(\log n)$	yes
CAN [22]	$O(rn^{1/r})$	nr	-	$rn^{1/r}$	yes
Pastry [23]	$O(\log^2 n)$	$O(n \log n)$	-	$O(\log n)$	yes
This Paper (Tapestry)	$O(\log^2 n)$	$O(n \log n)$	-	$O(\log n)$	yes
Awerbuch, Peleg[1]	-	$O(n \log^3 n)$	$O(\log^2 n)$, general	$O(\log^2 n)$	no
RRVV[21]	maybe $O(\log^3 n)$	$O(n \log^3 n)$	$O(\log^3 n)$, general	$O(\log^2 n)$	yes
PRR [20]	-	$O(n \log n)$	$O(1)$, special	$O(\log n)$	yes
PRR + This Paper	$O(\log^2 n)$	$O(n \log n)$	$O(1)$, special	$O(\log n)$	yes
PRR v.0 + This Paper	-	$O(n \log^2 n)$	$O(\log^3 n)$, general	$O(\log^2 n)$	no

Table 1: Comparison of Object Location Systems. In this table, n is the number of nodes. For simplicity, we assume that the network diameter is polynomial in n , and that assume the number of objects is $O(n)$. Both stretch and hops refer to an object search. Insert cost shows the number of hops or messages needed for node insertion; a - means the system does not provide an algorithm. In RRVV, the number of changes needed is polylogarithmic, but they do not give an algorithm to make the changes. In most cases, the time for insertion is given with high probability. In some cases, various messages can be sent in parallel; we did not allow for this optimization in stating the bounds in this table.

directory entries. Awerbuch and Peleg route with within a polylogarithmic factor of optimal for general network topologies, but do not balance the load. Unfortunately, both the PRR and Awerbuch-Peleg schemes assume full knowledge of the participating nodes, or, equivalently, they assume that the network is static. The RRVV scheme balances the load, bounding the space at every node, and while only a polylogarithmic number of nodes need change when a node enters or leaves the network, they also do not give a method to find the nodes that need to be updated.

There is also an abundance of theoretical work on finding compact routing tables [2, 9, 19, 28] whose techniques are closely related to those in this paper. See [11] for a survey. A recent and closely related paper is that of Thorup and Zwick, who showed that a sampling based scheme similar to that of PRR could be used to find small stretch routing tables and/or answer approximate distance queries in arbitrary metric spaces.³

Most of the recent work on peer-to-peer networks ignore stretch. Chord [26] constructs a distributed lookup service using a routing table of logarithmic size. Nodes are arranged into a large virtual circle. Each node maintains pointers to predecessor and successor nodes, as well as a logarithmic number of “chords” which cross greater distances within the circle. Queries are forwarded along chords until they reach their destination. CAN [22] places objects into a virtual, high-dimensional space. Queries are routed along axes in this virtual space until they reach their destination. Pastry [23] is loosely based on the PRR scheme, routing queries via successive resolution of digits in a high-dimensional name space. While its overlay construction leverages network proximity metrics, it does not provide the same stretch as the PRR scheme in object location. A recent paper by Li and Plaxton [15] presents a simplified version of the PRR scheme that may perform well in practice. All of these schemes can find objects with a polylogarithmic number of application-level network hops, while ensuring that no node contains more than its share of directory entries. In addition, Chord and CAN have run-time heuristics to reduce object location cost, so they may perform well in practice. Finally, all of these systems support the introduction and removal of nodes.

Recent peer-to-peer applications can locate objects in a dynamic network. Gnutella [18] utilizes a bounded broadcast mechanism to search neighbors for documents. FreeNet [7] utilizes a chaotic routing scheme in which objects are published to a set of nearest neighbors and queries follow gradients generated by object pointers; the behavior of FreeNet appears to converge somewhat toward the PRR

scheme when a large number of objects are present.⁴ Neither of these techniques are guaranteed to find objects.

Table 1 summarizes related work alongside our contributions. Systems with no entry in the “Stretch, Metric” column do not consider stretch at all; those with “special” assume the metric space has a certain low-expansion property described in Section 3.

1.2 Results

Our goals are not only to derive the best possible asymptotic results, but also to analyze the simple schemes that are the basis of the PRR and the Tapestry algorithms. This paper includes three main results:

- We present a simplification of the PRR scheme for object location. We cannot prove that this object location scheme meets the same bounds on stretch as the PRR scheme; however, it appears to perform well in practice.
- We extend this scheme (as well as the PRR approach) to deal with a changing participant set. We allow nodes to arrive and depart while maintaining the ability to locate existing objects and publish new objects. This works for a slightly broader class of metric spaces than assumed by PRR.
- We observe that a static version of the PRR scheme can be used for general metric spaces (*i.e.*, spaces that do not meet the conditions assumed by PRR) to get results similar to those of Awerbuch and Peleg [1].

Table 1 gives a summary of some of the previous results along with ours. Note that our result for general metrics can be improved using results of Thorup and Zwick [27] to use only $O(n \log n)$ space.

Techniques: The crux of our method for inserting nodes into the network lies in an algorithm for maintaining nearest neighbors in a restricted metric space. Our approach is similar in spirit to that of Karger and Ruhl [13], who give an algorithm for answering nearest neighbor queries in a similarly restricted metric space.⁵

The idea behind both the nearest-neighbor algorithm of Karger and Ruhl and the one presented here is to find the nearest neighbor by repeatedly finding some node halfway between the current node and the query node. If this is done $\log n$ times, one finds the closest node. The restricted metric spaces considered in both these papers mean that there is a substantial fraction of nodes at about the right distance, so halving the distance can be implemented by sampling

³A network topology gives a metric space.

⁴This is a qualitative statement at this time.

⁵Clarkson also presented a very similar approach in[8].

from nodes within the correct radius. The difficulty is maintaining a structure to do the sampling in a dynamic network.

Karger and Ruhl suggested this general approach in [13] and then present a specific data structure to accomplish it. Their data structure uses a random permutation to maintain the random sampling—an approach is reminiscent of the Chord network infrastructure. Our search algorithm also aims to halve the distance at each step, but we build a different data structure with a different search algorithm. In particular, we use random names to build a tree (for load balancing purposes, many trees) on which we search. This set of trees is the same as the set of trees used in the object location system described in this paper, which means that our search algorithm can share the data structure with the object location algorithm.

Recently, Krauthgamer and Lee developed a simple and deterministic nearest-neighbor data structure that has applications in a broader class of metric spaces. Their data structure is somewhat similar to the one described in this paper, but it is not yet clear if their data structure can be distributed in a load-balanced way.⁶

We also prove that an alternate scheme by Plaxton, Rajaraman, and Richa (called PRR v.0 in Table 1) gives a low stretch solution for general metric spaces. This follows from arguments similar to those used by Bourgain [4] for metric embeddings. In particular, we show that this scheme leads to a covering of the graph by trees such that for any two nodes u and v at distance δ , they are in a tree of diameter $\delta \log n$. Indeed, by modifying the PRR scheme along the lines proposed by Thorup and Zwick [27] one can improve the space bounds by a logarithmic factor, but we do not address this issue here.

The remainder of this paper is divided as follows: Section 2 describes the details of Tapestry, highlighting differences with the PRR scheme and introducing concepts and terminology for the remainder of the paper. Section 3 describes how to solve the incremental nearest neighbor problem. Section 4 explains how this is used as part of inserting a node. Section 5 gives algorithms for deletion. Section 6 discusses the issues in applying these theoretical results to physical networks. Section 7 gives a simple proof that PRR v.0 scheme has polylogarithmic stretch for general metric spaces. Section 8 concludes.

2 The Tapestry Infrastructure

Tapestry [32, 30] is the wide-area location and routing infrastructure of OceanStore [14]. Tapestry assumes that nodes and objects in the system can be identified with unique identifiers (names), represented as strings of digits. Digits are drawn from an alphabet of radix b . Identifiers are uniformly distributed in the namespace. We will refer to node identifiers as *node-IDs* and object identifiers as *globally unique identifiers* (GUIDs). This means that every query has a unique destination GUID which ultimately resolves to a node-ID. For a string of digits α , let $|\alpha|$ represent the number of digits in that string.

Tapestry inherits its basic structure from the data location scheme of Plaxton, Rajaraman, and Richa (PRR) [20]. As with the PRR scheme, each Tapestry node contains pointers to other nodes (*neighbor links*), as well as mappings between object GUIDs and the node-IDs of storage servers (*object pointers*). Queries are routed from node to node along neighbor links until an appropriate object pointer is discovered, at which point the query is forwarded along neighbor links to the destination node.

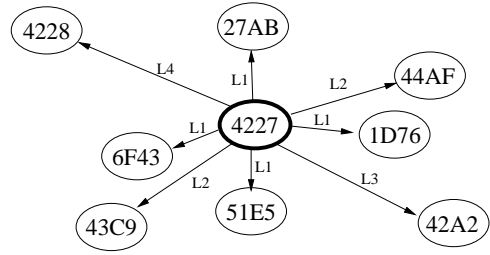


Figure 1: Tapestry Routing Mesh. Each node is linked to other nodes via *neighbor links*, shown as solid arrows with labels. Labels denote which digit is resolved during link traversal. Here, node 4227 has an L1 link to 27AB, resolving the first digit, an L2 link to 44AF, resolving the second digit, etc. Using the notation of Section 2.1, 42A2 is a (42, A) neighbor of 4227.

2.1 The Tapestry Routing Mesh

The Tapestry *routing mesh* is an overlay network between participating nodes. Each Tapestry node contains links to a set of neighbors that share prefixes with its node-ID. Thus, neighbors of node-ID α are restricted to nodes that share prefixes with α , that is, nodes whose node-IDs $\beta \circ \delta$ satisfy $\beta \circ \delta' \equiv \alpha$ for some δ, δ' . Neighbor links are labeled by their *level number*, which is one greater than the number of digits in the shared prefix, or $(|\beta| + 1)$. Figure 1 shows a portion of the routing mesh. For each *forward neighbor pointer* from a node A to a node B , there will be a *backward neighbor pointer* (or “backpointer”) from B to A .

Neighbors for node A are grouped into *neighbor sets*. For each prefix β of A ’s ID and each symbol $j \in [0, b - 1]$, the neighbor set $\mathcal{N}_{\beta,j}^A$ contains Tapestry nodes whose node-IDs share the prefix $\beta \circ j$. We will refer to these as (β, j) neighbors of A or simply (β, j) nodes. For each j and β , the closest node in $\mathcal{N}_{\beta,j}^A$ is called the primary neighbor, and the other neighbors are called secondary neighbors. When context is obvious, we will drop the superscript A . Let $l = |\beta| + 1$. Then, the collection of b sets, $\mathcal{N}_{\beta,j}^A$, form the level- l routing table. There is a routing table at each level, up to the maximum length of node-IDs. Membership in neighbor sets is limited by a constant parameter $R \geq 1$: $|\mathcal{N}_{\beta,j}^A| \leq R$, and of all the nodes that could be in the neighbor set, we choose the closest. Further, $|\mathcal{N}_{\beta,j}^A| < R$ implies $\mathcal{N}_{\beta,j}^A$ contains all (β, j) nodes. This gives us the following:

PROPERTY 1 (CONSISTENCY). *If $\mathcal{N}_{\beta,j}^A = \emptyset$, for any A , then there are no (β, j) nodes in the system. We refer to this as a “hole” in A ’s routing table at level $|\beta| + 1$, digit j .*

Property 1 implies that the routing mesh is fully connected. Messages can route from any node to any other node by resolving the destination node-ID one digit at a time. Let the source node be A_0 and destination node be B , with a node-ID equal to $\beta \equiv j_1 \circ j_2 \dots j_n$. If ϵ is the empty string, then routing proceeds by choosing a succession of nodes: $A_1 \in \mathcal{N}_{\epsilon,j_1}^{A_0}$ (first hop), $A_2 \in \mathcal{N}_{j_1,j_2}^{A_1}$ (second hop), $A_3 \in \mathcal{N}_{j_1 \circ j_2, j_3}^{A_2}$ (third hop), etc. This construction gives us locality, as described in the following property.

PROPERTY 2 (LOCALITY). *In both Tapestry and PRR, each $\mathcal{N}_{\beta,j}^A$ contains the closest (β, j) neighbors as determined by a given metric space. The closest neighbor with prefix $\beta \circ j$ is the primary neighbor, while the remaining ones are secondary neighbors.*

Property 2 yields the important locality behavior of both the Tapestry and PRR schemes. Further, it yields a simple solution to the *static nearest-neighbor problem*: Each node A can find its nearest neighbor by choosing from the set $\bigcup_{j \in [0, b-1]} \mathcal{N}_{\epsilon,j}^A$, where ϵ

⁶Personal communication with James R. Lee, April 8, 2003.

represents the empty string. Section 3 will discuss how to maintain Property 2 in a dynamic network.

2.2 Routing to Objects with Low Stretch

Tapestry maps each object GUID, ψ , to a set of *root nodes*: $\mathcal{R}_\psi = \text{MAPROOTS}(\psi)$. We call \mathcal{R}_ψ the *root set* for ψ , and each $A \in \mathcal{R}_\psi$ is a *root node* for ψ . It is assumed that $\text{MAPROOTS}(\psi)$ can be evaluated anywhere in the network.

To function properly, $\text{MAPROOTS}(\psi)$ must return nodes that exist. The size of a root set, $|\mathcal{R}_\psi| \geq 1$, is small and constant for all objects. In the simplest version of Tapestry, $|\mathcal{R}_\psi| = 1$. In this case, we can speak of *the root node* for a given node ψ . For this to be sensible, we must have the following property:

PROPERTY 3 (UNIQUE ROOT SET). *The root set, \mathcal{R}_ψ , for object ψ must be unique. In particular, $\text{MAPROOTS}(\psi)$ must generate the same \mathcal{R}_ψ , regardless of where it is evaluated in the network.*

Storage servers *publish* the fact that they are storing a replica by routing a publish message toward each $A \in \mathcal{R}_\psi$. Publish messages are routed along primary neighbor links. At each hop, publish messages deposit *object pointers* to the object. Unlike the PRR scheme, Tapestry maintains *all* object pointers for objects with duplicate names (*i.e.*, copies). Figure 2 illustrates publication of two replicas with the same GUID. To provide fault-tolerance, Tapestry assumes that pointers are *soft-state*, that is, pointers expire and objects must be republished (published again) at regular intervals. Republishing may be requested if something changes about the network.

Queries for object ψ route toward *one* of the root nodes $A \in \mathcal{R}_\psi$ along primary neighbor links until they encounter an object pointer for ψ , then route to the located replica. If multiple pointers are encountered, the query proceeds to the closest replica to the current node (*i.e.*, the node where the object pointer is found). At the beginning of the query, we select a root randomly from \mathcal{R}_ψ . Figure 3 shows three different location paths. In the worst case, a location operation involves routing all the way to root. However, if the desired object is close to the client, then the query path will be very likely to intersect the publishing path before reaching the root.

In the PRR scheme, queries route by examining all secondary neighbors before proceeding along the primary link toward the root. The number of secondary neighbors is set according to their metric space, but bounded by a constant. The following theorem shows an important property shared by PRR and Tapestry.

THEOREM 1. *PRR and Tapestry can perform location-independent routing, given Property 3.*

PROOF. The publishing process ensures that all members of \mathcal{R}_ψ contain mappings (object pointers) between ψ and every server which contains ψ . Thus, a query routed toward any $A \in \mathcal{R}_\psi$ will (in the worst case) encounter a pointer for ψ upon reaching A . \square

OBSERVATION 1. (Fault Tolerance) *If $|\mathcal{R}_\psi| > 1$ and the names in \mathcal{R}_ψ are independent of one another, then we can retry object queries and tolerate faults in the Tapestry routing mesh.*

In a general metric space, it is difficult to make claims about the performance of such a system. PRR restrict their attention to metric spaces with a certain even-growth property: they assume that for a given point A , the ratio of the number of points within $2r$ of A and the number of points within distance r of A is bounded above and below by constants. (Unless all points are within $2r$ of A .) Given this constraint, [20] shows the average distance traveled in locating an object is *proportional* to the distance from that object, that is, queries exhibit $O(1)$ stretch. Their data structure, however, ignores many issues important in practical systems. Tapestry is a simplification that is easier to implement and seems to provide low stretch in practice [32, 30].

2.3 Surrogate Routing

The procedures for *publishing* and *querying* documents outlined in Section 2.2 do not require the actual membership of \mathcal{R}_ψ to be known. All that is required is to be able to compute the next hop toward the root from a given position in the network. As long as this *incremental* version of $\text{MAPROOTS}()$ is consistent in its behavior, we achieve the same routing and locality behavior as in Section 2.2. Assume that $\text{INCRMAPROOTS}(\psi, B)$ produces an ordered list of the next hop toward the roots of ψ from node B .

In the PRR scheme, $\text{MAPROOTS}(\psi)$ produces a single root node A which matches in the largest possible number of prefix bits with ψ . Ties are broken by consulting a global order of nodes. The PRR scheme specifies a corresponding $\text{INCRMAPROOTS}()$ function as follows: the neighbor sets, \mathcal{N} , are supplemented with additional *root links* that fill holes in the routing table. To route a message toward the root node, PRR routes directly to ψ as if it were a node in the Tapestry mesh. Assuming that the supplemental root links are consistent with one another, every publish or query for document ψ will head toward the same root node.

We call this process *surrogate routing*, since it involves routing toward ψ as if it were a node, then adapting when the process fails. Roots reached in this way are considered *surrogate roots* of ψ .

Localized Routing Decisions: In a dynamic network, maintenance of routing pointers can be problematic. In Tapestry, we have chosen to focus on Property 1 as our primary consistency constraint. Thus, in contrast to the original PRR scheme, we do not maintain extra route links to aid in locating root nodes. Instead, all routing decisions are made based on the current routing table, the source and destination GUIDs, and information collected along the route by the query (*e.g.*, the number of digits resolved so far).

We highlight two variants on localized routing; others are certainly possible. Both of them proceed by routing one digit at a time toward the destination GUID, that is, each network hop resolves one additional digit toward the destination. Since there is no backtracking, these schemes are guaranteed to complete.

- **Tapestry Native Routing:** We route one digit at a time. When there is no match for the next digit, we route to the next filled entry in the same level of the table, wrapping around if needed. For example, if the next digit to be routed is a 3, and there is no entry, try 4, then 5, and so on.
- **Distributed PRR-like Routing:** We route one digit at a time as follows:
 1. *Before first hole:* Route one digit at a time as above.
 2. *At first hole:* Route along an existing neighbor link that matches the desired digit in as many significant bits as possible. If there is more than one such route, pick the route with the numerically higher digit.
 3. *After first hole:* Always pick a routing-table entry with the numerically highest available digit.

This technique routes to the root node with the numerically largest node-ID that matches the destination GUID in the most significant bits.

For both of these schemes, routing stops if the current node is the only node left at and above the current level in the routing table; the resulting node is the root node. Such localized schemes are simpler than PRR under inserts and deletes. In addition, the Tapestry Native Routing scheme may have better load balancing properties.

THEOREM 2. *Suppose Property 1 holds. Then the Tapestry version of surrogate routing will produce a unique root.*

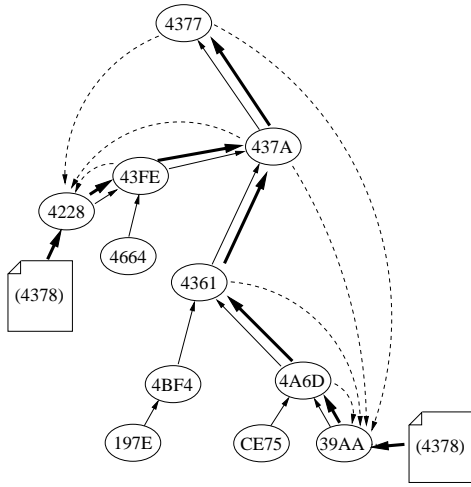


Figure 2: Publication in Tapestry. To publish object 4378, server 39AA sends publication request towards root, leaving a pointer at each hop. Server 4228 publishes its replica similarly. Since no 4378 node exists, object 4378 is rooted at node 4377.

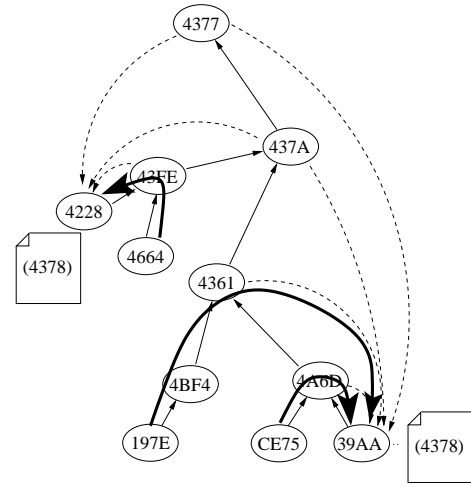


Figure 3: Routing in Tapestry: Three different location requests. For instance, to locate GUID 4378, query source 197E routes towards the root, checking for a pointer at each step. At node 4361, it encounters a pointer to server 39AA.

PROOF. Proof by contradiction. Suppose that messages for an object with ID X end routing at two different nodes, A and B . Let β be the longest common prefix of A and B , and let i be the length of β . Then, let A' and B' be the nodes that do the $(i+1)$ st routing step; that is, the two nodes that send the message to different digits. Notice that after this step, the first $(i+1)$ digits of the prefix remain constant in all further routing steps. Both $\mathcal{N}_{\beta,*}^{A'}$ and $\mathcal{N}_{\beta,*}^{B'}$ must have the same pattern of empty and non-empty entries. That is, if $\mathcal{N}_{\beta,j}^{A'}$ is empty, then $\mathcal{N}_{\beta,j}^{B'}$ must also be empty, by Property 1. So both A' and B' send the message on a node with the same $(i+1)$ st digit, a contradiction. \square

A similar proof is possible for the distributed PRR-like scheme. Localized routing may introduce additional hops over PRR; however, the number of additional hops is independent of n and in expectation is less than 2 [32]. Notice the following:

OBSERVATION 2. (Multiple Roots) Surrogate routing generalizes to multiple roots. First, a pseudo-random function is employed to map the initial document GUID ψ into a set of identifiers $\psi_0, \psi_1, \dots, \psi_n$. Then, to route to root i , we surrogate route to ψ_i .

2.4 A Comparison between Tapestry and PRR

In this section, we quickly outline the differences between Tapestry and PRR. With the exception of some changes to the maintenance of object pointers, the key ways that Tapestry differs from PRR revolve around eliminating the requirements of a static network and adding provisions to handle faults in a graceful manner, both of which are required for a system deployable on real networks. To further clarify the contributions of this work, a few of the key differences are noted here.

First, Tapestry nodes keep pointers to all copies of a given object. In PRR, a node stores at most one pointer to each object, regardless of how many copies are in the network. As a result, deleting an object is easier in Tapestry than in PRR, but comes with an increase in storage. Further, this property allows Tapestry applications to exploit multiple object replicas; queries can be multicast to every object of a given name or to a “close” subset of objects.

Second, PRR makes use of secondary neighbors in its core object location algorithms, while Tapestry uses them primarily for fault-resilience. When searching for an object, PRR searches on

the primary and secondary neighbors before taking an additional hop towards the object root. It is worth noting that this is equivalent to publishing on all the secondary neighbors, but only searching the primary neighbors. In contrast, Tapestry keeps for every entry, a small number (in the current implementation, two) of additional backup neighbor links for fault-resilience. All publishing and search takes place only on the primary links.

Third, as described above, surrogate routing is different in the two systems. First, PRR assumes a static set of nodes with precomputed surrogates, while Tapestry’s algorithms maintain the surrogate routes as part of the insertion process. PRR can be extended, however, with our dynamic algorithms while maintaining the same surrogate routing scheme. Tapestry surrogate routing does slightly better at load balancing of objects across the surrogate roots.

Finally, Tapestry is implemented and running, and downloads are available.

3 Building Neighbor Tables

Building the neighbor table is perhaps the most complex and interesting part of the insertion process, so we present it first. The problem is to build the neighbor sets, $\mathcal{N}_{\beta,j}^A$ for a new node A . These sets must satisfy Properties 1 and 2. This can be seen as solving the nearest neighbor problem for many different prefixes. One solution is to simply use the method of Karger and Ruhl [13] many times, once for each prefix. This would essentially require each node to participate in $O(\log n)$ Karger-Ruhl data structures, one for each level of the neighbor table.

The method we present below has lower network distance than a straightforward use of Karger and Ruhl (although the same number of network hops) and incurs no additional space over the PRR data structures.

As in [20], we adopt the following network constraint. Let $\mathcal{B}_A(r)$ denote the ball of radius r around A ; i.e., all points within distance r of A , and $|\mathcal{B}_A(r)|$ denote the number of such points. We assume:

$$|\mathcal{B}_A(2r)| \leq c |\mathcal{B}_A(r)|, \quad (1)$$

for some constant c . PRR also assume that $|\mathcal{B}_A(2r)| \geq c' |\mathcal{B}_A(r)|$, but that assumption is not needed for our extensions. Notice that our expansion property is almost exactly that used by Karger and Ruhl [13]. We also assume the triangle inequality in network dis-

```

method ACQUIRENEIGHBORTABLE (NewNodeName, NewNodeIP, PSurrogateName, PSurrogateIP)
1  $\alpha \leftarrow \text{GREATESTCOMMONPREFIX}(\textit{NewNodeName}, \textit{PSurrogateName})$ 
2  $\textit{maxLevel} \leftarrow \text{LENGTH}(\alpha)$ 
3 list  $\leftarrow \text{ACKNOWLEDGEDMULTICAST}[\text{on } \textit{PSurrogateIP}] (\alpha, \text{SENDID}(\textit{NewNodeIP}, \textit{NewNodeName}))$ 
4 BUILDTABLEFROMLIST(list,  $\textit{maxLevel}$ )
5 for  $i = \textit{maxLevel} - 1$  to 0
6   list  $\leftarrow \text{GETNEXTLIST}(\textit{list}, i, \textit{NewNodeName}, \textit{NewNodeIP})$ 
7   BUILDTABLEFROMLIST(list,  $i$ )
end ACQUIRENEIGHBORTABLE

method GETNEXTLIST (neighborlist,  $\textit{level}$ , NewNodeName, NewNodeIP)
1 nextList  $\leftarrow \emptyset$ 
2 for  $n \in \textit{neighborlist}$ 
3   temp  $\leftarrow \text{GETFORWARDANDBACKPOINTERS}(n, \textit{level})$ 
4   ADDTOTABLEIFCLOSER [on  $n$ ] (NewNodeName, NewNodeIP)
5   nextList  $\leftarrow \text{KEEPCLOSESTK}(\textit{temp} \cup \textit{nextList})$ 
6 return nextList
end GETNEXTLIST

```

Figure 4: Building a Neighbor Table. A few words on notation: FUNCTION [**on destination**] represents a call to run FUNCTION **on destination**, variables in *italics* are single-valued, and variables in **bold** are vectors. The Acknowledged-Multicast function is described in Figure 8.

tance, that is

$$d(X, Y) \leq d(X, Z) + d(Z, Y)$$

for any set of nodes X, Y , and Z . Our bounds in terms of network latency or network hops and ignore local computation in our calculations. None of the local computation is time-consuming, so this is a fair measure of complexity.

3.1 The Algorithm

Figure 4 shows how to build neighbor tables. In words, suppose that the longest common prefix of the new node and any other node in the network is α . Then we begin with the list of all nodes with prefix α . (We explain how to get this list in the next section.) We proceed by getting similar lists for progressively smaller prefixes, until we have the closest k nodes matching the empty prefix.

Let a level- i node be a node that shares a length i prefix with α . Then, to go from the level- $(i+1)$ list to the level- i list, we ask each node on the level- $(i+1)$ list to give us all the level- i nodes they know of (we ask for both forward and backwards pointers). Note that each level- i node must have at least one level- $(i+1)$ node in its neighbor table, so following the backpointers of all level- $(i+1)$ nodes gives us all level- i nodes. We then contact these nodes, and sort them according to their distance from the inserting node. Each node contacted this way also checks to see if the new node should be added to its own table (line 4). We then trim this list, keeping only the closest k nodes. If $b > c^2$, then Lemma 1 says there is some $k = O(\log n)$ such that with high probability, the lists at each level contain exactly the k closest nodes.

We then use these lists to fill in the neighbor table. This happens in line 7 of ACQUIRENEIGHBORTABLE. More precisely, recall that level i of the table consists of nodes with the prefix $\alpha_{i-1} \circ j$, where α_{i-1} is the first $(i-1)$ digits of the node's prefix. To fill in level i of the neighbor table, we look in the level- $(i-1)$ list. For $j \in [0, b-1]$, we keep the closest R (α_{i-1}, j) nodes (R is defined in Section 2.1).⁷

⁷We note that while the algorithm presented here is here is sensitive to failures, a slight modification can make the algorithm substantially more robust.

3.2 A Proof of Correctness

Theorems 3 and 4 prove that with high probability, the above algorithm correctly creates the new node's neighbor table and correctly updates the neighbor tables of the existing nodes. Theorem 3 uses Lemmas 1 and 2 to show that the new node's table gets built correctly, and Theorem 4 argues that the other nodes' tables are updated correctly.

The following lemma shows that if GETNEXTLIST is given the k closest level- $(i+1)$ nodes, it finds the k closest level- i nodes.

LEMMA 1. *If c is the expansion constant of the network, and $c^2 < b$, then given a list of the closest k level- $(i+1)$ nodes, we can find the k closest level- i nodes, for $k = O(\log n)$. In particular, if $k \geq \frac{24(a+1)b \log n}{(1-c^2/b)^2}$, the failure probability is bounded by $1/n^a$.*

PROOF. Let δ_i be the radius of the smallest ball around the new node containing k level- i matches. We would like to show that any node A inside the ball must point to a level- $(i+1)$ node within δ_{i+1} of the new node. If that is the case, then we will query A 's parent, and so find A itself.

For the rest of the proof to work, we need that at least one of the k level- i nodes is also a level- $(i+1)$ node. The probability this is not true is $(1 - 1/b)^k \leq \exp(-k/b) \leq \exp(-(a+1) \log n) \leq 1/(n^{a+1})$. For the remainder of the proof, we assume at least one of the k level- i nodes is also an level- $(i+1)$ node. Then the distance between A and its nearest level- $(i+1)$ node is no more than $2\delta_i$, since both A and the level- $(i+1)$ node are within the ball of radius δ_i . By the triangle inequality, the distance between the new node and A 's parent is no more than $2\delta_i + \delta_i = 3\delta_i$. (See Figure 5.) This means that as long as $3\delta_i < \delta_{i+1}$, A must point to a node inside δ_{i+1} . Since we query all level- $(i+1)$ in δ_{i+1} , this means we will query A 's parent, and so find A .

To complete the proof, we need $3\delta_i < \delta_{i+1}$ with high probability. This is the expected behavior; given a ball with k level- i nodes, doubling the radius twice gets no more than $c^2 k$ nodes, and so, no more than $k(c^2/b)$ level- $(i+1)$ nodes. Since $c^2/b < 1$, this means that the quadrupled ball has less than k level- $(i+1)$ nodes, or equivalently, the ball containing k level- $(i+1)$ nodes is at least three (really, four) times the size of the ball with k level- i nodes. The following turns this informal argument into a proof.

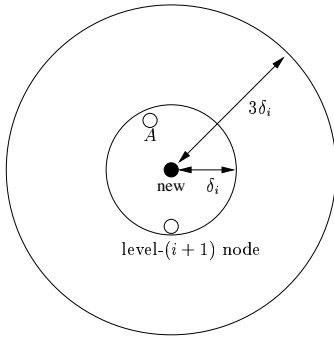


Figure 5: Figure for Theorem 3. If $3\delta_i$ is less than δ_{i+1} , then A must point to a node within d_{i+1} .

First, recall that $c^2/b < 1$. Pick λ' and λ as follows:

$$\begin{aligned}\lambda &= \frac{1}{2}(1 - c^2/b) < \frac{1}{2} \\ \lambda' &= \lambda(2 - c^2/b) < 1\end{aligned}$$

Note that $\lambda' > \lambda$ and $(1 - \lambda')b/c^2 = b/c^2(1 - 2\lambda) + \lambda = 1 + \lambda$. Notice that we can write k as a function of λ ; in particular $k = \frac{6(a+1)b \log n}{\lambda^2}$.

Now, let l be $(1 - \lambda')^{-1}kb^i$. This is the required number of nodes such that one expects $(1 - \lambda')^{-1}k$ of the nodes to be level- i nodes.

Let L_{real} be the random variable representing total volume of the ball (*i.e.*, the number of nodes in the ball) containing k level- i nodes. If $3\delta_i < \delta_{i+1}$, then we are done, so in the rest of the proof, we argue that the probability that $3\delta_i \geq \delta_{i+1}$ is small.

We use the fact that $\Pr[3\delta_i \geq \delta_{i+1}]$ is the same as

$$\begin{aligned}\Pr[3\delta_i \geq \delta_{i+1} | L_{\text{real}} > l] \cdot \Pr[L_{\text{real}} > l] \\ + \Pr[3\delta_i \geq \delta_{i+1} | L_{\text{real}} \leq l] \cdot \Pr[L_{\text{real}} \leq l].\end{aligned}$$

We will bound one term from each product.

We show $\Pr[L_{\text{real}} > l] \leq 1/n^{a+1}$. Let X_m be a random variable representing the number of level- i nodes in m nodes. Notice that the $\Pr[L_{\text{real}} > l]$ is bounded from above by $\Pr[X_l < k]$, since if $L_{\text{real}} > l$, then it must be that the closest l nodes to the new node do not contain k level- i nodes.

But

$$\Pr[X_l < k] = \Pr[X_l < (1 - \lambda')E[X_l]].$$

Using a Chernoff bound, this is less than

$$\exp(-\lambda'^2 E[X_l]/2) \leq \exp(-\lambda'^2 k/2) \leq \exp(-\lambda^2 k/2).$$

Substituting for k , this becomes

$$\exp\left(-\frac{6(a+1)b \log n \lambda^2}{2\lambda^2}\right) \leq \exp(-3(a+1) \log n) \leq 1/n^{a+1}.$$

We show $\Pr[3\delta_i \geq \delta_{i+1} | L_{\text{real}} \leq l] \leq 2/n^{a+1}$. Consider the ball of radius $3\delta_i$ around the new node. If this ball contains k level- $(i+1)$ nodes (δ_{i+1} is smaller than $3\delta_i$), then the ball of radius $4\delta_i$ must also contain at least k level- $(i+1)$ nodes.

However, we know the volume (that is, the number of nodes) of this ball is less than $c^2 l$ by Equation 1 and the fact $L_{\text{real}} \leq l$. Let Y_m be the number of $i+1$ nodes in m trials. Then, rewriting our goal with this notation, we wish to bound $\Pr[Y_{c^2 l} \geq k | X_l \geq k]$. (Notice that $Y_{c^2 l}$ is not independent of X_l .) We can write that the $\Pr[A | B] = \frac{\Pr[A \cap B]}{\Pr[B]} \leq \frac{\Pr[A]}{\Pr[B]}$, so it

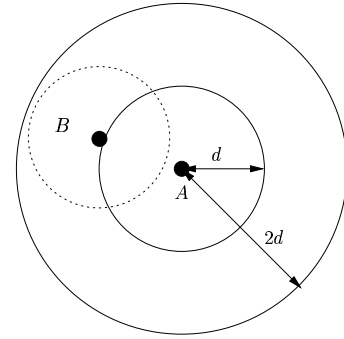


Figure 6: Figure for Theorem 4. The larger ball around A contains $O(\log n)$ nodes, while the smaller ball contains none.

suffices to bound $\frac{\Pr[Y_{c^2 l} \geq k]}{\Pr[X_l \geq k]}$. We already bounded the denominator, so next we wish to bound $\Pr[Y_{c^2 l} \geq k]$.

Since

$$E[Y_{c^2 l}] = \frac{c^2}{b} E[X_l] = \frac{kc^2}{b(1 - \lambda')} = \frac{k}{1 + \lambda'}.$$

we get that

$$\Pr[Y_{c^2 l} \geq k] = \Pr[Y_{c^2 l} \geq (1 + \lambda)E[Y_{c^2 l}]] \leq \exp(-\lambda^2 E[Y_{c^2 l}]/3).$$

But $E[Y_{c^2 l}] = k/(1 + \lambda) \geq k/2$, so $\Pr[Y_{c^2 l} \geq k] \leq \exp(-\lambda^2 k/6)$. Substituting k , we get that this is bounded by $\exp(-(a+1) \log n) = 1/n^{a+1}$.

$$\text{So } \Pr[3\delta_i \geq \delta_{i+1} | L_{\text{real}} \leq l] \leq \frac{\Pr[Y_{c^2 l} \geq k]}{\Pr[X_l \geq k]} \leq \frac{1/n^{a+1}}{1 - 1/n^{a+1}} \leq \frac{2}{n^{a+1}}. \text{ (So long as } n^{a+1} > 2)$$

Recall that we wish to bound $\Pr[3\delta_i \geq \delta_{i+1}]$, and we know that

$$\begin{aligned}\Pr[3\delta_i \geq \delta_{i+1}] &= \\ &\Pr[3\delta_i \geq \delta_{i+1} | L_{\text{real}} > l] \Pr[L_{\text{real}} > l] \\ &+ \Pr[3\delta_i \geq \delta_{i+1} | L_{\text{real}} \leq l] \Pr[L_{\text{real}} \leq l] \\ &\leq 1 \cdot 1/n^{a+1} + 2/n^{a+1} \cdot 1 \\ &\leq 1/n^a\end{aligned}$$

where the last step follows so long as $n \geq 3$. \square

Next, we show that given the k closest nodes matching in i digits, we can fill in level $(i+1)$ of the neighbor table if k is large enough. This means that given a list containing the k closest nodes with prefix β , for $k = O(\log n)$, there are at least $R(\beta, j)$ nodes for each $j \in [0, b-1]$. (Recall that R was defined in Section 2.1.)

LEMMA 2. For $k = O(\log n)$ and $R = o(\log n)$ the list of the closest k β nodes to a given node A contains $\bigcup_j \mathcal{N}_{\beta, j}^A$ with high probability. In particular, for $k \geq 8(a+1)b \log n$, the probability it does not is less than $1/n^a$.

PROOF. For any given j , let X_j be a random variable representing the number of (β, j) nodes in the list. In expectation, this is $k/b = 8(a+1) \log n$. We want to bound $\Pr[X_j \leq R]$, and since $R = o(\log n)$, we have $R \leq \frac{1}{2}k/b$.

Thus,

$$\Pr[X_j \leq R] \leq \Pr[X_j \leq \frac{1}{2}E[X_j]] \leq \exp\left(-\left(\frac{1}{2}\right)^2 \frac{E[X_j]}{2}\right).$$

The last step uses a Chernoff bound. We can then simplify the last equation and say that $\Pr[X_j \leq R] \leq 1/n^{a+1}$.

Now, we apply a union bound over all the b possible j to get that the probability that any j has less than R nodes in the list is bounded by b/n^{a+1} , and since we can assume $b < n$, this gives us a bound of $1/n^a$, which is the desired result. \square

Finally, we can combine these two lemmas to prove the following theorem:

THEOREM 3. *If c is the expansion constant of the network, $b > c^2$ (where b is the digit size), and $R = o(\log n)$, then there is a $k = O(\log n)$ for which the algorithm of Figure 4 will produce the new node's correct neighbor table with probability $1/n^a$ for any constant a .*

PROOF. By Lemma 1, there is a k_1 such that the probability that the i th list is incorrectly generated from the $(i+1)$ st list is less than $1/2n^{a+1}$. Since there are $\log n$ levels, the probability that any level fails is less than $\frac{\log n}{2n^{a+1}} \leq 1/(2n^a)$.

Second, by Lemma 2, for some $k_2 = O(\log n)$, the probability that we are unable to fill the neighbor table with $R = o(\log n)$ neighbors from lists of length k_2 is less than $1/(2n^{a+1})$. Since there are $\log n$ levels to fill, the probability that any of the levels is left unfilled is bounded by $\frac{\log n}{2n^{a+1}} \leq 1/2n^a$.

If we choose $k = \max(k_1, k_2)$, the probability that either of the lists are not correct or the table cannot be filled is bounded by $1/(2n^a) + 1/(2n^a) = 1/n^a$. This proves the theorem. \square

The new node also causes changes to the neighbor tables of other nodes. We instruct any node that is a candidate for the new node's table to check if adding the new node could improve its own neighbor table. This happens in line 4 of Figure 4. It remains to show that with high probability, line 4 of Figure 4 updates all nodes that need to be updated. In particular, we show that there is a $k = O(\log n)$ such that any node that needs to update its level- i link is one of the closest k nodes of level- i with high probability.

THEOREM 4. *If a new node B is a (α, j) neighbor of A (so B is one of the R closest nodes to A with prefix $\alpha \circ j$), then with high probability, A is among the $k = O(\log n)$ closest α -nodes to B . In particular, for $k = 16abc \log n$, and $R = o(\log n)$, the probability A is not among the closest k nodes is $1/n^a$.*

PROOF. We will show that the probability A is not among the k closest α -nodes to B can be made arbitrarily small. Let $d = d(A, B)$ or the distance between A and B . Consider the ball around A of radius d . (Shown in Figure 6). Since B is in the neighbor table of A , there are less than R (α, j) nodes in this ball. Further, notice that the ball around B containing k α -nodes does not contain A (or else the proof is done), so its radius must be less than d . Finally, consider the ball around A of radius $2d$. It completely contains the ball around B .

If A is not among the closest k nodes, then the ball around A of radius d contains no more than R (α, j) nodes, while the ball around B of radius d contains k nodes of prefix α . We will show the probability of this is very small. Since $R = o(\log n)$, for sufficiently large n , we can assume that $R \leq 2a \log n$.

Instead of arguing directly about the ball around B , we will argue about the ball of radius $2d$ around A , since $\mathcal{B}_A(2d) \supset \mathcal{B}_B(d)$. More precisely, if we let $|\mathcal{B}_A(d)|_\alpha$ denote the number of α nodes in the ball of radius d around A , then we want to argue that the probability that $|\mathcal{B}_A(d)|_{(\alpha, j)} \leq R$ and $|\mathcal{B}_A(2d)|_\alpha \geq k$ is small. To do this, we have two cases, depending on $|\mathcal{B}_A(d)|$. Let l_{real} be the number of nodes in the smaller ball around A (or $|\mathcal{B}_A(d)|$) and let $l = 8(a/p) \log n$ where p is the probability a node is a (α, j) -node. (Note that l_{real} is not a random variable.) Finally, let $k = 16abc \log n$. (Recall that b is the base of the logarithm and c is the expansion constant of the network.)

Case 1: $l_{\text{real}} > l$. Intuitively, this case is unlikely because if the smaller ball around A has many nodes, we expect there to be many (α, j) nodes.

More formally, let X_m be the number of (α, j) nodes found in m trials. Then $X_{l_{\text{real}}}$ is the random variable representing the number of (α, j) nodes found in the closest l_{real} nodes to A . In this case, we want to bound the probability that $X_{l_{\text{real}}} \leq R$. Then we can say that $\Pr[X_{l_{\text{real}}} \leq R] \leq \Pr[X_l \leq R]$. Further $E[X_l] = p(8a/p) \log n = 8a \log n$, so $\Pr[X_l \leq R] \leq \Pr[X_l \leq (1 - \frac{1}{2})E[X_l]]$. Using a Chernoff bound, this is less than $\exp(-\frac{1}{2}E[X_l]/2)$, and substituting, this is less than $\exp(-\frac{1}{2}8a \log n/2) \leq \exp(-a \log n) = 1/n^a$.

Case 2: $l_{\text{real}} \leq l$. In this case, we argue that the ball of radius $2d$ around A contains more than k nodes with probability less than $1/n^a$. But since the ball around B of radius d is contained in this ball, this will imply that the ball around B of radius d contains more than k nodes with probability less than $1/n^a$.

Let Y_m a random variable representing the number of α nodes in m trials. We wish to bound

$$\Pr[|\mathcal{B}_A(2d)|_\alpha \geq k],$$

but this is less than or equal to $\Pr[Y_{cl} \geq k]$, since $\mathcal{B}_A(2d)$ contains at most cl nodes. Then

$$\Pr[|\mathcal{B}_A(2d)|_\alpha \geq k] \leq \Pr[Y_{cl_{\text{real}}} \geq k] \leq \Pr[Y_{cl} \geq k].$$

Recalling that $E[Y_{cl}] = (pb)c(8a/p) \log n$, $k = 2E[Y_{cl}]$, so again using a Chernoff bound, we can write

$$\Pr[Y_{cl} \geq 2E[Y_{cl}]] \leq \exp(-\frac{8}{3}abc \log n) \leq \exp(-a \log n).$$

Since the probability of each case is bounded by $1/n^a$, the overall probability is also bounded by $1/n^a$, completing the proof. \square

To make sure the probability of failing to get either the new node's table or correctly update the tables of established nodes is less than $1/n^a$ for some a , we combine the results in the following way. Let k_1 be large enough that the probability a mistake is made in building the neighbor table is less than $1/(2n^a)$ (this is possible by Theorem 3), and choose k_2 large enough that the probability that the algorithm misses an update to another node's table is less than $1/(2n^a)$ (possible by Theorem 4). Finally, choose $k = \max(k_1, k_2)$, and the probability that the algorithm of Figure 4 fails to perform the correct updates will be less than $1/(2n^a) + 1/(2n^a) \leq 1/n^a$.

3.3 Running Time

Since each node has an expected constant number of pointers per level, the expected time of this algorithm is $O(k) = O(\log n)$ per level or $O(\log^2 n)$ overall. (We are concerned with network traffic and distance and hence ignore the cost of local computation.)

The number of backpointers is less than $O(\log n)$ per level per node with high probability, so we get a total time of $O(\log^3 n)$ with high probability. But this analysis can be tightened. Using the techniques of Theorems 3 and Theorem 4, one can argue that with high probability, all the visited level- i nodes are within a ball of radius $4\delta_{i+1}$. Further, again with high probability, there are only $O(\log n)$ level- i nodes within $4\delta_{i+1}$. This means we visit only $O(\log n)$ nodes per level, or $O(\log^2 n)$ nodes overall.

Further, notice that $\delta_i \leq \frac{1}{3}\delta_{i+1}$. Suppose the number of nodes touched at each level is bounded by q . We know (by the above) that $q = O(\log n)$. The total network latency is bounded by:

$$\sum_i \delta_i q = q \sum_i \delta_i$$

Since the δ_i are geometrically decreasing, they sum to $O(d)$, where d is the network diameter, so the total latency for building neighbor tables is $O(qd) = O(d \log n)$.

4 Node Insertion

Next, we describe the overall insertion algorithm, using the nearest neighbor algorithm as a subroutine. We would like the results of the insertion to be the same as if we had been able to build the network from static data. This means in addition to correctly updating the neighbor tables, maintaining the following invariant:

PROPERTY 4. *If node A is on the path between a publisher of object O and the root of object O , then A has a pointer to O .*

In this section, we show that if Property 1, Property 2, and Property 4 hold, then we can insert a node such that all three hold after the insertion, and the new node is part of the network. It may, however, happen that during a node insertion one or both of the properties is temporarily untrue. In the case of Property 1, this can be particularly serious since some objects may become temporarily unavailable. Section 4.3 shows how the algorithm can be extended to eliminate this problem.

Figure 7 shows the basic insertion algorithm. First, the new node contacts its surrogate; that is, the node with the ID closest to its own. Then it gets a copy of the surrogate's neighbor table. These first two steps could be combined, if desired. Next, the node contacts the subset of nodes that must be notified to maintain Property 1. These are the nodes that have a hole in their neighbor table that the new node should fill. We use the function `ACKNOWLEDGEDMULTICAST` (detailed in Section 4.1) to do this. As a final step, we build the neighbor tables, as described in Section 3. To reduce the number of multicasts, we can use the multicast in step 4 of the insertion algorithm to get the first **list** of the nearest neighbor algorithm. Finally, notice that once the multicast is finished, the node is fully functional, though its neighbor table may be far from optimal.

We would also like to maintain Property 4. This means that *all* nodes on the path from an object's server to the object's root have a pointer to that object. Once again, there are two failure cases, one of correctness, where not fixing the problem means that the network may fail to locate an object, and one of performance, where not fixing the problem may increase object location latency.

The function `LINKANDXFERROOT` from Figure 7 takes care of correctness by transferring object pointers that should be rooted at the new node and deleting pointers that should no longer be on the current node. If we do not move the object pointers, then objects may become unreachable. Performance optimization involves redistributing pointers and will be discussed in Section 4.2.

4.1 Acknowledged Multicast

To contact all nodes with a given prefix we introduce an algorithm called *Acknowledged Multicast*, shown in Figure 8. This algorithm is initiated by the arrival of a multicast message at some node.

A multicast message consists of a prefix α and a function to apply. To be a valid multicast message, the prefix α must be a prefix of the receiving node. When a node receives a multicast message for prefix α , it sends the message to one node with each possible extension of α ; that is, for each j , it sends the message to one (α, j) node if such a node exists. One of these extensions will be the node itself, so a node may receive multicast messages from itself at potentially many different levels. We know by Property 1 that if an (α, j) node exists, then every α -node knows at least one such node. Each of these nodes then continues the multicast. When a node cannot forward the message further, it applies the function.

Because we need to know when the algorithm is finished, we require each recipient to send an acknowledgment to its parent after receiving acknowledgments from its children. If a node has no children, it sends the acknowledgment immediately. When the initiating node gets an acknowledgment from each of its children, we know that all nodes with the given prefix have been contacted.

THEOREM 5. *When a multicast recipient with prefix α sends acknowledgment, all the nodes with prefix α have been reached.*

PROOF. This is a proof by induction on the length of α . In the base case, suppose node A receives a multicast message for prefix α and A is the only node with prefix α . The claim is trivially true.

Now, assume the claim holds for a prefix α of length i . We will prove it holds for a prefix α of length $i+1$. Suppose node A receives a multicast for a prefix of length $i+1$. Then A forwards the multicast to one node with each possible one-digit extension of α (i.e., $\alpha \circ j$ for all $j \in [0, b-1]$). Once A receives all those acknowledgments, all nodes with prefix α have been reached. Since A waits for these acknowledgments before sending its own, all nodes of prefix α have been reached when A sends its acknowledgment. \square

These messages form a tree. If you collapse the messages sent by a node to itself, the result is in fact a spanning tree. This means that if there are k nodes reached in the multicast, there are $k-1$ edges in the tree. Alternatively, each node will only receive one multicast message, so there are no more than $O(k)$ such messages sent. Each of those links could be the diameter of the network, so the total cost of a multicast to k nodes is $O(dk)$. Note that there is a variant of this algorithm that does not require maintaining state at all the participating nodes, but this is beyond the scope of this paper.

4.2 Redistributing Object Pointers

Recall that objects publish their location by placing pointers to themselves along the path from the server to the root. From time to time, we re-establish these pointers in an operation called republish. This section describes a special version of republish that maintains Property 4. This function is used to rearrange the object pointers any time the routing mesh changes the expected path to the root node for some object (e.g., when a node's primary neighbor is replaced by a closer node). This adjustment is not necessary for correctness, but does improve performance of object location.

If the node uses an ordinary republish (simply sending the message towards the root), it could leave object pointers dangling until the next timeout. For example, if the disappearance of node A changes the path from an object to its root node so that the path skips node B , then node B will still be left with a pointer to the object. Further, the simple republish may do extra work updating pointers that have not changed.

Instead, a node with a new forward route sends the object pointer up the new path. The new path and the old path will converge at some node, where a delete message is sent back down the old path, removing outdated pointers. This requires maintaining a last-hop pointer for each object pointer. Figure 9 shows the two methods that are needed to implement this procedure.

Notice, however, that Property 4 is not critical to the functioning of the system. If a node should use `OPTIMIZEOBJECTPTRS` but does not, then performance may suffer, but objects will still be available. Further, timeouts and regular republishes will eventually ensure that the object pointers are on the correct nodes.

4.3 Keeping Objects Available

While a node is inserting itself, object requests that would go to the new node after insertion may either go to the new node or to a pre-insertion destination. Figure 10 shows how to keep objects

```

method INSERT (gatewayIP, NewNodeIP, NewNodeName)
1 (PSurrogateIP, PSurrogateName)  $\leftarrow$  ACQUIREPRIMARYSURROGATE (gatewayIP, NewNodeName)
2  $\alpha \leftarrow$  GREATESTCOMMONPREFIX(NewNodeName, PSurrogateName)
3 GETPRELIMNEIGHBORTABLE [on PSurrogateIP] ()
4 ACKNOWLEDGEDMULTICAST [on PSurrogateIP] ( $\alpha$ , LINKANDXFERROOT[NewNodeIP, NewNodeName])
5 ACQUIRENEIGHBORTABLE (NewNodeName, NewNodeIP, PSurrogateIP, PSurrogateIP)
end INSERT

```

Figure 7: Node Insertion Routine. The insertion process begins by contacting a gateway node, which is a member of the Tapestry network. It then transfers object pointers and optimizes the neighbor table.

```

method ACKNOWLEDGEDMULTICAST( $\alpha$ , FUNCTION)
1 if NOTONLYNODEWITHPREFIX( $\alpha$ )
2   for  $i = 0$  to  $b - 1$ 
3     neighbor  $\leftarrow$  GETMATCHINGNEIGHBOR( $\alpha \circ i$ )
4     if neighbor exists
5        $\mathcal{S} \leftarrow$  ACKNOWLEDGEDMULTICAST [on GETIP(neighbor)] ( $\alpha \circ i$ , FUNCTION)
6   else
7     apply FUNCTION
8   wait  $\mathcal{S}$ 
9 SENDACKNOWLEDGEMENT()
end ACKNOWLEDGEDMULTICAST

```

Figure 8: Acknowledged Multicast. It runs FUNCTION on all nodes with prefix α .

available during this process: if either node receives a request for an object it does not have, it forwards the request to the other node.

If an inserting node receives a request for an object it does not have, it sends the request back out, routing as if it did not know about itself. That is, if the new node fills a hole at level i , it sends out a message at level- i to one of the surrogate nodes. The surrogate then routes the message as it would have if the new node had not yet entered the network.

If a pre-insertion root receives a request for an object pointer that has already been moved to the new node, it should forward the request to the new node. But we want to do this in such a way that the surrogate does not need to keep state to show which nodes are inserting. So we require all nodes to “check the routing” of an object request or publish before rejecting it: the nodes test whether the object made a surrogate step that it did not need to make. If it finds out it did make a surrogate step instead of going to the new node, the old root node redirects the message to the new node.

To make this work properly, we require that the old root not delete pointers until the new root has acknowledged receiving them. If this is done, then one of the two nodes is guaranteed to have the pointer. No matter which node receives the request, before or after the transfer of pointers, the node servicing the request either has the information to satisfy the query or else it forwards the query to the other node, which can satisfy it using local information.

As a final point, it is possible for a request for a non-existent object to loop until the insertion is complete. We can get around this by including information in the message header about where the request has been. This allows the system to detect and prevent loops, and since the number of hops is small, this is not an unreasonable overhead.

4.4 Simultaneous Insertion

In a wide-area network, insertions may not happen one at a time. If two nodes are inserted at once, each may get an older view of the network, so neither node will see the other. Suppose A and B are inserted simultaneously. There are three possibilities:

- A ’s and B ’s insertions do not intersect. This is the most likely case; A need only know about $O(\log^2 n)$ nodes with high

probability so the chance that B is one of them is small.

- For some (α, j) , B should be one of the (α, j) neighbors of A , but A has some more distant (α, j) neighbor instead.
- For some (α, j) , B is the only possible neighbor.

In the first case nothing needs to be done. In the second case, if B fails to get added to A ’s neighbor table, then the network still satisfies all object requests, but the stretch may increase. Local optimization mitigates this problem. If an exact answer is desired, we can rerun the neighbor table building algorithm after a random amount of time.

The third case is a much greater cause for concern, since if A has a hole where B should be, Property 1 would no longer hold. This could mean that some objects become unavailable. This problem could be solved by reinserting the node, but before the reinsertion occurs, objects may be unavailable. This is a serious problem, and this section presents our solution. (Recent work by Liu and Lam [17] also addresses this problem; their solution has the advantage that it does not require that non-joining nodes maintain state about on-going joins.)

We start with a definition:

DEFINITION 1. Assume that we start with a consistent Tapestry network. A core node is a node that is completely integrated in this network, that is, it has no holes in its neighbor table that can be filled by other nodes in the network, and it cannot fill holes in the neighbor tables of nodes in the network.

Together, the core nodes all satisfy Property 1. By this definition, a node could be a core node without meeting locality Property 2. The goal of this section is to prove that when a node finishes its multicast, it becomes a core node, and that all the nodes that were core nodes before its multicast finishes remain core nodes. We also add the requirement that any multicast, including those used in insertion, must start at a core node.

Two operations are *simultaneous* if there is a point in time when both operations are ongoing. This straight forward definition is important from a systems standpoint. It is a bit imprecise, however,

```

method OPTIMIZEOBJECTPTRS (sender, changedNode, objPtr, level)
1   oldsender  $\leftarrow$  GETOLDSENDER(objPtr)
2   if oldsender  $\neq$  null and oldsender  $\neq$  sender
3     OPTIMIZEOBJECTPTRS [on NEXTHOP(objPtr, level)] (self, changedNode, objPtr, level + 1)
4   if oldsender  $\neq$  changedNode
5     DELETEPOINTERSBACKWARD [on oldsender] (objPtr, changedNode, level - 1)
end OPTIMIZEOBJECTPTRS
method DELETEPOINTERSBACKWARD (changedNode, objPtr, level)
1   oldsender  $\leftarrow$  GetOldSender(objPtr)
2   DELETE(objPtr)
3   if oldsender  $\neq$  changedNode
4     DELETEPOINTERSBACKWARD [on oldsender] (objPtr, changedNode, level - 1)
end DELETEPOINTERSBACKWARD

```

Figure 9: OptimizeObjectPtrs and its helper function.

```

method OBJECTNOTFOUND (objectID)
1 if (Inserting)
2   level  $\leftarrow$  LENGTH(GREATESTCOMMONPREFIX(NewNodeName, PSurrogateName))
3   FINDOBJECT [on PSurrogateName] (objectID, level)
4 elseif not ROUTINGCONSISTENTWITHNEIGHBORS(objectID)
5   RETRYROUTING(objectID, Neighbors)
6 endif
end OBJECTNOTFOUND

```

Figure 10: Misrouting and route correction. Misrouting and route correction are used to keep objects available even during insertion.

since two multicasts could be simultaneous and yet be indistinguishable from sequential multicasts without the use of a global clock. We would like to distinguish between cases where all core nodes see evidence consistent with a sequential ordering and cases where there can be no such agreement. To this end, we say that two multicasts *conflict* if there are two nodes that receive the multicasts in different orders.

THEOREM 6. *Suppose A inserts. When it is done with its multicast, A 's table has no holes that can be filled by core nodes. Further, there are no core nodes with holes that A can fill. These statements are true even when other insertions proceed simultaneously.*

This is a proof by induction. We order the nodes by when they finish their multicasts. The we say that all nodes that have finished before A satisfy the theorem, and we prove the same is true of A . (Note that we are not assuming that there are no ongoing multicasts when A starts its multicast.) We start by proving a series of lemmas. Our first lemma is simple, but important. Lemma 3 states that simultaneously inserting nodes cannot interfere with one another's access to core nodes.

LEMMA 3. *Nodes in S , the set of core nodes, can be reached by a given multicast even in the presence of ongoing or completed insertions of other nodes.*

PROOF. Proof by contradiction. Theorem 5 says that all core nodes can reach one another. Suppose there is a multicast that misses node $X \in S$. Let B be the node that should have sent the multicast towards X but did not. Further, suppose that the prefix B received with the multicast was α . If B did not send the multicast towards X (that is, send it to (α, j) where $\alpha \circ j$ is a prefix of X 's ID), it must have been because it did not have a (α, j) node in its table. But this is not possible:

Case 1: B has not yet finished its multicast. Since B is supposed to send the multicast to X , we know that B and X share prefix α . Further, since we know that X was in the network before B began its multicast, B 's multicast consists of nodes with prefix α .

But this means that B would only have filled (α, j) entries, so it could not possibly have been contacted with a prefix smaller than $\alpha \circ j$. Contradiction.

Case 2: B has finished its multicast and is a core node. By Theorem 6, since X is an (α, j) node, B must have such a node in its table. Contradiction.

Although Theorem 6 uses Lemma 3, Case 2 is not circular: Node B was inserted *before* the point in time that we use it here. \square

It remains to deal with the case where two insertions conflict. We first introduce the notion of a pinned pointer. An (α, j) pointer to node A stored at node X is *pinned* when there are nodes whose multicasts through (α, j) have arrived at X but have not been acknowledged.

When a multicast for a new node filling an (α, j) slot arrives at some node X , X puts the new node in the table as a pinned pointer and sends the multicast to one unpinned pointer and all pinned pointers. When X receives acknowledgments from all recipients, X unlocks the pointer. Finally, X must keep at least one unpinned pointer and all pinned pointers. If this is done, then X will reach all (α, j) nodes it knows about without having to store them all. Intuitively, the unpinned pointer can reach all other unpinned pointers so unpinned pointers are all equivalent, while the pinned pointers are not well-enough connected to be reachable via multicast.

LEMMA 4. *A multicast through an unpinned (α, j) pointer at node X reaches all other nodes that have or had unpinned (α, j) pointers at node X .*

The proof is similar to other multicast arguments. Ideally, each multicast will see the other as completed. To enforce this condition, if any node gets a multicast from A , and notices that the hole for A is already filled, it contacts all nodes it has seen that fill that hole. As above, it contacts one unpinned pointer and all the pinned pointers.

Next, we deal with the case when A and B fill the same hole.

```

method ACKNOWLEDGEDMULTICAST( $\alpha$ , FUNCTION, holebeingfilled, watchlist, NewNodeIP)
1 watchlist  $\leftarrow$  CHECKFORNODESANDSEND(watchlist, NewNodeIP)
2 if NOTONLYNODEWITHPREFIX( $\alpha$ )
3   for  $i = 0$  to  $b - 1$ 
4     neighbor  $\leftarrow$  GETMATCHINGNEIGHBOR( $\alpha \circ i$ )
5     if neighbor exists
6        $\mathcal{S} \leftarrow$  ACKNOWLEDGEDMULTICAST [on GETIP(neighbor)] ( $\alpha \circ i$ , FUNCTION, holebeingfilled, watchlist, NewNodeIP)
7   else
8     apply FUNCTION
9      $\mathcal{S} \leftarrow$  MULTICASTTOFILLEDHOLE(holebeingfilled, FUNCTION, watchlist, NewNodeIP)
10 wait  $\mathcal{S}$ 
11 SENDACKNOWLEDGEMENT()
end ACKNOWLEDGEDMULTICAST

```

Figure 11: Acknowledged multicast with the watch list. This the version of acknowledged multicast used to handle simultaneous insertions.

```

method DELETESelf()
1   for pointer in { backpointers }
2     level = GETLEVEL(pointer)
3     LEAVINGNETWORK [on GETIP(pointer)] (selfID, level, GETNEAREST(pointer, level))
4   for pointer in { neighbors  $\cup$  backpointers }
5     REMOVELINK [on GETIP(pointer)] (selfID)
end DELETESelf

```

Figure 12: Voluntary Delete. This shows what a node should do when it leaves the network.

LEMMA 5. Suppose A and B fill the same hole. Then with the modification described above, if A 's multicast conflicts with B 's, A will get B 's multicast message before B 's multicast is finished.

PROOF. Let X be a node that gets A 's multicast before B 's. Then when X gets B 's multicast, it forwards it on to A , since they fill the same hole. Finally, since X does not send an acknowledgment until A returns an acknowledgment of B 's multicast, A has been informed by the time B 's multicast finishes. We then apply this same argument with the roles of A and B reversed. \square

We are not yet done. Consider when the $\alpha \circ i$ hole and the $\alpha \circ j$ hole are both being filled by two different nodes (with $i \neq j$). Then the $\alpha \circ i$ node may not get the $\alpha \circ j$ multicast and vice versa, even though their multicast sets are the same.

So, we further modify the multicast. The starting node sends down a "watch list" of prefixes for which it knows no matching node. This can be represented as a bit vector. When the inserting node sends this to the surrogate, it is a zero for every entry in the neighbor table. Each receiving node checks the watch list to see if it can fill in any blank on the list. If it can, it sends the relevant node to the originator of the multicast, marks the entry as found, and continues the multicast. From this description, it may sound as if we are sending a lot of information; in fact, we will be sending very little, since most of the lower levels of the table will be filled by the surrogate in the first step, and most of the upper levels of the table will be zero. In the normal case, we send only a few levels of the neighbor table, and each level is sixteen bits. This new version is shown in Figure 11. Using this new multicast, we get Lemma 6.

LEMMA 6. Let A be an α node, and let B be an (α, j) node. Then if the core α nodes get multicast messages from both A and B , the (α, j) slot on A will not be a hole. (The core nodes are those that have finished their multicasts when the latter of A and B start its multicast.)

PROOF. There are two cases.

Case 1: One α node, X , gets B 's multicast first and then A 's. In this case, when A 's multicast arrives on X , X checks A 's watch list, and if the watch list has an (α, j) hole B can fill, X has that hole filled and so will be able to notify A that it too can fill that hole. If there is no hole in the watch list, then A has already found such a node.

Case 2: All core α nodes gets A 's multicast first. This means that A gets the multicast about B .

This completes the proof. \square

Finally, we put everything together and prove Theorem 6.

PROOF. Consider a node B , and let α be the longest shared prefix between A and B .

Case 1: If A and B fill different holes on the same level (i.e., A fills an (α, i) hole and B fills (α, j) hole for $i \neq j$), then they multicast to the same prefix α . By Lemma 3, we know these nodes are reached, and we can apply Lemma 6, once with A in the theorem being A of the lemma, and once with A in the theorem as B in the lemma.

Case 2: If A and B fill different holes on different levels, then there are core α nodes in the network, and by Lemma 3 we know these nodes are reached. Given that, we again apply Lemma 6.

Case 3: If A and B fill the same hole on the same level, then there might not be a core node with prefix α so the preceding arguments fail. In this case, we rely on Lemma 5, which says that if the two multicasts are not serialized, each will find out about the other before their multicasts complete.

This completes the proof. \square

This algorithm can result in multiple multicasts to a new node. However, this is uncommon and it is rare that the new node will be anything other than a leaf in the tree (i.e. the new node will not forward the multicast). Further, the new node can easily suppress duplicate multicast messages.

4.5 Running Time Analysis

Without objects, the total network traffic required for node insertion is $O(d \log b)$ in expectation, where d is the diameter of the network. The total number of hops is $O(\log^2 n)$ with high probability.⁸ Finding the surrogate is no more costly than searching for an object pointer, and [20] argues that finding an object pointer requires $O(d)$ network traffic (and $O(\log n)$ hops). Multicast takes time $O(kd)$ where k is the number of nodes reached. But k will be small in expectation, and bounded by $\log n$ with high probability. Finally, building the neighbor tables takes $O(\log^2 n)$ messages. If there are m objects that should be on the new node, then the cost of republishing all those objects is at most $O(md)$. This gives a total traffic of $O(md \log n)$ for object pointer relocation.

5 Node Deletion

In this section, we present algorithms that help maintain our invariants when nodes leave the network. We consider two cases: *voluntary* and *involuntary* delete. A *voluntary* delete occurs when a node informs the network that it is about to exit. This is the preferred mode of deletion that permits the infrastructure to maintain availability of objects by fixing neighbor links and object pointers. An *involuntary* delete occurs when a node ceases to participate in the network without warning, due to a node failure, a network failure, or an attack. Note that it is unreasonable to hope that all deletes are voluntary deletes, and we present an algorithm for this case only for completeness. In real networks, nodes and links will typically fail without warning, so involuntary delete is the common case. We discuss this case in Section 5.2.

5.1 Voluntary Delete

When node A decides to leave the network, it ideally removes itself in a way that gives the infrastructure time to adapt its routing mesh and object pointers to maintain object availability. A begins by sending its intention to leave the network to all nodes on its backpointers list (all nodes which currently point to A somewhere in their routing table). Along with this notification, A sends along a potential replacement for itself on each routing level. On each such node (N), links to A are marked as “leaving.”

Removing the link to A could leave N with an incorrect hole in its routing table (breaking Property 1). This problem is mitigated by any existing secondary pointers backing up A and by potential replacements A sends with its notification. Node N may still wish to run the nearest-neighbor algorithm to tune the neighbor table.

When this initial notification is received by node N , it republishes any local object pointers which normally route through A as if A did not exist. Any incoming queries still route normally to A while it is marked as “leaving.” Publish operations, however, route to both A and its replacement. See Figure 12.

After node A sends out its initial notification messages, it examines local object pointers for which it is the root, and forwards them on to their respective surrogate nodes. Once all of these objects have found new root nodes and acknowledgments are received by A , objects that were rooted at A are now reachable through new surrogates. Thus availability is guaranteed. Node A then sends out a final delete notification to its backpointers, telling them to delete

⁸By determining k dynamically, we can show that this cost is $O(\log n)$ with high probability.

A from their routing tables completely. After all such nodes have responded, A disconnects. If it is allowable for objects to be temporarily unavailable, much of this work can be skipped, and the notification can happen in one phase rather than two.

5.2 Involuntary Delete

The other class of node deletions occurs as a result of any failure which prevents a node from performing normal operations as a Tapestry node. For simplicity, we consider here only complete failures such as network partitions, hardware failures, or complete system halts. In these scenarios, we would like the rest of the Tapestry network to detect this node’s failure and recover as much as possible to maintain object availability and full reachability of the routing mesh.

We propose that unexpected deletes be handled lazily. That is, when a node N notices some other node is down, it does everything it can to fix its own state, but does not attempt to dictate state changes to any other node. In the process of fixing its state, however, N may hint to other nodes that their state may be out of date. Deletion can be detected by soft-state beacons [32] or when a node sends a message to a defunct node and does not get a response.

When node N detects a faulty node, it should first remove the node from its neighbor table and find a suitable replacement. If this produces a hole in the table, N will have to find a replacement, to ensure Property 1 is maintained. Otherwise, N has several options depending on how good the replacement must be. It can find a replacement using a simple local search algorithm; that is, asking its remaining neighbors for their nearest matching nodes. This is not guaranteed to give the closest replacement node. Alternatively, the nearest neighbor algorithm can be repeated. In any case, it should also use OPTIMIZEOBJECTPOINTERS on all object pointers that would have gone through the deleted node.

To ensure Property 1, if deleting the node leaves a hole in its routing table, we must either find a replacement, or determine that none exists. To do so, we could use a multicast to all nodes sharing the same prefix of N and the dead node. While this is a workable solution, the multicast algorithm assumes all tables are complete, and may not reach a given node if some table along the path is incomplete. We can do slightly better. Liu and Lam [17] present a notification algorithm similar to multicast that solves this problem. Their algorithm has the property that if the node is in the table of *any* contacted node, then it will be returned to the node starting the multicast. Furthermore, their notification algorithm also requires that only the starting node to maintain state. This makes it ideal for this application.

One concern is that if a node N disappears, every node that used to point to N might start such a search, causing more traffic than is desirable. At the cost of centralization, we can pick one node (say the surrogate of the departed node) to perform the search, and have the surrogate return the answer when done.

Unfortunately, this does not maintain object availability. Objects rooted at the deleted node may become unavailable until a republish arrives at the node’s surrogate. In fact, a network partition may result in an inconsistent deletion; we do not address this here.

6 Realistic Deployment

In previous Sections we presented algorithmic solutions for maintaining Properties 1 and 2 during membership changes to the overlay network. These algorithms assume that the network satisfies the expansion property and that the distances are unchanging. In this section, we reconsider these assumptions and other issues in the context of physical networks such as the Internet, and provide mechanisms to improve performance.

One key assumption is the expansion property. Castro, Druschel,

Hu and Rowstron [6] address the question of the expansion property by running simulations on several different topology models. Their results suggest that even when the expansion property does not hold, it can be useful in guiding the design of overlay networks.

6.1 Locality

Tapestry as an object location infrastructure is flexible in the sense that it does not replicate data, only references to data. In systems that do not consider stretch (and consider instead, for example, network hops), one can always add a level of indirection and view data items as pointers at a cost of only one additional hop, so this distinction may not seem significant. Notice, however, that adding a level of indirection can drastically change the stretch properties even when it does not change significantly change the number of network hops.

Further, an object location system that does not allow arbitrary object placement is much less flexible. Systems such as Tapestry allow applications to choose their own data placement policies, (such as placing replicated objects near hotspots) and automatically make use of close resources when possible. This reduces network traffic, bandwidth, and may improve reliability as well, since less links are traversed on an access to a nearby object.

6.2 Physical Network Topologies

Several research projects have tried to create accurate models of the current Internet. One of the more widely accepted models is the transit-stub network model [29]. Transit-stub networks may or may not have the expansion property, depending on the layout of nodes inside the stub. Ideally, we want to provide good performance on such networks regardless of the intra-stub layout.

Because latency differences between intra-stub paths and inter-stub paths can be an order of magnitude or greater, we would like to ensure that an object locate request never leaves the originating stub if there is a copy of the object somewhere inside the stub. We propose an optimization for object publication and locate operations, that makes an effort to limit the operation to the local stub domain.

We assume that Tapestry nodes can detect whether the next hop is within the same stub network (possibly by setting a local network latency threshold). When object publication is about to route out of the local network, it spawns off a “local branch” publish message which treats the local network as its entire domain. While the original routes out to the wide-area, the local publish uses surrogate routing to route to a local root, where it terminates. If an object exists in the local network, then the request will terminate at or before the local object root. Otherwise, it examines the current node’s ID to determine when surrogate routing started, and resumes at that hop to continue with normal object location.

For example, suppose the node 1224 is trying to find an object labeled 1234. If a 123X node exists, but lies outside the local network, node 1224 uses surrogate routing inside the local stub, and tries to send the message to the closest 124X node in the stub. If none exists, it sends to the closest 125X node in the stub, and so on, until it finds a definitive local root node for 1234. If the object is not found before or at the local stub root, then node 1224 resumes normal routing outside the stub (to 123X). On the publish side, if a local node is publishing the object with ID 1234 and discovers the next hop is outside of the local network, it forwards two publish messages for the same ID. One continues as a normal publish outside of the network, while the other is restricted to the local-area, and publishes with surrogate routing to a local object root.

The net effect of this optimization is that queries for objects in the same stub network are always resolved without routing outside the stub network. The tradeoff is that queries for remote objects

will pay the price of additional surrogate routing hops inside the stub network. Because these are local hops, and surrogate routing lasts less than 2 hops in expectation [32], systems where there is network locality between clients and objects they query should see significant performance gains as a result of eliminating wide-area network traffic.

6.3 Continual Optimization

The other assumption we make is that the network distance metric, network latency for example, is constant over time. In the real Internet, routes frequently change due to a variety of factors, some of which are configuration changes to Border Gateway Protocol (BGP), changes in packet forwarding policy between Internet Service Providers (ISPs) or recalculation of IP routes after router failures inside Autonomous Systems (ASs). Therefore, network distance can change over time, potentially thwarting our efforts to provide locally optimal routes at each hop. In this section, we discuss some heuristics that should work well in practice to better adapt our network mesh to the underlying network distance instability.

There are a range of optimizations which trade off computation and network traffic for more up-to-date network distances. The minimum is to do nothing; the other extreme calls for periodic repetitions of the entire neighbor table algorithm. Only slightly more complicated than doing nothing is to change the primary neighbor to another node in the neighbor set. Recall that for any entry, there are R nodes; we periodically adjust which of these R nodes is the primary.

A third option is for a node to store all the nodes contacted in its search for a nearest neighbor. This is at most $O(\log^2 n)$ nodes. Then we can optimize one level or one entry at a time. To optimize level i , for example, we contact all the level- $(i + 1)$ nodes that were contacted during the original building process, and rebuild level- i . Since this should be an infrequent operation, the additional storage can be flushed to disk to minimize memory overhead. The frequency can be set for a given type of network, or can be dynamically triggered when nodes discover significant changes in their expected node to node performance.

Another option is to use local sharing of information. That is, periodically, a node sends to its level- i neighbors a copy of its level- i neighbor table. The receiving nodes repeat distance measurements to compare those neighbors with its own, replacing further away nodes where appropriate. This is the same idea as the heuristic neighbor table building algorithms in [32, 23].

In all these cases, when a new primary neighbor has been chosen, the node needs to move some object pointers. This can be done efficiently as described in Section 4.2.

6.4 Reliable Messaging and Sybil Attacks

Finally, we mention two more significant factors to consider in real deployment of systems such as Tapestry. First, we reconsider our assumption that all messages arrive reliably at their destinations in order. In the Internet, packets can be dropped, connections can be broken, and communication reliability cannot be guaranteed. This is exacerbated by the large scale nature of systems like Tapestry.

To provide resilience against unpredictable failures, we use soft-state. Soft-state is state that is not needed for correctness; in our case, using soft-state means that data is periodically refreshed by its source, requiring no explicit state maintenance or failure-recovery modes. For example, nodes use periodic heartbeat messages to detect node and link failures, while objects are republished at regular intervals in order to maintain high availability. Our use of soft-state allows us to quantitatively limit loss of availability due to failures and most attacks, while significantly simplifying our algorithms.

In addition to failures, we also need to be aware of malicious

users and potential attacks on the system (Sit and Morris [25] describe some of these attacks). One particularly devastating attack occurs when a single attacker impersonating a large number of physical nodes, in an attempt to gain control over key positions in the routing mesh. By generating enough randomized node IDs, an attacker would eventually obtain IDs that position it around an intended target node in the mesh. Attacks can then be carried out to deny service or manipulate and spoof communication.

These attacks, referred to as Sybil Attacks [10], are extremely hard to detect and circumvent. Given enough physical resources, an attacker can be hard to overcome. One viable solution is to use a centralized certificate authority to allocate node IDs. Its centralized nature, however, might lead to problems in scalability. This problem, like several others related to security, remain unsolved and is the subject of active research. (See Castro Druschel, Ganesh, and Rowstron [5] for some approaches to this problem.)

7 Object Location in General Metric Spaces.

For any metric space S , we show a way to route to an object such that the stretch is polylogarithmic with $O(|ID| \log^2 n)$ average space, where $|ID|$ is the size of an object ID. We remark that this is the strawman scheme proposed by Plaxton, Rajaraman, and Richa [20] without load balancing, and is similar to the scheme of Thorup and Zwick [27]. The proof is reminiscent of the metric embedding results of Bourgain [4], and Linial, London, and Rabinovich [16].

Let $S_{i,j}$ be a sample of the metric space such that each node is chosen with probability $2^i/n$, and let $i \in [1, \log n]$ and $j \in [0, c \log n]$. Pick a single node at random to be in $S_{0,0}$. Each node in the network stores the closest node in $S_{i,j}$ for each pair i, j . Also, each node in $S_{i,j}$ stores a list of all objects located at nodes which point to it.

Suppose node X wants to find object Y . Starting with $i = \log n$, X asks (for all j in parallel) its representative in the set $S_{i,j}$ if it knows of Y . If one of them does, it returns the pointer to Y . If this fails, it tries $S_{i-1,j}$ for all j . Recall that there is one node in $S_{0,0}$, so this will always find the object, if it exists. The following theorem is key to showing stretch bounds.

THEOREM 7. *Let i^* be the largest i such that there is some $S_{i,j}$ that points to both X and Y . We will show that $d(S_{i^*,j}, X) \leq d(X, Y) \log n$ with high probability. Moreover, the average space used by the data structure is $O(\log^2 n)$.*

PROOF. Let $\mathcal{B}_X(r)$ be the ball around X of radius r , that is, all the nodes within distance r of X . Now, consider a sequence of radii such that $r_k = kd$ for $k \in [1, \log n]$ and $d = d(X, Y)$. If $|\mathcal{B}_X(r_k) \cap \mathcal{B}_Y(r_k)| \geq \frac{1}{2} |\mathcal{B}_X(r) \cup \mathcal{B}_Y(r_k)|$ we call r_k good. We now show that if there exists a good r_k the theorem holds.

Let $r = r_k$ be a good radius. Then consider i such that

$$2^{\log n - i} \leq |\mathcal{B}_X(r) \cup \mathcal{B}_Y(r)| \leq 2^{\log n - i + 1}.$$

When $|\mathcal{B}_X(r) \cap \mathcal{B}_Y(r)|$ is $\frac{1}{2} |\mathcal{B}_X(r) \cup \mathcal{B}_Y(r)|$, for a given j , with constant probability there will be exactly one member of $S_{i,j}$ in the intersection and no other member in the union. We view each j as a trial, and since we have $c \log n$ trials, with high probability at least one will succeed. And if there is $s \in S_{i,j}$ that points to both X and Y , when X queries s , X will get a pointer to Y , so $i^* = i$.

We now argue that some r_k is good. Suppose that r_k is bad. Then $|\mathcal{B}_X(kd) \cap \mathcal{B}_Y(kd)|$ is less than half of $|\mathcal{B}_X(kd) \cup \mathcal{B}_Y(kd)|$. Notice that $\mathcal{B}_X(kd) \cap \mathcal{B}_Y(kd)$ contains

$|\mathcal{B}_X((k-1)d) \cup \mathcal{B}_Y((k-1)d)|$, and since

$$\begin{aligned} |\mathcal{B}_X(kd) \cup \mathcal{B}_Y(kd)| &\geq 2 |\mathcal{B}_X(kd) \cap \mathcal{B}_Y(kd)| \\ &\geq 2 |\mathcal{B}_X((k-1)d) \cup \mathcal{B}_Y((k-1)d)|, \end{aligned}$$

we can say that

$$|\mathcal{B}_X(kd) \cup \mathcal{B}_Y(kd)| \geq 2 |\mathcal{B}_X((k-1)d) \cup \mathcal{B}_Y((k-1)d)|.$$

But this can happen at most $\log n$ times, since $|\mathcal{B}_X(r_1) \cap \mathcal{B}_Y(r_1)| \geq 2$ (since it contains X and Y) and the network has only n nodes.

Finally, if at any point $|\mathcal{B}_X(r_k) \cup \mathcal{B}_Y(r_k)|$ contains the whole network, then let $i^* = 0$, and since there is only one element of each $S_{0,0}$, it will clearly be pointed to by X and have a pointer to Y . \square

To get the stretch bound, notice that if $d(S_{i^*,j}, X) \leq d(X, Y) \log n$, the total distance traveled on level i is $d(X, Y) \log^2 n$, and the latency (waiting time) is $d(X, Y) \log n$. Since there may be $\log n$ levels, this means the total latency is proportional to $d(X, Y) \log^2 n$ and the total distance traveled proportional to $c \cdot d(X, Y) \log^3 n$.

Note that we have assumed implicitly that the distance to the nearest $S_{i+1,j}$ is always less than the distance to $S_{i,j}$. This may not be strictly true. To make it true, we can require that $S_{i,j} \subset S_{i+1,j}$. Doing this would change the probability of a point being in $S_{i,j}$ only slightly, so the result still holds.

To provide load balancing, we let i range over all possible ID prefixes, and only search i 's that are prefixes of Y 's ID. This results in a very large table size. We do not know how to efficiently maintain this data structure.

8 Conclusion

We illustrate how to adapt to arriving and departing nodes in Tapestry, a location-independent overlay routing infrastructure with routing locality. This adaptation involves an efficient, distributed solution to the nearest-neighbor problem as well as a distributed algorithm for maintaining the prefix-based routing mesh. One of the salient properties of our system is that objects remain available, even as the network changes. Further, the cost of integrating new nodes is similar to that of systems that do not provide routing locality. The result is an infrastructure that provides deterministic location, routing locality, and load balance – even in a changing network.

9 Acknowledgments

We would like to thank Sean Rhea for his careful reading of this paper. We would also like to thank Kevin Chen, Megan Thomas, and the members of the OceanStore group and the anonymous reviewers for their insightful comments.

10 REFERENCES

- [1] AWERBUCH, B., AND PELEG, D. Concurrent online tracking of mobile users. In *Proc. of SIGCOMM* (Sept. 1991), pp. 221–233.
- [2] AWERBUCH, B., AND PELEG, D. Routing with polynomial communication-space trade-off. *SIAM Journal on Discrete Mathematics* 5, 2 (1992), 151–162.
- [3] BOLOSKEY, W. J., DOUCEUR, J. R., ELY, D., AND THEIMER, M. Feasibility of a serverless distributed file system deployed on an existing set of desktop PCs. In *Proc. of SIGMETRICS* (June 2000), pp. 34–43.

- [4] BOURGAIN, J. On Lipschitz embedding of finite metric spaces in Hilbert space. *Israel J. Math* 52 (1985), 46–52.
- [5] CASTRO, M., DRUSCHEL, P., GANESH, A., AND ROWSTRON, A. Secure routing for structured peer-to-peer overlay networks. In *Proceedings of the 5th Symposium on Operating Systems Design and Implementation* (Dec. 2002).
- [6] CASTRO, M., DRUSCHEL, P., HU, Y. C., AND ROWSTRON, A. Exploiting network proximity in peer-to-peer overlay networks. In *Proceedings of the International Workshop on Future Directions in Distributed Computing* (June 2002).
- [7] CLARKE, I., SANDBERG, O., WILEY, B., AND HONG, T. W. Freenet: A distributed anonymous information storage and retrieval system. *Lecture Notes in Computer Science* 2009 (2001), 46–66.
- [8] CLARKSON, K. L. Nearest neighbor queries in metric spaces. In *Proc. of the 29th Annual ACM Symp. on Theory of Comp.* (May 1997), pp. 609–617.
- [9] COWEN, L. J. Compact routing with minimum stretch. In *Proc. of the 10th Annual ACM-SIAM Symp. on Discrete Algorithms* (Jan. 1999), pp. 255–260.
- [10] DOUCEUR, J. The Sybil attack. In *Proceedings of IPTPS* (Mar. 2002), pp. 251–260.
- [11] GAVOILLE, C. Routing in distributed networks: Overview and open problems. *ACM SIGACT News - Distributed Comp. Column* 32, 1 (Mar. 2001), 36–52.
- [12] GOPAL, B., AND MANBER, U. Integrating content-based access mechanisms with hierarchical file systems. In *Proc. of ACM OSDI* (Feb. 1999), pp. 265–278.
- [13] KARGER, D., AND RUHL, M. Find nearest neighbors in growth-restricted metrics. In *Proc. of the 34th Annual ACM Symp. on Theory of Comp.* (May 2002), pp. 741–750.
- [14] KUBIATOWICZ, J., ET AL. OceanStore: An architecture for global-scale persistent storage. In *Proc. of ACM ASPLOS* (Nov. 2000), pp. 190–201.
- [15] LI, X., AND PLAXTON, C. G. On name resolution in peer-to-peer networks. In *Proceedings of the second ACM international workshop on Principles of mobile computing* (Oct. 2002), pp. 82–89.
- [16] LINIAL, N., LONDON, E., AND RABINOVICH, Y. The geometry of graphs and some of its algorithmic applications. In *IEEE Symposium on Foundations of Computer Science* (1994), pp. 577–591.
- [17] LIU, H., AND LAM, S. S. Neighbor table construction and update in a dynamic peer-to-peer network. In *Proceedings IEEE ICDCS* (May 2003).
- [18] ORAM, A., Ed. *Peer-to-Peer, Harnessing the Power of Disruptive Technologies*. O’Reilly Books, 2001.
- [19] PELEG, D., AND UPFAL, E. A tradeoff between size and efficiency for routing tables. In *Proc. of the 21th Annual ACM Symp. on Theory of Comp.* (may 1989), pp. 43–52.
- [20] PLAXTON, C. G., RAJARAMAN, R., AND RICHA, A. W. Accessing nearby copies of replicated objects in a distributed environment. In *Proc. of the 9th Annual Symp. on Parallel Algorithms and Architectures* (June 1997), pp. 311–320.
- [21] RAJARAMAN, R., RICHA, A. W., VÖCKING, B., AND VUPULURI, G. A data tracking scheme for general networks. In *ACM Symposium on Parallel Algorithms and Architectures* (July 2001), pp. 247–254.
- [22] RATNASAMY, S., FRANCIS, P., HANDLEY, M., KARP, R., AND SCHENKER, S. A scalable content-addressable network. In *Proc. of SIGCOMM* (Aug. 2001), pp. 161–172.
- [23] ROWSTRON, A., AND DRUSCHEL, P. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In *Proceedings of IFIP/ACM International Conference on Distributed Systems Platforms* (Nov. 2001).
- [24] ROWSTRON, A., AND DRUSCHEL, P. Storage management and caching in PAST, a large-scale, persistent peer-to-peer storage utility. In *Proceedings of SOSP* (Oct. 2001), pp. 188–201.
- [25] SIT, E., AND MORRIS, R. Security considerations for peer-to-peer distributed hash tables. In *Proceedings of the first International Workshop on Peer-to-Peer Systems* (Mar. 2002), pp. 261–9.
- [26] STOICA, I., MORRIS, R., KARGER, D., KAASHOEK, M. F., AND BALAKRISHNAN, H. Chord: A scalable peer-to-peer lookup service for internet applications. In *Proc. of SIGCOMM* (Aug. 2001), pp. 149–160.
- [27] THORUP, M., AND ZWICK, U. Approximate distance oracles. In *Proc. of the 33th Annual ACM Symp. on Theory of Comp.* (July 2001), pp. 183–192.
- [28] THORUP, M., AND ZWICK, U. Compact routing schemes. In *Proc. of the 13th Annual Symp. on Parallel Algorithms and Architectures* (July 2001), pp. 1–10.
- [29] ZEGURA, E. W., CALVERT, K., AND BHATTACHARJEE, S. How to model an internetwork. In *Proc. of IEEE INFOCOM* (Mar. 1996), vol. 2, pp. 594–602.
- [30] ZHAO, B. Y., HUANG, L., RHEA, S. C., STRIBLING, J., JOSEPH, A. D., AND KUBIATOWICZ, J. D. Tapestry: A global-scale overlay for rapid service deployment. *IEEE Journal on Selected Areas in Communications* (2003). Special Issue on Service Overlay Networks, to appear.
- [31] ZHAO, B. Y., JOSEPH, A., AND KUBIATOWICZ, J. Locality-aware mechanisms for large-scale networks. In *Proc. of Workshop on Future Directions in Distributed Comp.* (June 2002), pp. 80–83.
- [32] ZHAO, B. Y., KUBIATOWICZ, J. D., AND JOSEPH, A. D. Tapestry: An infrastructure for fault-tolerant wide-area location and routing. Tech. Rep. UCB/CSD-01-1141, UC Berkeley, Computer Science Division, Apr. 2001.