

Brocade: Landmark Routing on Overlay Networks

CS262A Fall 2001

Yitao Duan, Ling Huang

University of California, Berkeley

duan@cs.berkeley.edu, hlion@newton.berkeley.edu

Abstract

Peer-to-peer networks offer attractive features such as decentralized computing, fault tolerance, sharing and aggregating of resources (e.g., SETI@home). A common (and limiting) assumption on peer-to-peer network is that, in order to achieve fair sharing and fault tolerance, all computers in the network should be considered equal and there should be no server nor explicit structuring in the network. Many theories and algorithms used by existing peer-to-peer infrastructures (such as Tapestry [12], Chord [9], Pastry [7] and CAN [4]) comply with this tenet. However, such systems often show sub-optimal performance due to the asymmetry of nodes in reality and the lack of structure among them.

In this paper, we argue that a peer-to-peer system does not have to operate in a pure peer-to-peer fashion to obtain the advantages mentioned above. We believe that such system can benefit from some form of internal organization and differentiation among its nodes, thus enjoying the efficiency of a well-organized system, while still providing users with the P2P features.

As a proof of our belief, we proposed and simulated Brocade, a secondary overlay to be layered on top of these systems. Brocade exploits knowledge of underlying network characteristics and recognizes the differences among the nodes. Simulations were run on the GT-ITM [11] transit stub topology and showed encouraging results.

1 Introduction

Existing peer-to-peer overlay infrastructures such as Tapestry [12], Chord [9], Pastry [7] and CAN [4] demonstrated the benefits of scalable, wide-area lookup services for Internet applications. These architectures make use of name-based routing to route requests for objects or files to a nearby replica. Applications built on such systems [2, 3, 8], depend on reliable and fast message routing to a destination node, given some unique identifier.

Due to the theoretical approach taken in these systems, however, they assume that most nodes in the system are uniform in resources such as network bandwidth and storage. They also make no differentiation with regard to each node's position in the network. In addition, the routing algorithms used in these systems are decoupled from the underlying network topology. The result is messages being routed on the overlay with minimum consideration to actual network topology and differences between node resources.

In reality, nodes in a network are asymmetric in terms of processing power, network bandwidth and connectivity. Given this, a flat, uniform operation model suffers from two aspects: (1) it tends to overload the less powerful nodes while leaving the more powerful ones under utilized and (2) it doesn't take advantage of the aggregated knowledge about the network structure it could have used. While the first point seems obvious, the second one deserves some elaboration. It comes from the intuition that a system is more efficient if its components, even though they are equal in power, are organized. For example, a B+ tree offers far superior searching performance than sequential scan because elements are organized in the first case. The same argument applies to network. It would be clear in the following sections that even though we do not assume the existence of any powerful nodes, we could still benefit from better organization.

In this project, we propose Brocade, a secondary overlay to be layered on top of these systems. Brocade exploits knowledge of underlying network characteristics and recognizes the differences among the nodes. The secondary overlay builds a location layer between "supernodes," nodes that are situated near network access points, such as gateways to administrative domains. By associating local nodes with their nearby "supernode," messages across the wide-area can take advantage of the highly connected network infrastructure between these supernodes to shortcut

across distant network domains, greatly improve point-to-point routing distance and reducing overall network bandwidth usage.

In this paper, we present the initial architecture of a Brocade secondary overlay on top of a Tapestry network, and demonstrate its potential performance benefits by simulation. Section 2 briefly describes Tapestry routing and location, Section 3 describes the design of a Tapestry Brocade, and Section 4 present preliminary simulation results. Finally, we discuss related and future work and conclude in Section 5.

2 Tapestry Routing and Location

Our architecture leverages Tapestry, an overlay location and routing layer presented by Zhao, Kubiatiowicz and Joseph in [12]. The Tapestry location and routing infrastructure is one of several recent projects exploring the value of wide-area decentralized location services [4, 7, 9]. It allows messages to locate objects and route to them across an arbitrarily-sized network, while using a routing map with size logarithmic to the network namespace at each hop. We present here a brief overview of the relevant characteristics of Tapestry. A detailed discussion of Tapestry algorithms, its faulttolerant mechanisms and simulation results can be found in [12].

Each Tapestry node or machine can take on the roles of server (where objects are stored), router (which forward messages), and client (origins of requests). Also, objects and nodes have names independent of their location and semantic properties, in the form of random fixed-length bit-sequences represented by a common base (e.g., 40 Hex digits representing 160 bits). The system assumes entries are roughly evenly distributed in both node and object namespaces, which can be achieved by using the output of secure one-way hashing algorithms, such as SHA-1 [6].

2.1 Routing Layer

Tapestry uses local routing maps at each node, called neighbor maps, to incrementally route overlay messages to the destination ID digit by digit (e.g., $***8 \Rightarrow **98 \Rightarrow *598 \Rightarrow 4598$ where $*$'s represent wildcards). This approach is similar to longest prefix routing in the CIDR IP address allocation architecture [5]. A node N has a neighbor map with multiple levels, where each level represents a matching suffix up to a digit position in the ID. A given level of the neighbor map contains a number of entries equal to the base of the ID, where the i th entry in the j th level is the ID and location of the closest node which ends in $"i"+\text{suffix}(N, j-1)$. For example, the 9th entry of the 4th level for node 325AE is the node closest to 325AE in network distance which ends in 95AE.

When routing, the n th hop shares a suffix of at least length n with the destination ID. To find the next router, we look at its $(n+1)$ th level map, and look up the entry matching the value of the next digit in the destination ID. Assuming consistent neighbor maps, this routing method guarantees that any existing unique node in the system will be found within at most $\log_b N$ logical hops, in a system with an N size namespace using IDs of base b . Since every neighbor map level assumes that the preceding digits all match the current node's suffix, it only needs to keep a small constant size ($\log_b N$) entries at each route level, yielding a neighbor map of fixed size ($b * \log_b N$). Figure 1 shows an example of hashed-suffix routing.

2.2 Data Location

Tapestry employs this infrastructure for data location. Each object is associated with one or more Tapestry location roots through a distributed deterministic mapping function. To advertise or publish an object O , the server S storing the object sends a publish message toward the Tapestry location root for that object. At each hop along the way, the publish message stores location information in the form of a mapping $\langle \text{Object-ID}(O), \text{Server-ID}(S) \rangle$. Note that these mappings are simply pointers to the server S where O is being stored, and not a copy of the object itself. Where multiple objects exist, each server maintaining a replica publishes its copy. A node N that keeps location mappings for multiple replicas keeps them sorted in order of distance from N .

During a location query, clients send messages directly to objects via Tapestry. A message destined for O is initially routed towards O 's root from the client. At each hop, if the message encounters a node that




Figure 1: Tapestry routing example. Path taken by a message from node 0325 for node 4598 in Tapestry using hexadecimal digits of length 4 (65536 nodes in namespace).

contains the location mapping for O , it is redirected to the server containing the object. Otherwise, the message is forward one step closer to the root. If the message reaches the root, it is guaranteed to find a mapping for the location of O . Note that the hierarchical nature of Tapestry routing means at each hop towards the root, the number of nodes satisfying the next hop constraint decreases by a factor equal to the identifier base (e.g. octal or hexadecimal) used in Tapestry. For nearby objects, client search messages quickly intersect the path taken by publish messages, resulting in quick search results that exploit locality. These and other properties are analyzed and discussed in more detail in [12].

3 Brocade Base Architecture

In this section, we present the overall design for the Brocade overlay proposal, and define the design space for a single instance of the Brocade overlay. We further clarify the design issues by presenting algorithms for an instance of a Tapestry on Tapestry Brocade.

To improve point to point routing performance on an overlay, a Brocade system defines a secondary overlay on top of the existing infrastructure, and provides a shortcut routing algorithm to quickly route to the local network of the destination node. This is achieved by finding nodes which have high bandwidth and fast access to the wide-area network, and tunneling messages through an overlay composed of these "supernodes."

In overlay routing structures such as Tapestry [12], Pastry [7], Chord [9] and Content-Addressable Networks [4], messages are often routed across multiple autonomous systems (AS) and administrative domains before reaching its destination. Each overlay hop often incurs long latencies across multiples AS's and multiple IP hops inside a single network domain, while consuming bandwidth. To minimize both latency and network hops and reduce network traffic for a given message, Brocade attempts to determine the network domain of the destination, and route directly to that domain. A "supernode" acts as a landmark for each network domain. Messages use them as endpoints of a tunnel through the secondary overlay, where messages would emerge near the local network of the destination node.

Before we examine the performance benefits, we address several issues necessary in constructing and utilizing a Brocade overlay. We first discuss the construction of a Brocade: how are supernodes chosen and how is the association between a node and its nearby supernode maintained? We then address issues in Brocade routing: when and how messages find supernodes, and how they are routed on the secondary overlay.

3.1 Brocade Construction

The key to Brocade routing is the tunneling of messages through the wide area between landmark nodes (supernodes). The selection criteria are that supernodes have significant processing power (in order to route large amounts of overlay

traffic), minimal number of IP hops to the wide-area network, and high bandwidth outgoing links. Given these requirements, gateway routers or machines close to them are attractive candidates.

The actual decision on deploying supernodes can be influenced by network administration policy. Given a selection of supernodes, we face the issue of determining one-way mappings between supernodes and normal tapestry nodes for which they act as landmarks in Brocade routing. One possibility is to exploit the natural hierarchical nature of network domains. Each network gateway in a domain hierarchy can act as a Brocade routing landmark for all nodes in its subdomain not covered by a more local subdomain gateway. We refer to the collection of these overlay nodes as the supernode's cover set. An example of this mapping is shown in Figure 2. Supernodes keep up-to-date member lists of their cover sets, which are used in the routing process, as described below.




Figure 2: Example of Supernode Organization

A secondary overlay can then be constructed on supernodes. Supernodes can have independent names in the Brocade overlay, with consideration to the overlay design, e.g. Tapestry location requires names to be evenly distributed in the namespace.

3.2 Brocade Routing

Here we describe mechanisms required for a Tapestry-based Brocade, and how they work together to improve long range routing performance. Given the complexity and latency involved in routing through an additional overlay, three key issues are: how messages are filtered so that only long distance messages are directed through the Brocade overlay, how messages find a local supernode as entry to the Brocader, and how a message finds the landmark supernode closest to the message destination in the secondary overlay.

3.2.1 Selective Utilization

The use of a secondary overlay incurs a non-negligible amount of latency overhead in the routing. Once a message reaches a supernode, it must search for the supernode nearest to the destination node before routing to that domain. Consequently, only messages that route outside the reach of the local supernode benefit from Brocade routing.

We propose a naive solution by having each supernode maintain a listing of all Tapestry nodes in its cover set. We expect the node list at supernodes to be small, with a maximum size on the order of tens of thousands of entries. When a message reaches a supernode, the supernode can do an efficient lookup (via hashtable) to determine whether the message is destined for a local node, or whether Brocade routing would be useful. To reduce message traffic at supernodes, we expect normal Tapestry nodes to maintain a "cache" of proximity measure of previously contacted destinations, and use it to determine whether forwarding to the local supernode for Brocade routing is worthwhile.

3.2.2 Routing Towards Supernodes

For a message to take advantage of Brocade routing, it must be routed to a supernode on its way to its destination. How this occurs plays a large part in how efficient the resulting Brocade route is. There are several possible approaches. We discuss three possible options here, and evaluate their relative performance in Section 4.

Naïve A naive approach is to make Brocade tunneling an optional part of routing, and invoke it, by supernode, only when a message reaches a supernode as part of normal routing. The advantage is simplicity. Normal nodes need to do nothing to take advantage of supernodes in the overlay infrastructure. They do not need to keep any information about supernodes, nor do they have to change their routing strategy. The disadvantage is that, due to the small supernodes to normal nodes ratio in the network, the probability of a message encountering a supernode in its early stage of Tapestry routing (when it can benefit from Brocade overlay) is very low. As a result, messages can traverse several overlay hops before encountering a supernode, or never encountering a supernodes, reducing the effectiveness of the Brocade overlay.

IP-snooping In an alternate approach, supernodes can "snoop" on IP packets to determine if they are Tapestry messages. If so, supernodes can parse the message header, and use the destination ID to determine if Brocade routing should be used. The intuition is that because supernodes are situated near the edge of local networks, any Tapestry message destined for an external destination will likely cross its path. This also has the advantage that the source node sending the message need not know about the Brocade supernodes in the infrastructure. The disadvantage is difficulty in implementation, and possible limitations imposed on regular traffic routing by header processing.

Directed A final solution is to require overlay nodes to remember the location of their local supernode. This information can easily be passed along as the node is inserted into the overlay. This state can be maintained easily by soft-state protocols, such as a periodic beacon from the local supernode to its cover set. Nodes can use a local cache to remember approximate distances to nodes they have sent messages to before, and use this information to decide whether a message is destined for a local node. If so, it routes the message normally. Otherwise it sends the message directly to its supernode for processing. The supernode, having an up-to-date list of its cover set, can determine precisely whether this message is actually destined for local node and send it back to its domain if so, thus correcting any false negative an overlay node may cause due to its limited cache size. The supernode can also send feedback to the source node as a hint to update the latter's cache. This is a proactive approach that takes full advantage of any potential performance benefit Brocade routing can offer. It does, however, require some fault-tolerant mechanism to inform local nodes of a replacement should a supernode fails.

3.2.3 Landmark Routing on Brocade

Once a message destined for a distant node arrives at some supernode near the sender, Brocade needs to determine the supernode closest to the message destination node. There is no restriction on how this is done. In the context of wide-area lookup services systems [12, 9, 7, 4], we propose to leverage the existing infrastructure. Recall that each supernode maintains a list of all overlay nodes inside its cover set. These overlay nodes can be treated as the objects or services the supernode possesses. Therefore determining the supernode closest to the message destination node is essentially an object location process at the Brocade level and any one of the mechanisms like Tapestry [12], Chord [9], Pastry [7] and CAN [4] can be used to accomplish this.

In this project, we investigated the approach of organizing the Brocade overlay as a Tapestry location infrastructure. As described in Section 2.2 and [12], Tapestry location allows nodes to efficiently locate objects given their IDs. In the Brocade overlay, each supernode advertises the IDs of the overlay nodes in its cover set as IDs of objects it "stores." When a supernode tries to route an outgoing interdomain message, it uses Tapestry to search for an object with an ID identical to the message destination ID. By locating the object on the Brocade layer, a supernode has found the local supernode of the message destination, and forwards the message directly to that supernode. The destination supernode receives the message and initiates normal overlay routing to the destination.

Note these discussions make the implicit assumption that routing between domains takes significantly more time compared to routing between nodes in a local domain. This, in combination with the distance constraints in Tapestry, allows us to assert that intradomain messages will rarely, if ever, be routed outside the domain. This is because the destination node will almost always offer the closest node with its own ID. This also means that once a message arrives at the destination's supernode, it is likely to quickly route to the destination node.

3.2.4 Optimization on Brocade Object Location

While Brocade significantly reduces the number of hops a message passes through, its performance in terms of latency depends to a great extent on the efficiency of object location mechanism used on Brocade overlay. We have observed that if the latency on interdomain links is high, the resulting latency for tunneling a message through Brocade is also high if no optimization on object location is done. This is due to the fact that Tapestry object location requires routing toward root to get the location information [12]. Clearly directly applying Tapestry object location on Brocade is not optimal.

A key observation is that Brocade is isolated from Tapestry routing in that (1) a supernode does not utilize all the information it gathers while constructing Tapestry routing table and (2) once a message gets on Brocade, all progress made in Tapestry routing is discarded. Based on this, we have investigated various ways to improve object location efficiency on Brocade. We describe here one approach that we think to be the most promising -- Bloom Filter.

Recall that each supernode keeps an up-to-date member list of overlay nodes that belong to its domain. Locating the supernode that an overlay node (the destination of a message) belongs to is essentially a group membership query problem. Bloom Filter offers quick test of membership thus provides a natural solution.

The cost of Bloom Filter can be shown to be minimal. Suppose there are 100 supernodes and 1 million ordinary nodes in Tapestry, if we use 4 hash functions and with a false positive rate less than 5%, the total memory required of a supernode is only $7 * 1000,000 \text{ bits} = 7\text{Mb} = 0.8 \text{ MB}$.

A more serious concern with the use of Bloom Filter is update. Bloom Filter is known for the difficulty to update and a wide-area network is certainly dynamic. This problem can be alleviated by the following soft-state consistency approach:

Each supernode is responsible for updating its own bit array and propagating it to other supernodes. The system can tolerate a certain degree of inconsistency in that incorrect testing of membership only affect performance, in a mild manner, but not correctness. Given this, a supernode can delay updating its Bloom Filter until the number of nodes that change in its domain reaches a threshold. The propagation can be done using multicast to reduce network load (the compressed nature of Bloom Filter should make the traffic less burdensome).

4 Evaluation of Base Design

In this section, we present initial simulation results showing the routing performance improvement possible with the use of Brocade overlays. In particular, we simulate the effect Brocade routing has on point to point routing latency and bandwidth usage. For our experiments, we implemented a two layer Brocade system inside a packet-level simulator that used Tapestry as both the primary and secondary overlay structures. The packet level simulator measured the progression of single events across a large network without regard to network effects such as congestion or retransmission.

To drive our experiments, we used the GT-ITM [11] transit stub topology generator to generate networks of 5000 nodes. We constructed Tapestry networks of size 4096, and marked 16 transit stubs as Brocade supernodes. We then measured the performance of pair-wise communication paths using original Tapestry and several Brocade algorithms. We experimented with all three Brocade algorithms for finding supernodes (Section 3.2.2). We include four total algorithms: 1. original Tapestry, 2. naive Brocade, 3. IPsnopping Brocade, 4. directed Brocade. For Brocade algorithms, we assume the sender knows whether the destination node is local, and only uses Brocade for interdomain routing.

We use as our key metric a modified version of Relative Delay Penalty (RDP) ([1]). Our modified RDP attempts to account for the processing of an overlay message up and down the protocol stack by adding 1 hop unit to each overlay node traversed. Each data point is generated by averaging the routing performance on 100 randomly chosen paths of a certain distance. In the RDP measurements, the sender's knowledge of whether the destination is local explains the low RDP values for short distances, and the spike in RDP around the average size of transit stub domains.

We measured the hop RDP of the four routing algorithms. For each pair of communication endpoints A and B, hop RDP is a ratio of # of hops traversed using Brocade to the shortest hop distance between A and B. As we can see from Figure 3, all Brocade algorithms improve upon the original Tapestry point to point routing. As expected, naive Brocade offers minimal improvement. Meanwhile, IP snooping improves the hop RDP substantially, while directed Brocade provides the most significant improvement in routing performance. For paths of moderate to long lengths, directed Brocade reduces the routing overhead by more than 50% to near optimal levels (counting processing time). The small spike in RDP for IP snooping and directed Brocade is due to the Tapestry location overhead in finding landmarks for destinations in nearby domains.

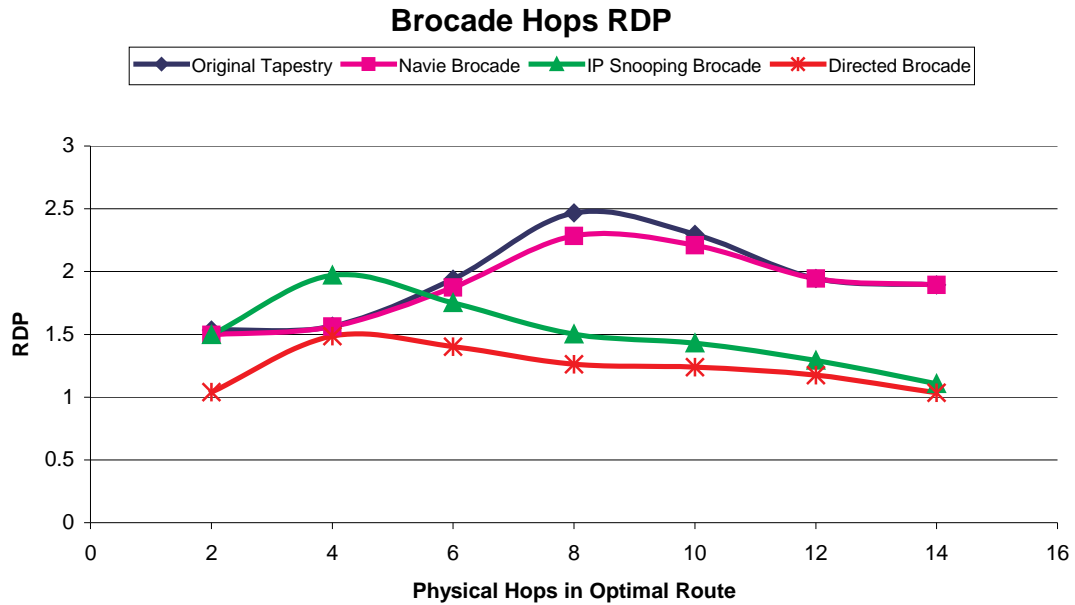


Figure 3: Hop-based RDP

Figure 3 makes a simple assumption that all physical links have the same latency. To account for the fact that interdomain routes have higher latency, Figure 4 shows an RDP where each interdomain hop counts as 3 hop units of latency. We see that IP snooping and directed Brocade still show the drastic improvement in RDP found in the simplistic topology results. We note that the spike in RDP experienced by IP snooping and directed Brocade is exacerbated by the effect of higher routing time in interdomain traffic making Tapestry location more expensive. We also ran this test on several transit stub topologies with randomized latencies direct from GT-ITM, with similar results.

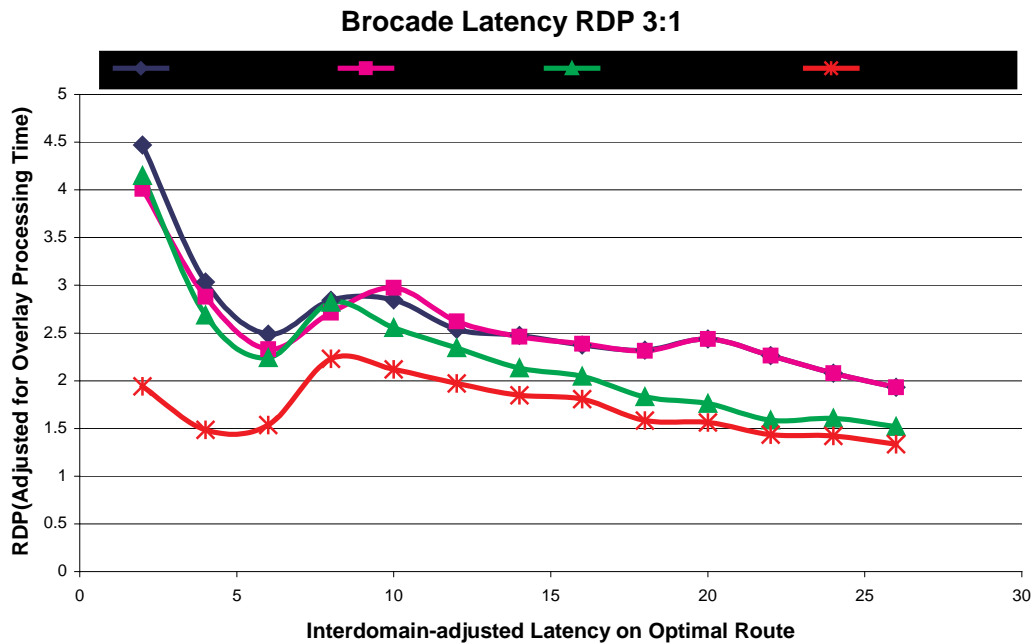


Figure 4: Weighted latency RDP, ratio 3:1

In addition, we examine the effect of Brocade on reducing overall network traffic, by measuring the aggregate bandwidth taken per message delivery, using units of (sizeof(Msg) * hops). The result in Figure 5 shows that IP snooping Brocade and directed Brocade dramatically reduce bandwidth usage per message delivery. This is expected, since Brocade forwards messages directly to the destination domain, and reduces message forwarding on the wide-area.




Figure 5: Aggregate bandwidth used per message

Finally we examine the effect of applying Bloom Filter to optimize Brocade object location. Figure 6 and 7 show the results. Both figures assume latency on Brocade links is 8 times that of ordinary link. Figure 6 shows the message latency RDP on Brocade without any optimization. Due to the long latency on its links, Brocade performs worse than Tapestry for short distance routing. Figure 7, on the other hand, shows the effect of Bloom Filter optimization. Because of the high accuracy with which we test an overlay node’s cover set membership, it displays near optimal performance.

While certain decisions in our design are Tapestry specific, we believe similar design decisions can be made for other overlay networks [7, 4, 9], and these results should apply to Brocade routing on those networks as well.

5 Related Work and Status

In related work, the Cooperative File System [2] leverages nodes with more resources by allowing them to host additional virtual nodes in the system, each representing one quantum of resource in storage, computational power and network bandwidth. This quantification is directed mostly at storage requirements, and CFS does not propose a mechanism for exploiting network topology knowledge. Our work is also partially inspired by the work on landmark routing [10], where packets are directed to a node in the landmark hierarchy closest to the destination before local routing. Finally, we are not currently aware of any work on building secondary overlays on peer-to-peer infrastructures.




Figure 6: Brocade latency RDP without Bloom Filter




Figure 7: Brocade latency RDP with Bloom Filter

While we present here a naming and routing architecture with design decisions based on a Tapestry overlay, the Brocade overlay architecture can be generalized on top of any peer-to-peer network infrastructure. In particular, the presented architecture works without change on the Pastry [7] network. We are currently exploring Brocades on top of Content-addressable Networks [4] and Chord [9], and also designing a Brocade for integration into the OceanStore [3] wide-area storage system. The Brocade design is also work in progress. In order to get more useful simulation results, we are generating more realistic transit stub topologies, and examining more closely the effect of longer interdomain

routes on Brocade performance. Finally, we are trying to reduce the high overhead of Tapestry location on the Brocade supernode level by experimenting with more efficient mechanisms for locating landmark nodes.

In conclusion, we have proposed the use of a secondary overlay network on a collection of well-connected "supernodes," in order to improve point to point routing performance on peer-to-peer overlay networks. The secondary Brocade layer uses Tapestry location to direct messages to the supernode nearest to their destination. Simulations show that taking shortcuts across Brocade significantly improves routing performance and reduces bandwidth consumption for a large portion of the point to point paths in a wide-area overlay. We believe Brocade is an interesting enhancement that leverages network knowledge for enhanced routing performance.

6 Acknowledgement

This project could not have materialized without the guidance and help from Professor Anthony Joseph and Professor Joe Hellerstein. Many thanks to Ben Y. Zhao who fed us with background information about P2P network in general and Tapestry in particular and helped us clarify our ideas. In addition, Professor John D. Kubiawicz provided many valuable insights into the operation and optimization of Tapestry.

7 REFERENCES

- [1] CHU, Y., RAO, S. G., AND ZHANG, H. A case for end system mul-ticast. In Proceedings of ACM SIGMETRICS (June 2000), pp. 1-12.
- [2] DABEK, F., KAASHOEK, M. F., KARGER, D., MORRIS, R., AND STOICA, I. Wide-area cooperative storage with CFS. In Proceedings of SOSP (October 2001), ACM.
- [3] KUBIATOWICZ, J., ET AL. OceanStore: An architecture for global-scale persistent storage. In Proceedings of ACM ASPLOS (November 2000), ACM.
- [4] RATNASAMY, S., FRANCIS, P., HANDLEY, M., KARP, R., AND CHENKER, S. A scalable content-addressable network. In Proceedings of SIGCOMM (August 2001), ACM.
- [5] REKHTER, Y., AND LI, T. An architecture for IP address allocation with CIDR. RFC 1518, <http://www.isi.edu/in-notes/rfc1518.txt>, 1993.
- [6] ROBSHAW, M. J. B. MD2, MD4, MD5, SHA and other hash functions. Tech. Rep. TR-101, RSA Laboratories, 1995. version 4.0.
- [7] ROWSTRON, A., AND DRUSCHEL, P. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In Proceedings of IFIP/ACM Middleware 2001 (November 2001).
- [8] ROWSTRON, A., AND DRUSCHEL, P. Storage management and caching in PAST, a large-scale, persistent peer-to-peer storage utility. In Proceedings of SOSP (October 2001), ACM.
- [9] STOICA, I., MORRIS, R., KARGER, D., KAASHOEK, M. F., AND ALAKRISHNAN, H. Chord: A scalable peer-to-peer lookup service for internet applications. In Proceedings of SIGCOMM (August 2001), ACM.
- [10] TSUCHIYA, P. F. The landmark hierarchy: A new hierarchy for routing in very large networks. *Computer Communication Review* 18, 4 (August 1988), 35-42.
- [11] ZEGURA, E. W., CALVERT, K., AND BHATTACHARJEE, S. How to model an internet work. In Proceedings of IEEE INFOCOM (1996).
- [12] 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, University of California at Berkeley, Computer Science Division, April 2001.
- [13] Bloom, B. Space/time Trade-offs in Hash Coding with Allowable Errors. *Communications of the ACM*, 13 (7). 422-426.
- [14] Mitzenmacher, M., Compressed Bloom Filters. in *Twentieth ACM Symposium on Principles of Distributed Computing (PODC 2001)*, (Newport, Rhode Island, 2001).