

# A Study of the Performance Potential of DHT-based Overlays

Sushant Jain    Ratul Mahajan    David Wetherall

University of Washington

## Abstract

We study whether overlays based on the recent distributed hash tables (DHTs), which have attractive scaling properties, can deliver performance comparable to that of overlays based on measurements, which are constructed specifically to provide good quality paths. Our work is motivated by the use of DHTs for services such as multicast. We observe that multicast is already targeted by measurement-based overlays, and yet despite this convergence of purpose there is little understanding of how the two approaches compare. We simulate current and idealized versions of DHTs to highlight their potential for improvement over time. We compare three DHT-based overlays, CAN, Chord and Pastry, with Narada and NICE, which are representatives of measurement-based overlays, and with power-law random graphs (PLRGs) representing Gnutella.

We find that when configured with the same average out-degree, basic versions of DHTs have a latency stretch that is more than NICE and Narada by at least a factor of two, but have similar performance in terms of bandwidth hotspots. PLRGs perform well in terms of latency stretch but poorly in terms of bandwidth hotspots. We also find that considerable performance gain can be achieved in DHTs with routing heuristics and topology-aware overlay construction. Together, these two techniques have the potential to bring their performance at par with NICE, and hence are a promising direction.

## 1 Introduction

Overlays have become the preferred vehicle for providing new Internet services (e.g., CDNs [1, 12], application-level multicast [17, 4, 18, 9, 27, 7, 38, 21, 34], and P2P file sharing [19, 14, 10, 23, 11, 13]). As such, overlay construction protocols that provide good levels of performance, scalability, and robustness are of considerable importance. There has been a surge of interest in the area, along with rapid advances that have led to two main approaches for overlay construction: protocols based on measurements, and protocols based on the recently developed distributed hash tables (DHTs).

Measurement-based protocols, initially motivated by application-level multicast, use estimates of net-

work properties such as latency between overlay nodes to make an informed choice of overlay structure. Examples include Narada [17], RON [2], NICE [4], TAG [21] and Gossamer [9]. These algorithms provide good performance via high quality paths. Recent extensions have allowed them to scale up to tens of thousands of nodes by using hierarchy [4, 20].

On the other hand, DHT-based protocols begin with structure in mind, and may fold measurement information into that structure. They map the overlay nodes to a virtual space commonly known as the node identifier space, typically pseudo-randomly to achieve a balanced distribution. The overlay topology, which determines how the nodes connect to each other, is governed mainly by these node identifiers. Examples include CAN [25], Chord [32], Pastry [29] and Tapestry [37]. These algorithms are extremely scalable and were motivated by peer-to-peer (P2P) applications such as distributed file sharing, where millions of nodes may be involved.

While originally developed for different purposes, DHT-based overlays are now being targeted at some of the same applications as measurement-based overlays, most visibly application-level multicast [27, 7, 31, 38]. Here, efficient use of the network is a key concern, more so than for earlier DHT applications such as distributed indexing. Much recent work thus aims to better the performance of DHT-based overlays with improved construction heuristics [5, 36, 35, 26, 33]. However, despite this convergence of purpose and plenitude of work, there is no real understanding of how the two approaches compare.

In this paper, we study the performance potential of DHT-based overlays. We seek to determine whether DHT-based overlays can deliver performance that is comparable to measurement-based overlays, and thus whether research on improved heuristics is likely to be fruitful. Our approach is to sidestep ongoing improvements to DHTs, and rather than report on many proposed versions of DHT-based overlays that may quickly become outdated, we simulate basic and idealized DHT variants that use global knowledge to understand their potential

for improved performance as better heuristics are discovered. Through simulation, using the same topologies, scale and metrics, we compare CAN, Chord and Pastry with Narada and NICE, which are representatives of measurement-based overlays. We also study power-law random graphs (PLRGs) with flooding-based routing, representing Gnutella, to provide another point of comparison. Our study was conducted on a scale of 1000s of nodes, as against millions of nodes, which is the stated target of most DHTs. This scale is achievable by all of measurement-based overlays, DHTs and PLRGs, and in our experience even this scale provides useful insights into comparative performance of these overlays. To our knowledge, this is the first “apples to apples” comparison of these approaches.

We find that when configured with the same average out-degree, basic versions of CAN, Chord and Pastry have a latency stretch longer than NICE and Narada by a factor of two or more, depending on the scale. Somewhat surprisingly, all these algorithms showed similar performance in terms of bandwidth hotspots. PLRGs performed better than the basic versions of DHTs in terms of latency stretch due to the use of flooding as the routing mechanism in PLRGs, but poorly in terms of bandwidth hotspots due to highly variable node out-degrees. We also find that considerable latency performance gain can be achieved in DHTs with better routing heuristics and topology awareness. Together, these techniques have the potential to bring DHT performance at par with NICE, and thus are a promising direction. But it was more difficult to simultaneously achieve both good latency and good bandwidth performance for DHTs. As others [4, 20] we find that the hierarchy, which is used in NICE to help it scale, does not significantly degrade performance compared to Narada.

The paper proceeds as follows. The next section briefly describes the overlay algorithms we study. Section 3 discusses the metrics of interest when evaluating overlays, and Section 4 describes our experimental methodology. We present our results in Section 5, discuss related work in Section 6, and conclude in Section 7.

## 2 Background

This section provides an overview of the overlay construction algorithms that we study. An overlay is built by forming virtual links or tunnels between the participating nodes; a tunnel between a pair of nodes usually traverses multiple links in the underlying network. Given a set of nodes, the goal of an

overlay construction algorithm is to select the virtual links that will be used as the overlay topology, and to determine how to route over that topology.

### 2.1 Measurement-based Overlays

Measurement-based overlays are constructed primarily using active measurement of network properties such as latency between overlay nodes. Several algorithms for building these kind of overlays exist, most of which target multicast services [17, 4, 34, 24, 21, 18, 9]. We study Narada [17] and NICE [4], both of which are optimized for latency.

In Narada, the goal is to settle into a flat (no hierarchy) overlay topology that minimizes the latency between nodes while keeping a small number of tunnels per node. This is accomplished by choosing an initial set of neighbors, and periodically adding new tunnels and dropping existing ones. The tunnel addition and deletion process is governed by its utility to the two end points. The utility metric is computed using the reduction in distance to other nodes the tunnels brings about. This computation requires periodic exchange of routing tables between neighbors and every node probing every other node. This poses a barrier to scalability, but provides all nodes with near global knowledge leading to a highly optimized overlay.

To tackle the scalability problem with a flat measurement based overlay, NICE creates a hierarchy of node clusters. At the bottom of hierarchy, nodes are partitioned into clusters of fixed size. Each cluster has a representative node, roughly at the topological center of the cluster that is determined by nodes in the cluster probing each other. The representative node of a lower-level cluster is a member in the next level of the hierarchy. This process of partitioning into clusters and adding another level is recursively repeated, yielding a tree topology. The per node network bandwidth required for protocol maintenance in NICE is  $O(\log(n))$ , compared to  $O(n)$  in Narada, where  $n$  is the number of nodes in the overlay.

For the purpose of our study, Narada provides a baseline for an overlay with high quality paths at small scale. NICE provides an indication of the performance that can be maintained when additional structure is imposed to scale to larger sizes.

### 2.2 DHT-based Overlays

DHT-based overlays are constructed using algorithms for distributed hash tables (DHTs) [25, 32, 29, 37]. These were originally developed to provide highly scalable and fully distributed indexing services for peer-to-peer file sharing, but have

since been applied to other services such as multicast [27, 31, 7, 38]. We study CAN [25], Chord [32] and Pastry [29] in this class of protocols.

In CAN, nodes are mapped pseudo-randomly to a virtual  $d$ -dimensional Cartesian space, which wraps around at the edges and has no resemblance to the underlying physical topology. Every node has  $2 \times d$  neighbors, corresponding to the adjacent nodes in each dimension. Routing is achieved by following a path through the Cartesian space that increasingly progresses from source to destination. The average path length is  $O(dn^{1/d})$ , where  $n$  is the number of nodes in the overlay and the per node network bandwidth required for protocol maintenance is  $O(d)$ .

In Chord, every node is pseudo-randomly assigned an  $m$  bit identifier. For simplicity of explanation, assume  $n = 2^m$ , where  $n$  is the number of nodes in the overlay. Every node is connected to  $m$  neighbors (i.e.  $\log(n)$  neighbors) with identifiers that are spaced at distances of  $2^0, 2^1, 2^2 \dots 2^{m-1}$  from its identifier in identifier space using modulo arithmetic. A packet from node  $n_s$  to node  $n_d$  is forwarded to the neighbor of  $n_s$  that is arithmetically closest to  $n_d$  but less than or equal to  $n_d$ . The above process ensures a route with maximum length  $\log(n)$  between any two nodes. The per node network bandwidth required for protocol maintenance for Chord is  $O(\log(n))$ .

In Pastry, nodes are pseudo-randomly mapped to an  $m$ -bit identifier space in base  $2^b$ . The routing table of a Pastry node is a matrix with  $m/b$  rows, and  $2^b$  columns. The node in cell  $(r, c)$  shares the first  $r$  digits with the local node and has the last digit equal to  $c$ . Routing is accomplished by each node along the way forwarding a message for key  $k$  to the node in its routing table with the longest matching prefix; ties are broken using arithmetic closeness to  $k$ . Both the average number of hops required to route a message and per node bandwidth required for overlay maintenance in Pastry is  $O(\log(n))$ .

What we have described above are the basic versions of CAN, Chord and Pastry. Several modifications have been proposed for both the algorithms [25, 32, 26, 11, 5]. Some of these target performance, while others target application-specific aspects such as data availability, resiliency and hotspot management. Since the primary focus of our work is performance, we study only the performance enhancing heuristics (Section 4.1) and ignore the rest.

CAN, Chord and Pastry represent different families of algorithms, all of which are being actively developed [27, 26, 31, 11]. Tapestry [37] is similar to

Pastry, and we believe our results are relevant for Tapestry as well.

### 2.3 Random Power-Law Overlays

We study one more class of overlays: power-law random graphs (PLRGs). It has been observed that the topology of “naturally emerging overlays” such as Gnutella, which are formed when users create links to other nodes, are characterized by a node out-degree distribution that obeys a power law. Routing in such overlays is accomplished through flooding over all links. These overlays are attractive for their simplicity and high tolerance to random failures [30]. We study them to provide another point of comparison in the space.

### 3 Overlay Performance Metrics

In this section, we describe the metrics used in our study. Defining appropriate metrics for comparing overlays is not straightforward because the applications they are designed for may have conflicting requirements. However, the performance of an overlay with respect to its latency overhead and its use of network resources is always a concern. For this reason two criteria have been widely used in overlay evaluation: relative delay penalty (RDP) and link stress [17, 4, 27, 31, 18, 5, 7]. We use both of these, along with a third, *Load Balance Ratio*, that measures the distribution of the routing responsibility in the overlay.

Ideally, an overlay would deliver latencies and bandwidths between nodes that match those of the underlying network. However, there is a penalty for routing between overlay nodes. If there is no virtual link between two overlay nodes, the path between them through the overlay will be longer than the path through the underlying network. If several virtual links pass over an underlying physical link, the link will experience higher load and in case of multicast the same message may travel over it multiple times. This latency overhead is measured using RDP, and the bandwidth penalty using link stress.

#### 3.1 Relative Delay Penalty

Relative delay penalty (RDP) is a measure of the additional packet delay introduced by the overlay on the delivery of a message between two nodes. It is defined as the ratio of the latency experienced when sending data using the overlay to the latency experienced when sending data directly using the underlying network.

As defined, RDP is measured between a pair of nodes and thus provides a set of values across various pairs. To gauge the level of performance of the

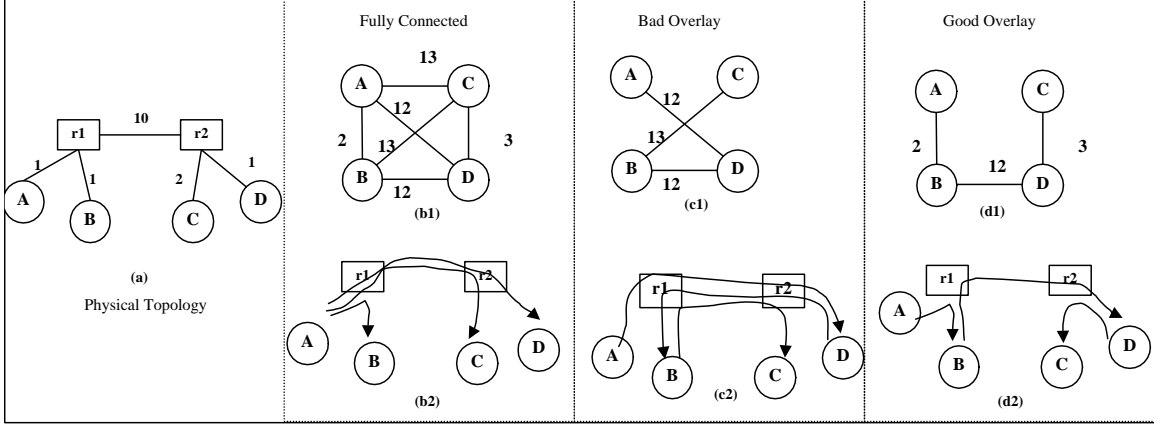


Figure 1: Example overlay topologies. The physical network is shown in the first pane followed by three possible overlays. The logical topology is shown on the top, and the paths taken when A sends a message to B, C and D is shown on the bottom.

entire overlay, we use the 90th percentile from the RDP distribution. Various authors have used the same criteria for several reasons [17, 31, 27]. First, the 90th percentile serves to bound the delay multiplier that will be seen by most nodes in practice, whereas the average by itself does not convey any sense of the variation. Second, choosing the 90th percentile rather than the worst case RDP hides sensitivity to simulation parameters. Finally, RDP can easily be very high when the physical latency is small, and using the 90th percentile instead of worst case RDP filters out these outliers [17].

### 3.2 Link Stress

Link stress is defined as the number of tunnels that send traffic over a physical link. Links with high stress are potential bandwidth hotspots in the system. In case of unicast, it is a measure of load experienced by the link; more tunnels means higher load. In case of multicast, it becomes a measure of excess bandwidth consumption induced by the overlay; more tunnels means more duplicate messages. For multicast, stress is a function of both the topology and the multicast tree. Two approaches to implement multicast exist – flooding [27] and tree-based [31, 7, 38]. For DHTs and measurement-based overlays, we consider only the latter as it is known to be more efficient [8]. Here the multicast tree is the union of the unicast routes from the source to all destinations. For PLRGs we use flooding, the default mechanism.

Stress as defined above is a distribution over all physical links for each multicast tree i.e per source. To gauge the level of performance of the entire overlay, we use the 90th percentile of the distribution given by the worst stress for each multicast tree. Using worst case distribution conveys a bound on the

stress seen by any link, and using the 90th percentile reduces the sensitivity to simulation parameters.

An ideal overlay should have both low RDP and low stress. Unfortunately, it may be difficult to simultaneously achieve both objectives. For instance, consider the physical network and three possible overlays topologies shown in Figure 1. Figure 1(b1) shows a fully connected overlay topology, in which the RDP between all pairs is 1 but the stress on links close to the end hosts is high. To communicate with B, C, and D simultaneously, A must send three packets over its physical access link, leading to a stress of 3. Next consider the overlay topology in Figure 1(c1), which has low stress on links close to the end hosts. In this case, all overlay tunnels go over the high latency physical link  $r1-r2$ , leading to high RDP between most pairs. Figure 1(d) shows a good overlay with both acceptably low RDPs and low stresses.

### 3.3 Load Balance

In an overlay, nodes also act as routers and forward packets between other nodes. Ideally an overlay would not place a much higher forwarding load on one node compared to other nodes<sup>1</sup>. Otherwise, with increasing workload the overloaded node would soon become a hotspot for the system, leading to degraded performance. For example, logical topologies such as a star or a tree are not well balanced compared to a ring topology because they have key nodes that are used for most point-to-point communications.

<sup>1</sup>Here we are assuming homogeneous nodes. Although this may not be a realistic assumption in practice, existing protocols are not designed to take into account heterogeneity of members [28]. Any variation in forwarding load is thus unintended.

To measure the distribution of the forwarding load, we define a metric called the *Load Balance Ratio*, which is computed as follows. For each node in the overlay, the routing load is the number of source-destination pairs between which it forwards messages. Load balance ratio is the ratio of the maximum routing load to the median routing load. This measures how much worse the maximally loaded node is compared to the halfway loaded node. This metric is less relevant for multicast (where all nodes forward one message for a given source). Rather, it assesses the performance of overlays when used more generally for multiple point-to-point communications.

For the purpose of our study, load balance ratio exposes how performance-based routing concentrates traffic in non-uniform ways. One of the favorable arguments behind DHT-based approaches is that they are better load balanced because of their regular geometric structure and use of randomization. But heuristics that improve performance can interfere with this. We also note that load balance distinguishes topologies and routing protocols that are only suited for multicast from those that are more widely applicable. For example, protocols such as NICE that use hierarchy are a good fit for multicast but would place extremely high load on nodes high in the tree if used for general unicast communications.

### 3.4 Other Considerations

There are several other performance measures that we do not explore in this paper. We do not measure the overhead of the overlay protocols, either in terms of the amount of state or traffic that is needed to maintain the overlay. It is well known that pure DHT algorithms make only local measurements and scale extremely well, while schemes such as Narada perform global measurements and scale relatively poorly. Our emphasis instead is on the performance levels that can be achieved at a given scale for which the overhead of the algorithms under study has been deemed acceptable. For example, NICE is able to scale to 1000s of nodes (because its overhead is logarithmic with overlay size), while Narada is not. Thus, for overlays this large, NICE is our only option among measurement-based overlays for comparison with DHTs.

We also do not measure protocol dynamics, such as maintaining connectivity in face of failures. Overlays are expected to operate with members leaving dynamically, and different overlay construction algorithms may be disrupted in different ways; in the

extreme, partitions are possible. While these dynamic properties are important, our first priority is to understand the static performance potentials of the different algorithms.

## 4 Methodology

In this section, we describe our experimental methodology. We first describe the variations of the DHT-based algorithms that we compare. We then describe our simulation set-up.

Recall that our goal is to understand whether DHT-based overlays will be able to match the level of performance of measurement-based overlays, which are constructed specifically to provide good quality paths. One difficulty is that different heuristics are being continuously proposed to enhance the performance of DHT-based overlays [26, 25, 35, 31, 33, 5], and we do not wish our results to quickly become irrelevant by being tied too closely to specific heuristics. Instead, we side-step this race by taking advantage of simulation as a tool to report on the performance of both the current versions and idealized versions with performance enhancing heuristics that use global knowledge.

### 4.1 DHT Heuristics

We study performance-enhancing variations to DHTs along two dimensions – those that attempt to find better paths over a given overlay, and those that construct the overlay itself in a topology-aware manner. We describe each in turn.

#### 4.1.1 Routing Heuristics

In the simplest version of DHTs, routing proceeds using only the geometric properties of the overlay algorithm. We refer to this as the *Base* variant. It is possible to improve performance by routing across the overlay in a manner that takes latency into account. The structure of the overlay itself remains unaffected. We study the heuristic specified in [25] for CAN and [11] for Chord. Here, the chosen next hop is the one out of all possible neighbors that results in the maximum progress towards the destination, where progress is defined as the ratio of the distance in identifier space to the physical latency. We refer to these versions of CAN and Chord as *Proximity* variants. In Pastry, the next hop is unique and fixed and therefore it does not have a corresponding Proximity routing variant. Comparison between *Base* and *Proximity* shows the value of the current routing heuristics.

Other heuristics have been proposed, e.g., smallest physical latency for CAN is described as an improvement in [8]. However, rather than simulate

all the different heuristics we can find, we stick with the above as a yardstick (since it applies to both CAN and Chord and is the subject of most published results) and define a new variant intended to provide an upper bound on how well any future heuristic can perform. This variant, *Shortest-Path*, comes from the observation that the maximum gain achievable by any routing heuristic is that of shortest path routing (implemented using distance vector or link state algorithm). This algorithm requires global knowledge, unlike proposed heuristics, and so may not be a good choice at large scales. However, it can readily be computed in our simulation setting to provide a bound on the performance of better heuristics that will inevitably be proposed. That is, comparison between *Proximity* and *Shortest-Path* shows how much room routing heuristics have for improvement. Another advantage is that the *Shortest-Path* variant is independent of the actual overlay algorithm and therefore it applies to all three DHTs.

#### 4.1.2 Topology Awareness

In the simplest version of DHTs, both mapping of nodes to identifiers and choice of tunnels (when multiple choices are possible as in Pastry) is pseudo-random to provide a well-balanced structure. This overlay construction process, which we refer to as the *Random* variant, is unaware of the underlying topology.

In a *topology-aware* overlay heuristics are used to create an overlay that reflects the underlying network topology. The intuition is that if the overlay and the underlying network topology closely mirror one another, then the overlay paths will closely follow network paths and good performance will result. We now describe how to achieve topology awareness in Pastry, CAN and Chord.

In Pastry *topology-awareness* can be achieved through neighbor selection [5]. A routing table entry can be filled by any node that matches the criteria for that cell. Topology-aware construction chooses the closest such node. Achieving perfect awareness using this method requires global knowledge, though reasonable approximations that rely only on limited information exchange are possible [5]. In this paper, we only consider the global knowledge approach.

The above neighbor-selection approach is not applicable to CAN and Chord because in both these algorithms topology is completely defined once the identifier assignment is done. Instead topology awareness is achieved through intelligent identifier assignment. For instance, identifier assignment based

4	3	2	X
X	X	1	X
X	X	X	X
5	X	X	X

(a)  
*Random*

4	3	2	
	5	1	X
		X	

(b)  
*Topology aware*

Figure 2: Topology awareness in a 2-d CAN. A new node 5, which is physically closest to 1, joins the overlay. X denotes the possible assignments for 5 in case of *random* and *topology-aware* overlay construction. The *topology-aware* case yields an assignment in which 5 is geometrically close to 1.

on proximity to landmarks has been proposed as a heuristic [26] to achieve a topology aware mapping in CAN.

However, it is generally the case that heuristics for topology-aware identifier assignment are not as well defined nor studied as heuristics for improving routing. To understand how much topology awareness itself can improve performance, assuming that good heuristics will be found we would like to use an analogue to our Shortest-Path variation above, rather than simulate many possible heuristics. However no such globally optimal assignment has been defined for either CAN or Chord. So, we use a greedy assignment based on global knowledge to construct a good assignment. Our greedy assignment<sup>2</sup> rule works as follows: the identifier of a new node joining the overlay is chosen so that it is a neighbor of the node that is closest to it in the underlying network. That is, the rule makes the overlay neighbors of a node similar to the neighbors of the node in the underlying topology. Figure 2 illustrates the concept for a 2-dimensional CAN. This also requires global knowledge (or at least a knowledge of distances to all nodes in the vicinity) and so may not be a good implementation choice at large scales. Interestingly, a similar heuristic has been proposed in parallel with our work [33].

#### 4.2 Simulation Setup

Since our primary interest is in understanding performance under static environments, we simulate overlay construction using centralized algorithms.

<sup>2</sup>We believe that our approach based on global knowledge would yield a better mapping than using a fixed number of landmarks. The performance of the latter is also sensitive to the choice and number of landmarks.

The only exception is Narada, which we simulated using an event-driven simulator because Narada does dynamic evaluations to improve the overlay over time. To generate PLRGs overlays, we first generated power-law topologies using Brite [22], and then randomly mapped the nodes in this graph to nodes in overlay. For DHTs pseudo-random hash functions are used to provide a well-balanced structure. However, since we are using simulation we created well-balanced structures directly by a priori partitioning the space uniformly. This is a conservative simplification that is consistent with our goal of presenting the DHT-based overlays in their best light.

An important parameter that affects the performance is the average out-degree (number of neighboring nodes) of a node. It determines number of links in the overlay and therefore directly affects performance. For example, an overlay with more links will have lower RDPs because of shorter paths but higher stress because more overlay links traverse a physical link. We study the effect of out-degree in Section 5.4. The average out-degree is a configurable parameter in all overlays we study except Chord. In Chord the outdegree is  $\log_2(\text{overlay\_size})$ . In Pastry the average degree depends on  $b$  and varies as  $b \log_2^b(\text{overlay\_size})$ . We fixed  $b$  as 1 to make Pastry’s average degree comparable to that of Chord. The values of degree we study range from 4 to 12.

We used the transit-stub model of GT-ITM [15] to generate the physical network topology. GT-ITM also assigns latencies to links in the physical topology. Additional nodes were attached to the stub nodes to represent hosts connected to lower level routers. Overlay nodes were picked randomly from these hosts.

The overlay algorithms that we compare are summarized in Table 1. We study one version of each of Narada [17], NICE [4], and PLRG. Narada and NICE build their topologies using latency measurements as described earlier, and they use shortest path routing using latency as the metric. PLRG forms a random topology where the node outdegree distribution follows a power law, and then flood over this topology; this represents Gnutella-like overlays.

For each parameter setting we ran 9 simulations – three different physical network topologies, each with three random seeds. Each topology had 4,040 backbone nodes and 20,200 stub nodes. Simulated overlay sizes were varied from 64 to 4096.

	Topology Awareness	Routing	Degree
Narada	-	<i>Shortest-Path</i>	4 – 12
NICE	-	<i>Shortest-Path</i>	4 – 12
CAN	<i>Topology-Aware, Random</i>	<i>Shortest-Path, Proximity, Base</i>	4 – 12
Chord	<i>Topology-Aware, Random</i>	<i>Shortest-Path, Proximity, Base</i>	–
Pastry	<i>Topology-Aware, Random</i>	<i>Shortest-Path, Base</i>	–
PLRG	-	<i>Flooding</i>	4 – 12

Table 1: Summary of comparisons performed.

## 5 Results

We now present the results of comparing different protocols and their variants on each of our three performance metrics. We first consider a comparison using an out-degree of 10 for all overlays except Chord and Pastry. An out-degree of 10 was found to be a good representative by experimenting with different out-degrees. The effect of changing out-degree on performance is considered in Section 5.4.

### 5.1 RDP

This section investigates the latency stretch of various overlays. We start with studying the impact of using heuristics in DHT-based overlays, and then compare all the overlays.

#### 5.1.1 Effect of Heuristics

Figure 3 and 5 shows the impact of routing heuristics in CAN on top of an overlay with *random* and with *topology-aware* overlay construction respectively. They plot the 90th percentile RDP for the *Base* version of CAN, CAN with *Proximity* routing, and CAN with *Shortest-Path* routing (best possible performance for a given overlay). Figure 4 and 6 shows the same results for Chord. Results from all 9 simulations are presented to show the variance and the line is drawn through the average. This representation is used in all plots unless otherwise specified.

For both CAN and Chord, improved routing brings about significant improvement in RDP (similar results were obtained using 50th percentile and 95th percentile RDP). *Proximity* routing does quite well; it reduces the RDPs to roughly halfway between *Base* and *Shortest-Path* routing. However, it does not match the performance of *Shortest-Path* routing because while the former is a greedy decision at each hop, the latter computes globally optimal

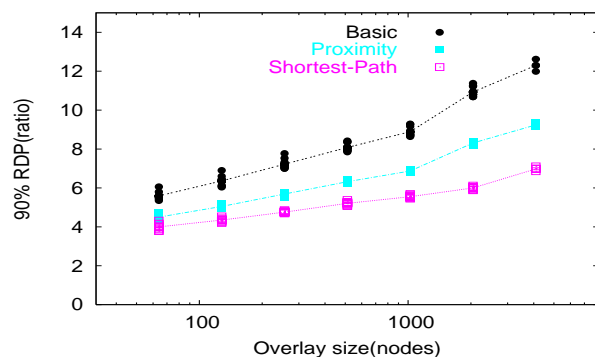


Figure 3: Effect of routing heuristics on RDP in CAN with *random* overlay construction.

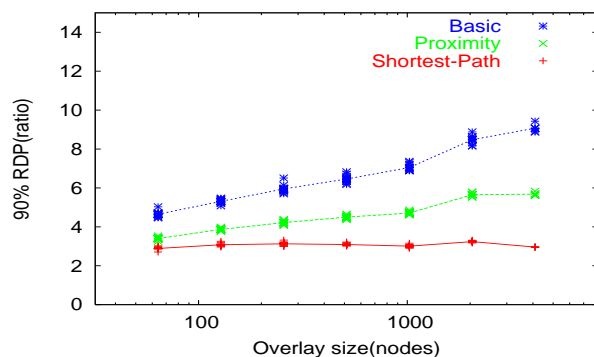


Figure 5: Effect of routing heuristics on RDP in CAN with *topology-aware* overlay construction.

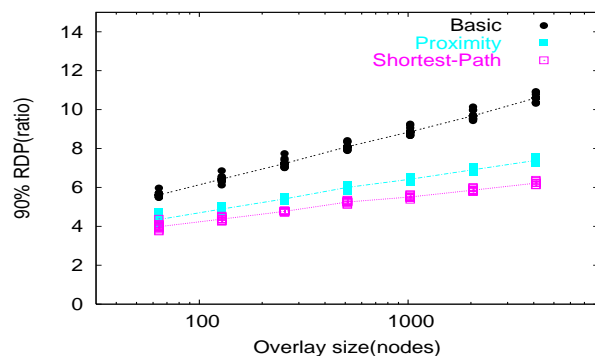


Figure 4: Effect of routing heuristics on RDP in Chord with *random* overlay construction.

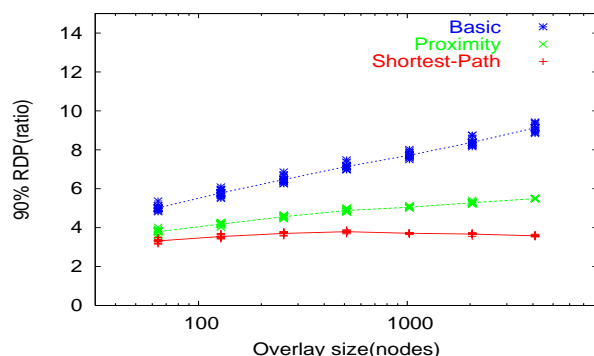


Figure 6: Effect of routing heuristics on RDP in Chord with *topology-aware* overlay construction.

Routing	Topology Awareness	
	Random	<i>topology-aware</i>
Base	<b>8.89</b>	7.05
Proximity	6.87	4.71
Shortest-Path	5.55	<b>3.02</b>

Table 2: Effect of heuristics on 90% RDP for 5-dimensional CAN with 1024 nodes

paths. The improvement in Chord using *Proximity* routing is slightly greater because as we increase the overlay size, the number of choices for the next hop increases, leading to better paths. As a result *Proximity* routing comes closer to *Shortest-Path* routing in Chord than in CAN.

To understand the effect of the composition of the different heuristics, we tabulate the 90th percentile RDP for a 5-dimensional CAN with 1024 nodes in Table 2. Observe that the effect of heuristics compose and the best performance is achieved by enabling both *Shortest-Path* routing and *topology-*

*aware* overlay construction. This improvement is substantial, from 8.89 to 3.02 (almost 70% reduction), and indicates the potential for improvement through more practical heuristics. By looking across columns in the table, one can deduce that being topology aware by itself brings about a significant performance gain.

Figure 7 shows the effect of both *topology-aware* construction and *routing* heuristics on RDP for Pastry. As before heuristics lead to substantial improvement.

### 5.1.2 Comparing All Protocols

We try to answer two questions now. First: How do DHTs with only scalable heuristics (such as *Proximity routing*) compare to measurement-based overlays, especially NICE since it has similar scalability? Second: How do DHTs compare to the other approaches when *Shortest-Path* routing along with *topology-aware* overlay construction is implemented?



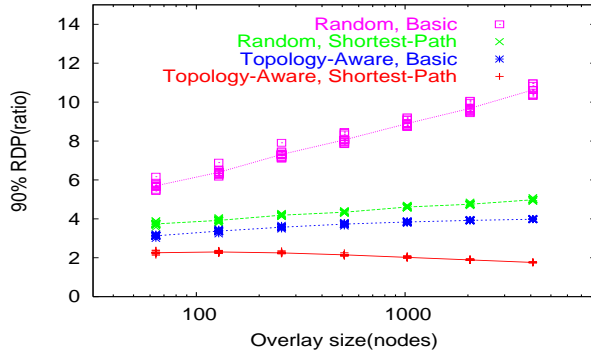


Figure 7: Effect of heuristics on RDP in Pastry.

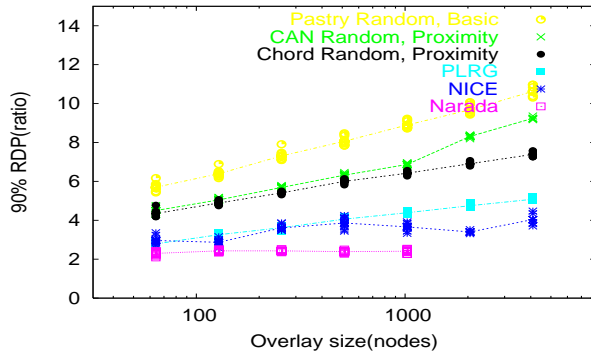


Figure 8: Variation of RDP with size. Simple practical versions of DHTs are shown – CAN and Chord with *random* overlay construction and *Proximity* routing, and Pastry with *random* overlay construction and *base* routing.

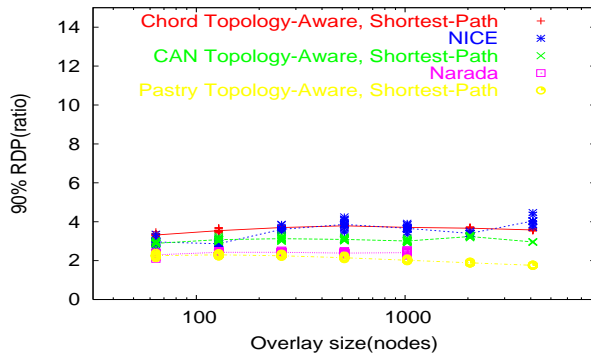


Figure 9: Variation of RDP with size for various overlays. Idealized versions of DHTs – with *topology-aware* overlay construction and *shortest-path* routing – are shown.

Figure 8 shows the 90th percentile RDP as a function of overlay size for different protocols<sup>3</sup>. *Proximity* routing is used for CAN and Chord with *random* overlay construction. For Pastry we show *random* overlay construction with *base* routing. These represent simple, efficient, and practical versions of these overlays. A distributed implementation of topology-awareness is shown to be a reasonable approximation in [6]. We believe that by the same token, similar implementations are possible for CAN/Chord because a new node being able to find the closest live node is a key assumption in both scenarios. Instead of assuming this to be a fact and for fairness of comparison, we chose to use the *random* overlay construction for all DHTs.

This is because as indicated in [6] topology-awareness in Pastry can be achieved in an efficient manner without having much affect on performance.

We make the following observations:

- The difference in RDP between NICE and the DHTs is large and grows with the overlay size. It is a factor of two for 1024 nodes.
- Flooding over PLRGs performs well, because with flooding the shortest-path between nodes is taken. This also indicates that without heuristics DHT overlays are not better than random overlays from a latency perspective.
- All DHTs have similar RDPs. The slightly worse performance for Pastry stems from the absence of a *Proximity* routing equivalent, and the slightly better performance of Chord for overlay sizes greater than 1024 stems from a higher average out-degree in Chord compared to CAN beyond this size,<sup>4</sup> which leads to shorter paths.
- The performance of NICE does not deteriorate with increasing overlay size even though the levels of hierarchy increase. Further, it is not much worse than that of Narada. Our findings agree with those of the authors of NICE [4].

Note that although Narada has the best RDPs, it is not scalable and therefore not applicable for larger overlays.

Figure 9 compares DHTs with *topology-aware* overlay construction and *Shortest-Path* routing with all other protocols. This represents the maximum performance one can achieve from a latency perspective

<sup>3</sup>For Narada, results are not shown for overlay sizes greater than 1024 nodes because simulation times were on the order of days.

<sup>4</sup>Average degree for Chord is  $\log(n)$  – for 1024 nodes average degree is 10 which is the same as the average degree configured for CAN.

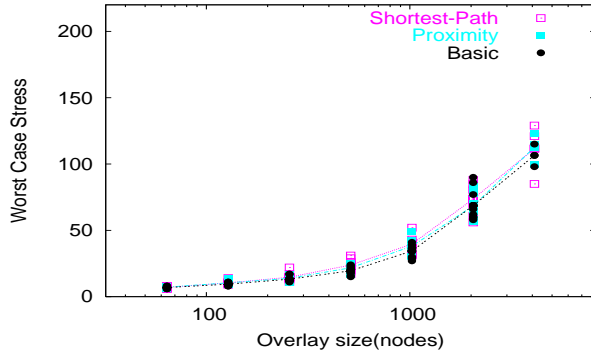


Figure 10: Effect of routing heuristics on link stress in CAN with *random* overlay construction.

for these overlays. There is no real qualitative difference between performance in this case. This is encouraging because it shows that RDPs comparable to measurement-based overlays can be achieved using improved versions of DHT-based overlays.

## 5.2 Stress

We now investigate how the various protocols compare on the metric of the worst case stress, and the impact on stress of RDP-enhancing heuristics in DHT-based overlays.

### 5.2.1 Effect of Heuristics

Figures 10 and 11 show the stress as a function of overlay size for 5-dimensional CAN with *random* and *topology-aware* overlay construction respectively. The three lines correspond to the three routing variants – *Base*, *Proximity* and *Shortest-Path*. The results for Chord were similar, and have been omitted due to space constraints.

Somewhat surprisingly, improved routing does not have much negative impact on stress. With improved routing, we expected stress near few well-placed nodes to go up. This is evident to some degree in the fact that the worst case stress does go up slightly. But at the same time, this increase is reasonably countered because improved routing leads to shorter paths, which means that fewer links are traversed. By comparing the results across *random* (Figure 10) and *topology-aware* (Figure 11), we can also see that topology awareness has little impact on stress.

Figure 12 shows the effect of heuristics on link stress for Pastry. As for CAN, *routing* heuristic had no significant impact on stress. However, *topology-awareness* with *base* routing had much worse (almost factor of 2) stress values. A key difference

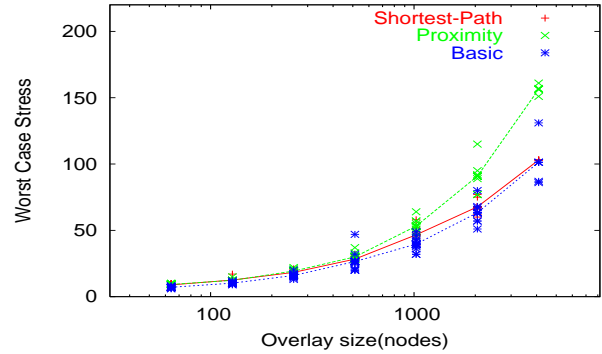


Figure 11: Effect of routing heuristics on link stress in CAN with *topology-aware* overlay construction.

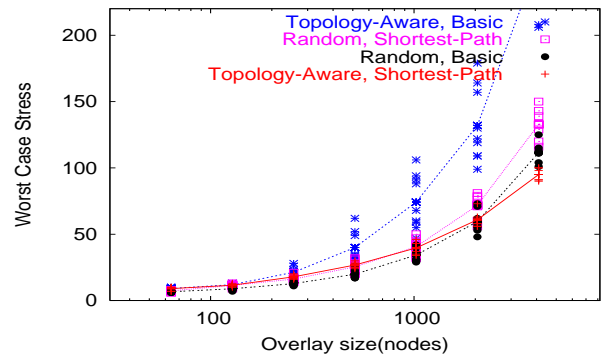


Figure 12: Effect of heuristics on link stress in Pastry.

between the topology-aware overlay construction mechanisms of Pastry and CAN/Chord is that in Pastry a few well placed nodes can be the neighbors of many nodes, whereas in CAN/Chord a node can be a neighbor of only a small number of other nodes. This has a direct consequence on link stress and load balancing. Hence, our results should be interpreted in terms of difference between the two methods to achieve topology-awareness – choosing node identifier assignment and choosing neighbors.

### 5.2.2 Comparing All Protocols

Figure 13 shows how stress varies with size for different algorithms. Since the performance of various variants of DHTs were largely similar, we show the stress only for the best version – *topology aware* overlay construction with *Shortest-Path* routing. The striking artifact in the graph is that PLRG has very high stress, more than five times worse than the other protocols for 1024 nodes. This is due to its use of flooding as the routing mechanism, and an uneven out-degree distribution among nodes.

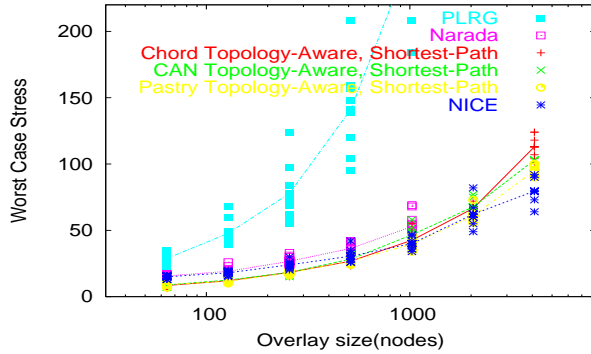


Figure 13: Variation of stress with size for various overlays. Variants of DHT with *topology aware* overlay construction and *shortest-path* routing are shown.

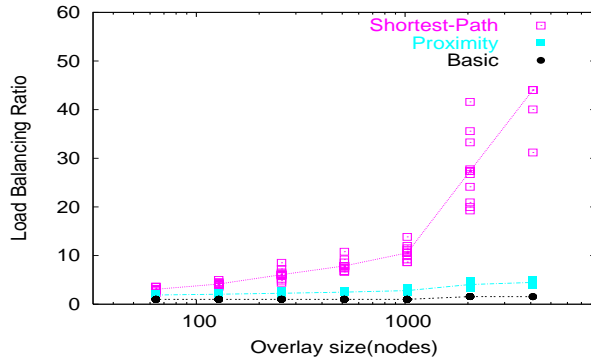


Figure 14: Effect of routing heuristics on load balancing in CAN with *random* overlay construction.

All other protocols exhibit similar stress values, although for large overlay sizes (over 2048 nodes) NICE has slightly smaller values. Large group sizes have higher density for a fixed topology size, which has the tendency to increase the stress on the backbone links in DHTs. In this situation, clustering of close nodes in NICE helps in reducing the stress on the backbone links.

In summary, for similar average out-degrees all protocols exhibited similar worst case stress properties. In case of CAN and Chord heuristics have no significant impact on stress properties. For Pastry however, *topology-awareness* with *base* routing had significantly higher stress than other variants.

### 5.3 Load Balancing

We now study the load balancing properties.

#### 5.3.1 Effect of Heuristics

Figure 14 shows the load balancing ratio for different overlay sizes for a 5-d CAN with *random* overlay

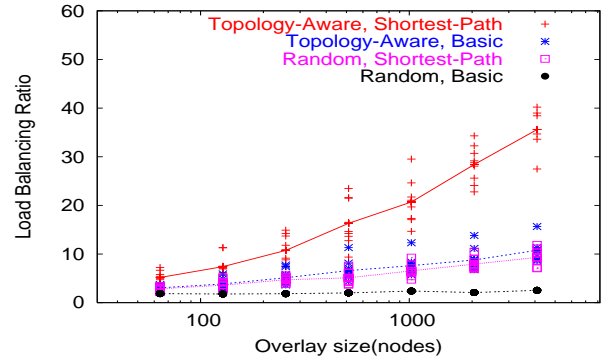


Figure 15: Effect of heuristics on load balancing in Pastry.

construction. As before, the three lines corresponds to the three routing variants. With *Base* routing, the load is balanced very evenly, with the ratio between one and two. This is a direct consequence of the regular structure and the random overlay construction in the DHT algorithms. With the *Proximity* routing, load balancing deteriorates slightly but grows very slowly with the overlay size. However, there is a significant degradation with *Shortest-Path* routing because there exist well-placed nodes in the topology that have a low latency to many nodes, and therefore are used very often independent of their node-ID. This raises the issue of usefulness of a routing heuristic that mimics *Shortest-Path* routing if balanced load is desired.

Similar behavior was observed for CAN with *topology-aware* overlay construction because being topology aware only changes the relative assignments and does not modify the DHT structure itself. Results for Chord were also similar; we omit them due to space constraints.

Figure 15 shows the load balancing ratio for different heuristics. As in CAN we find significant degradation with *Shortest-Path* routing. Also we can see that in the case of Pastry *topology-awareness* also deteriorated load balancing properties. This is for reasons similar to those for higher link stress mentioned in Section 5.2.1.

#### 5.3.2 Comparing All Protocols

Figure 16 shows the load balancing ratio for all the protocols. The *Shortest-Path* routing variant for DHTs is shown as it had the highest load. Note that the scale of y-axis differs from that in Figure 14. We can see that NICE has extremely high load balancing ratio (two orders of magnitude for overlays

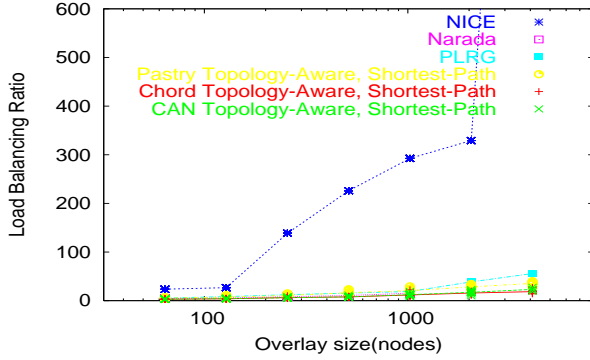


Figure 16: Variation of load balancing ratio with size for various overlays. Variants of DHTs with *topology-aware* overlay construction and *shortest-path* routing are shown.

bigger than 1024). This is because in NICE’s hierarchy, the *root* node is responsible for forwarding all packets whose source and destination lie in different sub-trees. On the other hand Narada, a measurement-based overlay with no hierarchy, performs similar to CAN and Chord with *Shortest-Path* routing. PLRG too has a high load balancing ratio because of highly uneven degree distribution. Note that the relatively sharper increase in NICE at some points (256 and 4096) is due to formation of unbalanced trees because of quantization effects.

In summary, our results point at the limitations of hierarchy for applications involving general unicast communications. DHT routing heuristics that mimic *Shortest-Path* routing can have significant negative impact on load balancing and thus may not be a suitable choice for some applications. Although *topology-aware* overlay construction by itself does not degrade load balancing in CAN/Chord, it degrades it in the case of Pastry.

#### 5.4 Effect of Out-Degree

In this section we study the effect of increasing the average out-degree of nodes in various overlays. Intuitively, increasing out-degree will decrease RDP because it shortens the number of overlay hops. But at the same time, stress could increase because the number of overlay links per physical link increases.

Node out-degree is controlled in CAN by varying the number of dimensions. We consider two versions of CAN: *i)* *Proximity* routing with *random* overlay construction; and *ii)* *Shortest-Path* routing with *topology-aware* overlay construction. Narada and PLRG algorithms can be specified an average out-degree, and out-degree in NICE is controlled

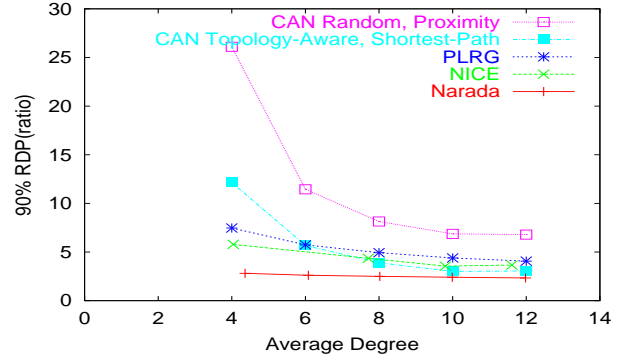


Figure 17: Effect of out-degree in various overlays. Two variants of CAN are shown – *random* overlay construction with *Proximity* routing and *topology-aware* overlay construction with *shortest-path* routing.

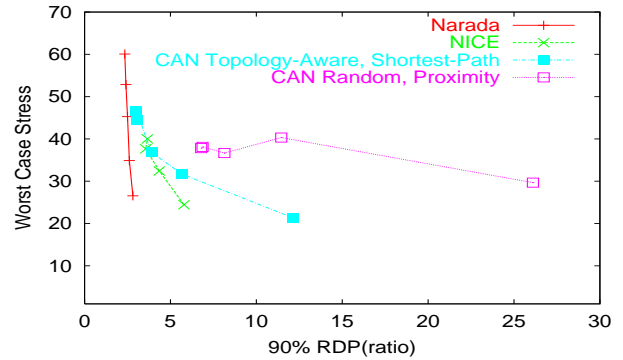


Figure 18: Variation of link stress with RDP. Two variants of CAN, same as those in Figure 17, are shown.

through cluster size. Chord and Pastry are not shown because one can not vary the out-degree of nodes independent of the overlay size.

The results in this section are presented for fixed overlay size of 1024 nodes and only averages over 9 simulations are plotted. Figure 17 shows the effect of increasing average out-degree on RDP for different protocols. Two observations can be made. First, there is a sharp reduction in RDP for CAN variants as the average degree is increased. For other protocols also, the RDP decreases with increasing degree but the gain is much less. Second, NICE was able to achieve good RDPs with a much lower average out-degree compared to CAN with *Proximity* routing over *random* overlay construction.

We now explore the trade-off between RDP and link stress. Figure 18 shows the relationship between

stress and RDP. It plots the measured stress for a given RDP obtained through varying the node out-degree (Figure 17). The worst case stress for PLRGs was extremely high, and hence has not been shown in this graph. All overlays exhibit the basic trade-off between RDP and link stress, though to varying degrees. By increasing average out-degree, RDP can be reduced, but that reduction comes at the cost of higher stress. However, note that the CAN variant with *Proximity* routing over *random* overlay construction lies further away from the origin than NICE. This means that it is harder to simultaneously achieve both low stress and low RDP in this variant of CAN than in NICE. At the same time CAN with *Shortest-Path* routing and *topology-aware* overlay construction performs comparably to NICE, which again points favorably towards the potential of optimization in DHT-based approaches.

## 6 Related Work

Many overlay construction schemes have been proposed, both measurement-based [17, 4, 21, 2, 24, 9, 34, 18] and DHT-based [25, 32, 29, 37]. Much ongoing work aims to improve the performance of DHT-based approaches [5, 36, 35, 26, 33, 31, 11] and the scalability of measurement-based approaches [4, 21]. At the same time, plenty of work in DHT-based approaches [27, 7, 38, 31] focuses on providing new services such as application-level multicast, for which the measurement-based overlays were initially intended. In this paper, we started with an observation of this convergence of purpose and scale, and investigated whether current and possible future versions of DHT-based overlays would be able to match the performance of measurement-based overlays.

Ratnasamy et. al outline the challenges facing DHT-based overlays [28]. They identify performance with respect to latency as a key open issue. Our work indicates that with good routing heuristics and topology-awareness, DHT-based overlays can match the performance of measurement-based overlays.

There has been very little work on studying how different overlay algorithms compare to each other. Our work and that of a few other researchers are first steps in this direction. Banerjee and Bhattacharjee [3] provide a survey of various overlay protocols. Our focus is however on empirical evaluation under identical environments. Castro et. al compare the performance of *tree-building* and *flooding* on top of CAN and Pastry [8]. They find that flooding has high overhead compared to tree-based approaches. Our results are similar in that we found that flooding on top of PLRGs has much higher overhead com-

pared to all other protocols we considered, which were tree-based.

## 7 Conclusions and Future Work

In this paper we studied the performance potential of DHT-based overlays by comparing their basic and idealized versions against measurement-based protocols. Specifically, we compared CAN, Chord and Pastry with Narada and NICE. Our key findings are:

- For the same average out-degree, basic versions of DHTs have a latency stretch that is longer than NICE and Narada by a factor of two or more, depending on the size of the overlay. The performance in terms of bandwidth hotspots is similar however.
- Considerable performance gains in latency can be achieved in DHTs with better routing heuristics and with topology-aware overlay construction.
- In most cases we observed no substantial effect of heuristics on link stress properties. The only case for which this was not true was with topology-aware overlay construction through latency sensitive choice of overlay tunnels, as is done in Pastry.
- Heuristics, especially routing heuristics, in DHTs lead to relatively higher load on nodes.
- Using routing heuristics and topology-aware overlay construction, the performance of DHTs can be brought at par with that of measurement-based overlays. This implies that DHT algorithms are a promising direction, with the potential to achieve not only scale but also performance.

As others [4, 20], we found that hierarchy does not significantly degrade performance from a latency perspective.

Our results are more applicable in scenarios where latency is the key optimization criteria. Comparison with bandwidth-oriented measurement-based protocols such as Overcast [18] and the enhanced version of Narada [16] would be required to understand whether DHTs, enhanced with bandwidth specific heuristics (we are not aware of any), can deliver comparable bandwidth performance. Overlay performance evaluation is in general sensitive to the choice of topology. Understanding the effect of choice of topology on performance from a comparison perspective is an open issue. Relative performance of these protocols under realistic dynamic environments such as rapid overlay membership changes is another front we hope to explore.

## References

- [1] Akamai home page. <http://www.akamai.com>.
- [2] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris. Resilient overlay networks. In *ACM SOSP*, 2001.
- [3] S. Banerjee and B. Bhattacharjee. A comparative study of application layer multicast protocols. In *Submitted for review*, 2002.
- [4] S. Banerjee, B. Bhattacharjee, and C. Kommareddy. Scalable application layer multicast. In *ACM SIGCOMM*, 2002.
- [5] M. Castro, P. Druschel, Y. C. Hu, and A. Rowstron. Exploiting network proximity in peer-to-peer overlay networks. In *FuDiCo*, 2002.
- [6] M. Castro, P. Druschel, Y. C. Hu, and A. Rowstron. Exploiting network proximity in peer-to-peer overlay networks. In *Technical Report MSR-TR-2002-82*, 2002.
- [7] M. Castro, P. Druschel, A.-M. Kermarrec, and A. Rowstron. Scribe: A large-scale and decentralised application-level multicast infrastructure". In *IEEE JSAC*, 2002.
- [8] M. Castro and et al. An evaluation of scalable application-level multicast built using peer-to-peer overlay networks. In *IEEE INFOCOM*, 2003.
- [9] Y. Chawathe, S. McCanne, and E. Brewer. An architecture for Internet content distribution as an infrastructure service, 2000.
- [10] I. Clarke, O. Sandberg, B. Wiley, and T. Hong. Freenet: A distributed anonymous information storage and retrieval system. In *Design Issues in Anonymity and Unobservability*, 2000.
- [11] F. Dabek, M. F. Kaashoek, D. Karger, R. Morris, and I. Stoica. Wide-area cooperative storage with cfs. In *ACM SOSP*, 2001.
- [12] Digital island home page. <http://www.digitalisland.com>.
- [13] R. Dingledine, M. J. Freedman, D. Hopwood, and D. Molnar. The free haven project: Distributed anonymous storage service. In *Design Issues in Anonymity and Unobservability*, 2000.
- [14] Gnutella. <http://www.gnutella.com>.
- [15] Gt-itm: Modeling topology of large internetworks. <http://www.cc.gatech.edu/projects/gtitm/>.
- [16] Y. hua Chu, S. G.Rao, S. Seshan, and H. Zhang. Enabling conferencing applications on the Internet using an overlay multicast architecture. In *ACM SIGCOMM*, 2001.
- [17] Y. hua Chu, S. Rao, and H. Zhang. A case for end system multicast. In *ACM SIGMETRICS*, 2000.
- [18] J. Jannotti, D. K. Gifford, K. L. Johnson, M. F. Kaashoek, and J. W. O'Toole. Overcast: Reliable multicasting with an overlay network. In *OSDI*, 2000.
- [19] Kazaa home page. <http://www.kazaa.com>.
- [20] D. Kostic and A. Vahdat. Latency versus cost optimizations in hierarchical overlay networks. Technical report, Duke University, 2001.
- [21] M. Kwon and S. Fahmy. Topology-aware overlay networks for group communication. In *ACM NOSS-DAV*, 2002.
- [22] A. Medina, I. Matta, and J. Byers. On the origin of power laws in Internet topologies. Technical Report 2000-004, Boston University, 2000.
- [23] Napster home page. <http://www.napster.com>.
- [24] D. Pendarakis, S. Shi, D. Verma, and M. Waldvogel. ALMI: An application level multicast infrastructure, 2001.
- [25] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker. A scalable content addressable network. In *ACM SIGCOMM*, 2001.
- [26] S. Ratnasamy, M. Handley, R. Karp, and S. Shenker. Topologically-aware overlay construction and server selection. In *INFOCOM*, 2002.
- [27] S. Ratnasamy, M. Handley, R. M. Karp, and S. Shenker. Application-level multicast using content-addressable network. In *NGC*, 2001.
- [28] S. Ratnasamy, S. Shenker, and I. Stoica. Routing algorithms for DHTs: Some open questions. In *IPTPS*, 2002.
- [29] A. Rowstron and P. Druschel. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In *Middleware*, 2001.
- [30] S. Saroiu, P. K. Gummadi, and S. D. Gribble. A measurement study of peer-to-peer file sharing systems. In *MMCN*, 2002.
- [31] I. Stoica, D. Adkins, S. Zhuang, S. Shenkerand, and S. Suranas. Internet indirection infrastructure. In *ACM SIGCOMM*, 2002.
- [32] I. Stoica, R. Morris, D. Karger, F. Kaashoek, and H. Balakrishnan. Chord: A peer-to-peer lookup service for Internet applications. In *ACM SIGCOMM*, 2001.
- [33] M. Waldvogel and R. Rinaldi. Efficient topology-aware overlay network. Research Report RZ-3436, IBM, 2002.
- [34] Yoid: Your own Internet distribution. <http://www.isi.edu/yoid/>.
- [35] B. Y. Zhao and et al. Brocade: Landmark routing on overlay networks. In *IPTPS*, 2002.
- [36] B. Y. Zhao, A. Joseph, and J. Kubiawicz. Locality-aware mechanisms for large-scale networks. In *FuDiCo*, 2002.
- [37] B. Y. Zhao, J. D. Kubiawicz, and A. D. Joseph. Tapestry: An infrastructure for fault-tolerant wide-area location and routing. Technical Report UCB/CSD-01-1141, UCB, 2001.
- [38] S. Q. Zhuang and et al. Bayeux: An architecture for scalable and fault-tolerant wide-area data dissemination. In *NOSSDAV*, 2001.