On Small World Graphs in Non-uniformly Distributed Key Spaces*

Sarunas Girdzijauskas, Anwitaman Datta, Karl Aberer
Ecole Polytechnique Fédérale de Lausanne (EPFL)
School of Computer and Communication Sciences
CH-1015 Lausanne, Switzerland
Email: {sarunas.girdzijauskas, anwitaman.datta, karl.aberer}@epfl.ch

Abstract

In this paper we show that the topologies of most logarithmic-style P2P systems like Pastry, Tapestry or P-Grid resemble small-world graphs. Inspired by Kleinberg's small-world model [6] we extend the model of building "routing-efficient" small-world graphs and propose two new models. We show that the graph, constructed according to our model for uniform key distribution and logarithmic outdegree, will have similar properties as the topologies of structured P2P systems with logarithmic outdegree. Moreover, we propose a novel model of building graphs which support uneven node distributions and preserves all desired properties of Kleinberg's small-world model. With such a model we are setting a reference base for nowadays emerging P2P systems that need to support uneven key distributions.

Keywords: *Distributed Hash Tables, Routing, Small-World graphs, Storage Load Balancing*

1. Introduction

After the success of the first generation of P2P systems, such as Napster and Gnutella, an immense research activity started to address drawbacks of the original P2P systems, such as centralization (Napster) and enormous bandwidth consumption (Gnutella). One outcome of this research was a variety of approaches to structured P2P overlay networks [1, 12, 13, 14] providing better scalability, accuracy and efficiency. Since then, the number of proposed solutions for structured P2P overlay networks has been growing rapidly, so that it is somewhat difficult to qualitatively and quantitatively compare them. Still it is not hard to notice that many of the proposed solutions share similar properties and are structured in a comparable way.

Among the wide range of proposed structured P2P overlay networks a majority can be characterized as *logarithmic-style P2P overlay networks* which, although being different in their maintenance algorithms, share the same structural properties and search algorithm characteristics. E.g. expected search cost in P2P systems like balanced P-Grid [1], Chord [14], Pastry [13]of them maintains on average $O(\log N)$ entries in their routing tables.

In this paper we first provide one way to better understand the common characteristics of these systems by relating them to the seminal work on small-world graph construction introduced by Kleinberg [6]. This applies in particular for randomized overlay network structures [5] and randomized variants of deterministic structures such as randomized Chord [7, 18]. The first contribution of our paper is to clarify the relationship among logarithmic structured overlays and Kleinberg's model. We introduce a modified version of Kleinberg's model where the outdegree of the overlay graph is logarithmic instead of constant. This not only provides a better insight into the nature of existing logarithmic-style overlay networks, but also a foundation to develop less constrained overlay network structures and to trade-off between search and maintenance cost by choosing the routing table sizes flexibly by varying from constant to logarithmic size.

In a second step we address within the same framework a problem that is receiving recently increasing interest in many data-oriented P2P applications. In this type of applications it is important to preserve semantic relationships among resource keys, such as ordering or proximity, to allow semantic data processing, such as complex queries or information retrieval. This implies that the construction of (efficient) overlay networks has to support the case of non-

* The work presented in this paper was (partly) carried out in the framework of the EPFL Center for Global Computing and supported by the Swiss National Funding Agency OFES as part of the European project Evergrow No 001935. The work presented in this paper was supported (in part) by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under grant number 5005-67322.

uniformly distributed resource keys while exhibiting good load-balancing properties. Examples of overlay networks that have been proposed to address non-uniform key distributions are CAN [12] and P-Grid [1]. Both CAN and P-Grid can partition the key-space upto any granularity, such that each partition has a balanced number of keys assigned to them. In doing so each of the CAN and P-Grid overlay networks sacrifice some desirable properties. Search efficiency in terms of the number of overlay hops can't be guaranteed in CAN for arbitrary partitioning of the key-space (zones). In contrast, P-Grid's randomization helps retaining routing efficiency [2], however peers require more than logarithmic routing states. Pursuing the idea of building the small-world overlay topology we propose a further extension that can handle skewed key distributions without loosing efficient routing properties. Thus in such an overlay network, both routing latency and number of routing states per peer stay $O(log(N))$ independent of the skew of the key-space partition.

By proposing the extension to Kleinberg's model we are providing a foundation for a novel type of structured overlay networks that would support load balancing for unbalanced key and workload distributions, to tradeoff routing table size with search cost, and is expected to be robust due to use of randomization. Given such a foundation for the structural aspect of overlay networks, an orthogonal research question will be the development of algorithms for constructing and maintaining this type of overlay networks. This work can build on the ample experience that has been developed over the last years for the various variants of structured overlay networks, and we wrap the paper outlining some possible design choices and difficulties in realizing a real overlay network.

2. Background

The Small-World phenomenon in social networks was discovered in the sixties [10]. Since then there were numerous works and proposals to explain and model small-world graphs. One approach for building small-world graphs was presented by Watts and Strogatz [16]. The idea was to randomly rewire a regular graph. Starting from a regular graph with constant outdegree, with probability parameter $p \in [0..1]$ at each node an edge is re-linked to another randomly chosen node. With the parameter $p = 1$, one obtains a completely random graph and with $p = 0$, the graph remains regular. When the probability p is between 0 and 1, one obtains a wide range of small-world graphs, that have properties of both regular and random graphs: high clustering coefficient and low diameter. Kleinberg proved [6] that among that wide range of small-world graphs, there exists only one class of small-world graphs in which decentralized (greedy) routing is most efficient. Kleinberg proposed different algorithms for constructing small-world graphs. The idea is to rewire the links to other nodes not randomly, but depending on the distance to the other node.

In Kleinberg's model nodes populate a regular k-dimensional lattice and each node has a *neighboring link* to the neighboring nodes that are a unit distance away from the given node. Each node also has a constant number of *long-range links* that are chosen among the whole set of nodes. Node u chooses to have a long-range link to v with probability inversely proportional to $d(u, v)^r$, where $d(u, v)$ is the distance between these nodes and r is a structural parameter. It was proven that to construct "routing-efficient" small-world graph (where greedy distance minimizing routing will perform best) is possible iff the structural parameter r is equal to the space dimension.

Several recent works employ various small-world properties for building P2P systems, e.g., Symphony [8] or SWAM [3] to name a few. There are other ongoing work in the area trying to extend Kleinberg's model and to use his ideas to improve the performance of P2P networks, like [17, 4], but to the best of our knowledge, none of these address the issue of skewed key-spaces.

2.1. Notations and definitions

In the following we will introduce a variation of Kleinberg's model which shows that the properties of standard logarithmic-cost overlay networks, i.e. logarithmic cost of routing (in terms of overlay hops) with logarithmic size routing tables, can be achieved under much weaker assumptions than usually made. Since many existing DHT proposals are based on one-dimensional key spaces (e.g., Chord, P-Grid, Pastry), we will give the result for this case, and more precisely for the case of an interval topology. Analogous result can be given for other topologies, in particular the ring topology. Unlike as in Kleinberg's proof, we relax our assumptions such that we do not need peers to be connected in a grid, but only that they are randomly distributed in the space according to some probability density function f.

Before introducing the model we provide some notations and definitions that we will use:

- G: the graph resembling P2P overlay network
- N: number of nodes (peers) in the P2P system (the graph G)
- R: the identifier (key) space where nodes (peers) are populated
- id: the identifier (key) of the node (peer)
- f: the probability density function setting the manner of how identifiers of the nodes (peers) are distributed in R

3. Extended Kleinberg's Small-World model for Uniform Key Distribution and Logarithmic Outdegree

First we extend Kleinberg's Small-World model for Uniform Key Distribution, i.e. $f = const$. We model a P2P overlay network as a directed graph $G = (P, E)$ with N nodes. Each peer of a P2P overlay network corresponds to a node in the graph and the routing table entries of this peer correspond to the outgoing edges from that node[1]. The nodes are embedded into the key-space R by uniformly randomly distributing them on the unit interval $[0; 1)$ such that each node u obtains an unique identifier $u_{id} \in [0; 1)$. The distance among two nodes u and v is given as

$$d(u, v) = |v_{id} - u_{id}|. \quad (1)$$

The edges E of the graph G can be classified into *neighboring edges* NE and *long-range edges* LE. Each node u has two neighboring edges: one to the left neighbor and one to the right neighbor. This condition makes G always connected. Different to Kleinberg's model we assume that a node has $\log_2 N$ long-range edges (instead of a constant number of long-range edges). Node u can have long-range edge to any node $v \in LE_u$ for which $|v_{id} - u_{id}| \geq \frac{1}{N}$ and v is chosen such that

$$P[v \in LE_u] \propto \frac{1}{d(u, v)}.$$

With the condition $|v_{id} - u_{id}| \geq \frac{1}{N}$ we restrict the choice of long-range edges to nodes that are not too close. Routing in such an overlay network is based on *greedy distance minimizing routing*. In each step a node u forwards a search request for a target key t to the node with the minimal distance to the target node t among all nodes reachable through an edge from u. We prove that under this model, the expected search cost in number of overlay hops is $O(\log_2 N)$ as in all logarithmic-cost P2P overlay networks. The proof is structurally the same as for Kleinberg's model, however, the bounds have to be derived differently due to the changed model.

Theorem 1 *The expected routing cost in the graph built according to "Model for uniform key distribution" using greedy distance minimizing routing is $O(log_2 N)$.*

Proof. The probability that a node u chooses a node v as a long range contact is $\frac{\frac{1}{d(u,v)}}{\sum_{v \in LE_u} \frac{1}{d(u,v)}}$. First we have to calculate the upper bound of $\sum_{v \in LE_u} \frac{1}{d(u,v)}$ for any node u. The sum can acquire its highest value when it is calculated for a node u which is at the center of the key-space. Thus if we measure the sum for $u_{id} = \frac{1}{2}$, it gives an upper-bound. The distance from u to the closest node will be at least $d_u \geq \frac{1}{N}$. We can calculate expected mean of inverse distance values from the node u to all the other nodes given probability density function $f(x)$ as $2 \int_{d_h}^{\frac{1}{2}} \frac{1}{x} f(x) dx$. Because nodes are distributed uniformly $f(x) = 1$ and $\sum_{v \in LE_u} \frac{1}{d(u,v)}$ is upper-bounded by:

$$N2 \int_{\frac{1}{N}}^{\frac{1}{2}} \frac{1}{x} dx = 2N \ln x|_{\frac{1}{N}}^{\frac{1}{2}} < 2N \ln N. \quad (2)$$

Hence, the probability that node u will choose v as one of its long-range links is at least

$$\frac{1}{d(u, v) 2N \ln N}. \quad (3)$$

Let us view the key-space as $\log_2 N$ partitions $A_1, A_2, .., A_{\log_2 N}$, where each partition A_j is populated by the nodes whose distance from the target node t is $[2^{-\log_2 N + j - 1}; 2^{-\log_2 N + j})$. During greedy routing after a node forwards the search request to node s we say that the message is at partition A_j if the distance between the current message holder s and the target t is within the range $2^{-\log_2 N + j - 1} \leq d(s, t) < 2^{-\log_2 N + j}$. We calculate the probability P_{next} that the current message holder has at least one long-range link to some node v in some partition A_l where $l < j$, i.e. the current message holder can forward the message closer to the target at least by one partition. There are at least $2N2^{-\log_2 N + j - 1}$ such nodes. The distance from the current message holder to the most distant node in the partition A_l is at most $2^{-\log_2 N + j - 1} + 2^{-\log_2 N + j} = 3 * 2^{-\log_2 N + j - 1}$. Using (3) we can determine that the probability that node u will choose some v in A_l as one of its long-range links is at least

$$\frac{2N2^{-\log_2 N + j - 1}}{3 * 2^{-\log_2 N + j - 1} \cdot 2N \ln N} = \frac{1}{3 \ln N}. \quad (4)$$

Since each node has $\log_2 N$ long-range links, P_{next} is on expectation at least

$$P_{next} \geq 1 - \left(1 - \frac{1}{3 \ln N}\right)^{\log_2 N} > 1 - e^{-\frac{1}{3 \ln 2}} = c. \quad (5)$$

Thus, when each node has $\log_2 N$ long-range links the lower bound of the probability that the message will be forwarded closer to the target partition does not depend on N and is a constant that we denote by c. The probability that the message will stay in the same partition when node u forwards it to the next node is at most $P_{same} \leq 1 - c$.

Let us denote by X_j the total number of hops and EX_j as the expected total number of hops that greedy distance minimizing routing will make within the partition A_j before jumping into some partition A_l that is closer to the target t, i.e., $l < j$. If N_{A_j} is the number of nodes in A_j then we have

1 We use the terms "peer" and "node" interchangeably.

$$EX_j = \sum_{i=0}^{N_{A_j}} iPr[X_j = i] < \sum_{i=0}^{\infty} i(P_{same})^i P_{next}$$
$$= \sum_{i=0}^{\infty} i(1-c)^i c = \frac{1-c}{c}. \quad (6)$$

There exist $\log_2 N$ partitions of the key-space and the expected number hops in each of them is less than $\frac{1-c}{c}$, so the expected total number of hops that the algorithm will need, including the long-range hops is at most $(\frac{1-c}{c} + 1)\log_2 N$. The expected number of nodes that algorithm will have to visit using neighboring edges from the partition A_1 to the target node t is $N \int_0^{\frac{1}{N}} f(x)dx = 1$. Therefore the total expected number of hops is $\frac{1}{c}\log_2 N + 1$, i.e. $O(\log_2 N)$.

q.e.d.

Note that a tighter bound can be derived by determining the expected number of long-range hops, and thus our derivation is gives a pessimistic upper-bound, which nonetheless suffice to prove that the expected cost is $O(\log_2 N)$.

3.1. Similarities with logarithmic-style P2P overlays

Notice the fact that EX_j is a small constant. This means that each logarithmic partition of the key-space is reached in a constant number of hops. This result can be explained by the fact, that such a small-world graph possesses nice "probabilistic partitioning" properties which are also widely exploited in traditional logarithmic-style P2P overlays. Indeed, in traditional logarithmic-style P2P overlays each peer u views the identifier space partitioned in $\log_2 N$ logarithmic partitions of identifier space where each partition is twice bigger than the previous one (or k times bigger if we consider base k logarithmic partitioning, e.g. in Pastry $k = 16$). The routing table of u in such systems contains $\log_2 N$ links to some node from every partition. E.g. in Chord [14] the chosen node will be with the smallest identifier of the given partition, in Pastry [13] and P-Grid [1] - any random node of the partition. While routing, the message in every next hop is being routed to a node which belongs to a partition, that is at least twice (k times) smaller than the previous partition where the previous message holder (node) used to be. Therefore we can imagine such a P2P network as a space where the message approaches the target with steps of logarithmically decreasing size.

Overlays based on a graph built according to the above mentioned variation of Kleinberg's model, will have a very similar topology and routing properties as logarithmic-style P2P overlays. Indeed we can partition the identifier space of any node u into $\log_2 N$ partitions $A_1, A_2.., A_{\log_2 N}$, where A_j consists of all nodes whose distance from u is between $2^{-\log_2 N + j - 1}$ and $2^{-\log_2 N + j}$ (every next partition is twice bigger as the previous one). It is interesting to observe, that in this case node u has almost equal probabilities to choose the long-range neighbor from each of these partitions. Therefore when each node chooses $\log_2 N$ long-range neighbors in the same way, they will be uniformly distributed among the partitions, whereas in logarithmic-style P2P overlays $\log_2 N$ neighbors would be chosen strictly from each partition. We can consider logarithmic-style P2P overlay topologies as one "special case" of small-world topology with stronger restrictions. This provides insight into how our modification of Kleinberg's model relaxes existing logarithmic-cost overlay networks where routing entries have to point to each logarithmic partition of the key-space. Hence the possibility to generalize and model the behavior of logarithmic-style P2P topologies from a small-world model point of view.

The feature of "Kleinbergian" graphs to model logarithmic P2P topologies suggests a more flexible manner of maintaining the networks. One of the possibilities would be to maintain a variable number of entries in routing tables for a tradeoff of logarithmic to polylogarithmic search cost, an observation that was also made in Symphony [8]. It also implies that the networks built according to "Kleinbergian" style would be more robust and resistant to network churn. Even in the case of connectivity loss, the routing cost will be at worst poly-logarithmic given we have at least one long-range link and the neighboring links intact.

4. Extended Small-World Model for Skewed Key Distribution

The main purpose of P2P overlay networks is to distribute resources among peers, such that resources can be be efficiently located and the workload is distributed as uniformly as possible among peers. In most standard P2P overlay networks uniform workload distribution is achieved by applying randomized hashing functions, such as SHA-1, to resource identifiers such that the hashed identifiers are uniformly distributed in the key-space. Then by also uniformly distributing peers in the key-space an approximately uniform load distribution is achieved. However, in many data-oriented P2P applications it is important to preserve relationships among resource keys, such as ordering or proximity, to allow semantic data processing, such as complex queries or information retrieval. Thus uniform key distribution cannot be assumed, and in order to achieve uniform workload peers will be distributed non-uniformly in the key-space. In addition, different resources might be associated with different workload patterns, e.g. query frequency, which require further adaptations in the distribution of the peers over the key-space.

In the following we show that the construction we introduced in the previous Section 3 can be extended to peers, distributed non-uniformly in the key-space, without loosing routing efficiency in terms of either the expected routing latency or the number of routing states per peer. This provides the theoretical foundation for developing a novel class of P2P overlay networks that are able to deal with non-uniform load distributions.

4.1. Model for skewed key distribution

We assume that there exists a mechanism that assigns peers according to a non-uniform distribution in the key-space adapting to the load-distribution (e.g., storage), such that the balanced number of data objects are assigned to each peer, irrespectively of their distribution in the key-space. Several examples of such mechanisms have been recently discussed in the literature [2, 15, 11]. Thus each peer acquires its identifier according to a non-uniform probability density function f. In order to account for the non-uniform peer distribution peers have to choose their long-range neighbors in graph G according to the following criterion: a peer u chooses peer v as long-range neighbor with a probability that is inversely proportional to the integral of probability density function between these two nodes, i.e.

$$P[v \in LE_u] \propto \frac{1}{\int_{u_{id}}^{v_{id}} f(x)dx}.$$

Using this criterion we claim that routing in the resulting overlay network is as efficient as in the case of uniform (balanced) key distribution.

Theorem 2 *The expected routing cost in the graph built according to the "Model for skewed key distribution" using greedy distance minimizing routing is* $O(log_2 N)$.

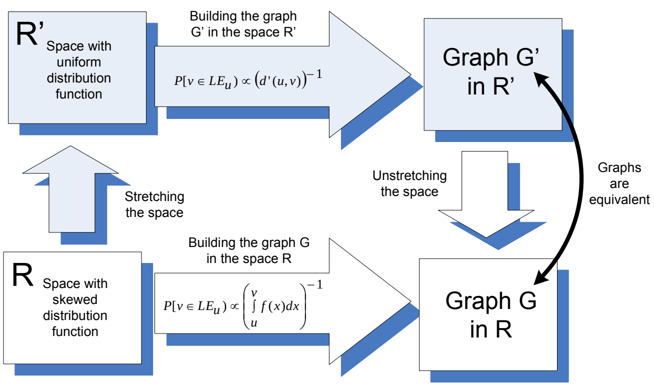

Figure 1. Normalization of the space

Proof. We have to show that by using the modified selection criterion for long-range links we are constructing a routing-efficient graph G. The schema of the proof is depicted in Figure 1. The idea underlying our model is to normalize the space R in such a way that the normalized space R' will have a uniform probability density function f'. Any node u with identifier u_{id} in the space R will have a corresponding identifier u'_{id} in the space R'. The value of identifier u'_{id} is chosen as $u'_{id} = \int_0^{u_{id}} f(x)dx$, such that

$$\int_0^{u_{id}} f(x)dx = \int_0^{u'_{id}} f'(x)dx$$

and peer identifiers are uniformly distributed in R'. The distance between two nodes u and v in the space R' can be represented as

$$d'(u'_{id}, v'_{id}) = \int_{u_{id}}^{v_{id}} f(x)dx. \quad (7)$$

As described in the previous section we already know how to construct a "routing-efficient" graph, i.e. choosing long-range links proportional to $\frac{1}{d'(u'_{id}, v'_{id})}$. As we have already proven, the resulting graph G' will be "routing-efficient", i.e. the expected search cost using greedy distance minimizing routing will be $O(\log_2 N)$.

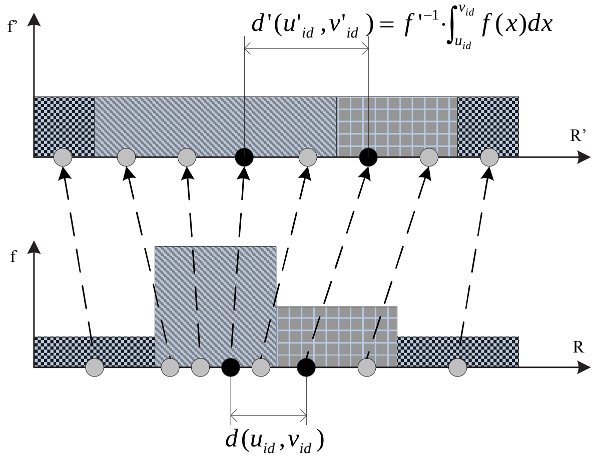

Figure 2. Mapping of nodes from R to R'

As shown in Figure 2 using the original criterion for selecting long-range links for uniform key distribution in space R', i.e. inverse proportional to $d'(u'_{id}, v'_{id})$, is equivalent to choosing long-range links directly in space R using the modified criterion, i.e. inverse proportional to $\int_{u_{id}}^{v_{id}} f(x)dx$. The resulting graph G in the original space R will have the same connectivity as the graph G' constructed in space R', although the peers have different identifiers. The search efficiency depends only on the connectivity of the graph, therefore the resulting graph G will have the same search efficiency as graph G', i.e. $O(\log_2 N)$. *q.e.d*

4.2. Building efficient structured overlay for non-uniformly distributed key-spaces

The adaptation of our model in practice is straightforward in the case where each peer in the P2P network knows the global key distribution, i.e. the probability density function f. In such a case the following network construction model can be applied.

While joining the network, some peer u generates a value according to probability density function f and assigns it as its identifier. The peer u contacts any known peer and issues a query with that identifier. When u gets an answer from some peer v (in this case v has the closest identifier to u), u announces to v that it will become its immediate neighbor. Both u and v correct in their routing tables of the immediate neighboring links. Because the peer u knows the function f it can calculate the boundaries of partitions $A_1, A_2.., A_{\log_2 N}$, where A_j, consists of all nodes whose distance from u is between $\int_0^{2^{-\log_2 N+j-1}} f(x)dx$ and $\int_0^{2^{-\log_2 N+j}} f(x)dx$. By randomly selecting the value in each of the partitions u issues $\log_2 N$ search queries and the peers that respond are added to its routing table as long-range neighbors. In such a way the peer u completely joins the network.

The task however is more complicated for a more realistic situation, where peers do not have information of the distribution f and have to acquire it locally, by interacting with other peers. Moreover, the distribution f may vary over time, further complicating the design of a practical system. In such a case, at each peer an iterative process of revising its routing table according to the current knowledge on f has to be employed. The above mentioned steps have to be repeated whenever a peer obtains more precise information about f. Such iterative process can be performed indefinitely if the function f changes over time in the system. In this way the topology would be always self-adjusted to the current conditions of the system.

As our future work we plan to investigate what maintenance cost is necessary under various models of network churn in order for each peer to be able to predict locally a good enough approximation of the probability density function f for the construction of "routing-efficient" P2P networks.

5. Conclusions

The work of Kleinberg on small-world graphs caused a stir in P2P community. It boosted the research towards investigating randomized topologies [7, 9] and even resulted in new proposals such as Symphony [8]. We used Kleinberg's model to provide a perspective on existing standard P2P overlay networks and to explain their nature. In this paper we introduced two variants of Kleinberg's model which allow to model a large class of P2P overlay networks. The first model for uniform key distribution and logarithmic outdegree allows us to better understand the behavior of logarithmic-style P2P overlay networks. The flexibility of Kleinberg's model demonstrates the possibility of making flexible logarithmic P2P topologies by allowing them to change routing table size from constant to logarithmic. With our second model we showed that with Kleinberg's principle of building "routing-efficient" networks we can build P2P topologies for skewed key distributions. As for now, we are working on the theoretical analysis of variation of the expected search cost and on models that can take into account an unstable P2P environment (nodes are allowed to fail). As our future work we will adapt our variant of overlay network, P-Grid, according to the proposed model and thus develop solutions for the construction and maintenance of the network. Doing that, an important aspect will be the study of various tradeoffs between, maintenance and search cost, load balancing characteristics and robustness of the resulting networks in a more generic setting as it can be done today with the relatively constrained topologies used.

References

[1] K. Aberer. P-Grid: A self-organizing access structure for P2P information systems. In *Proceedings of the Sixth International Conference on Cooperative Information Systems (CoopIS)*, 2001.

[2] K. Aberer, A. Datta, and M. Hauswirth. Multifaceted Simultaneous Load Balancing in DHT-based P2P systems: A new game with old balls and bins. In *Self-* Properties in Complex Information Systems, "Hot Topics" serices, Lecture Notes in Computer Science*. Springer Verlag(to be published), 2005.

[3] F. Banaei-Kashani and C. Shahabi. Swam: A family of access methods for similarity-search in peer-to-peer data networks. In *Information and Knowledge Management (CIKM'04), Washington D.C.*, 2004.

[4] M. Franceschetti and R. Meester. Navigation in small world networks, a scale-free continuum model. Under revision.

[5] K. Gummadi, R. Gummadi, S. Ratnasamy, S. Shenker, and I. Stoica. The Impact of DHT Routing Geometry on Resilience and Proximity. In *Proceedings of the ACM SIGCOMM*, 2003.

[6] J. Kleinberg. The Small-World Phenomenon: An Algorithmic Perspective. In *Proceedings of the 32nd ACM Symposium on Theory of Computing*, 2000.

[7] G. S. Manku. Routing networks for distributed hash tables. In *Proceedings of the twenty-second annual symposium on Principles of distributed computing (PODC)*, pages 133–142. ACM Press, 2003.

[8] G. S. Manku, M. Bawa, and P. Raghavan. Symphony: Distributed hashing in a small world. In *4th USENIX Symposium on Internet Technologies and Systems, USITS*, 2003.

[9] G. S. Manku, M. Naor, and U. Wieder. Know thy neighbor's neighbor: the power of lookahead in randomized p2p networks. In *36th ACM Symposium on Theory of Computing, STOC 2004, p 54-63*, 2004.

[10] S. Milgram. The small world problem. In *Psychology Today 1, 61*, 1967.

[11] Mayank Bawa Prasanna Ganesan and Hector Garcia-Molina. Online balancing of range-partitioned data with applications to peer-to-peer systems. In *Proceedings VLDB*, 2004.

[12] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker. A Scalable Content-Addressable Network. In *Proceedings of the ACM SIGCOMM*, 2001.

[13] A. Rowstron and P. Druschel. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In *IFIP/ACM International Conference on Distributed Systems Platforms (Middleware), Heidelberg, Germany*, 2001.

[14] I. Stoica, R. Morris, D. Karger, F. Kaashoek, and H. Balakrishnan. Chord: A Scalable Peer-To-Peer Lookup Service for Internet Applications. In *Proceedings of the ACM SIGCOMM*, 2001.

[15] X. Wang, Y. Zhang, X. Li, and D. Loguinov. On zone-balancing of peer-to-peer networks: Analysis of random node join. In *ACM SIGMETRICS*, 2004.

[16] D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. In *Nature, vol. 393, pp. 440-442*, 1998.

[17] H. Zhang, A. Goel, and R. Govindan. Using the small-world model to improve freenet performance. In *Proc. of IEEE Infocom*, 2002.

[18] H. Zhang, A. Goel, and R. Govindan. Incrementally improving lookup latency in distributed hash table systems. In *In ACM SIGMETRICS 2003, pages 114125*, 2003.