

RAQ: A Range-Queriable Distributed Data Structure

Hamid Nazerzadeh ^{*} Mohammad Ghodsi [†]

Abstract

Different structures are used in peer-to-peer networks to represent their inherently distributed, self-organized and decentralized memory structure. In this paper, a simple range-queriable distributed data structure, called RAQ, is proposed to efficiently support exact match and range queries over multi-dimensional data. In RAQ, the key space is partitioned among the network with n nodes, in which each element has links to $O(\log n)$ other elements. We will show that the look-up query for a specified key can be done via $O(\log n)$ message passing. In addition, RAQ handles range-queries in at most $O(\log n)$ communication steps. We also present a probabilistic variation of RAQ in which each node has $O(\log \log n)$ links on the average and look up and range queries are resolved with $O(\log^2 n)$ latency. At the end we evaluate RAQ through simulation.

1 Introduction

Distributed and peer-to-peer networks are significant components of recent research on networking. There is a simple idea behind the peer-to-peer networks: each node maintains its own index and searching mechanism compared to the traditional client-server architecture with global information. The significant growth in the scale of such networks, (e.g. Gnutella [4]), reveals the critical emerging need to develop decentralized searching methods. A peer-to-peer storage system can be considered as a large scale distributed decentralized data structure. We use the term *Queriable Distributed Data Structure* (QDS) to denote such a self-organized, decentralized, distributed, internet-scale structure which provides searching and data transferring services. In QDS, every node of the network is an element of the whole structure, which provides decentralized searching services over the data scattered among the nodes of the network. New file sharing systems such as Scour, FreeNet, Ohaha, Kazaa and Jungle Monkey are examples of QDS from current internet systems.

Distributed Hash Table (DHT) [15, 11] can be viewed as a QDS. In DHT systems, keys and data are stored in the nodes of the network using a *hash function*, in which data can be the addressing information (e.g. IP address of the server containing the data), rather than its actual data. Searching mechanism in these systems consists of two main phases: (1) hashing the key, and (2) querying the network to find the node that contains the key. This node resolves the query by providing the actual data or its addressing information. Similarly, some other systems like *SkipNet* [6] which are based on more theoretical data structures like *skip graphs* [2], allow more flexibility on the location of data on the nodes of the network.

In this paper, we propose RAQ, a range-queriable distributed data structure to resolve exact match and range queries over multi-dimensional data efficiently. In RAQ, the key space is partitioned among the n nodes of the network, in which each element has links to $O(\log n)$ other

^{*}Computer Engineering Department. Sharif University of Technology, <http://www.bamdad.org/~hamid>, email: nazerzadeh@sharif.edu .

[†]Computer Engineering Department. Sharif University of Technology, Tehran, Iran, <http://www.sharif.edu/~ghodsi>, email: ghodsi@sharif.edu .

elements of the network. We will show that the look-up query for a specified key can be done, in our structure, via $O(\log n)$ message passing. The bound on the out-degree of the nodes and the exact-match query cost are both comparable to those in DHT systems like Chord [15], CAN [11], Pastry [13] and Viceroy [9].

The main contribution of RAQ is that it is simple and can handle range-queries in multi-dimensional space. Our data structure supports such queries in at most $O(\log n)$ communication steps. *Split the Space, Duplicate the Query* is a novel approach used by the RAQ to resolve range-queries. This method anticipates the answer space of the query at the source and spreads the query only through the appropriate nodes by duplicating the query; each of these new queries addresses a reduced subspace.

Most other QDS systems do not support multi-dimensional range-queries, because they mostly use one-dimensional key space. CAN [11] supports multi-dimensional key space, but in spite of its similarity to RAQ's basic structure, the out-degree of node and its routing cost depend on the dimension of the key space. For a d dimensional space, the average routing path length in CAN is $(d/4)(n^{1/d})$ hops and individual nodes maintain $2d$ neighbors. This limitation forces the system to use hashing to reduce the dimension in some applications. But, since hashing destroys the logical integrity of the data, CAN cannot support range queries over multi-dimensional data efficiently.

We also present, PRAQ, a probabilistic variation of RAQ in which the connection rule of RAQ is modified in such a way that each node has $O(\log \log n)$ links on average. We will show that the look up and range queries can be handled in PRAQ via $O(\log^2 n)$ communication steps.

In this paper, we first overview the related works briefly, followed by the principal ideas and structures of RAQ. Query handling methods are discussed in sections section 5 and section 6, followed by the algorithms for joining and leaving a node. In section section 8, we describe the probabilistic model and at the end, go through the experimental results.

2 Related Works

Supporting range queries in QDS systems has been the subject of several recent works. In SWAM [3], for example, the key space is partitioned according to *Voronoi Tessellation*. By this property, and using links based on *Small-World Phenomenon* [8, 10], SWAM resolves k -nearest-neighbor search and range queries via $O(\log n + R)$ message passing, where R is the size of the answer. But, the number of links of each node grows exponentially with the dimensions of search space [14].

Prefix Hash Tree is a solution proposed by Rantasamy *et. al* [12] to face the problem of hashing used in DHTs that destroys the integrity of data. Their approach is based on distributed *trie*. Given a query, this system attempts to recognize the longest prefix of the query that appears as a trie-node. The complexity of this operation is $O(\log \log d \times \log n)$, where d is the size of the discrete domain.

Recently, some authors have addressed the problem from the load balancing viewpoint. Karger and Ruhl [7], offer an approach to overcome the known limitation of hash functions, by designing an algorithm to maintain load balance in non-uniform distributed key space. Aspens *et al.* [1] have proposed a similar solution based on *skip graphs*.

Gao and Steenkiste [5], present a QDS that relies on a logical tree data structure, the Range Search Tree (RST), to support range queries in one dimensional space. In this system, nodes in RST are registered in groups. To handle the range queries, queries are decomposed into a small number of sub-queries where the cost depends on the load factor of the data and query capacity of the nodes in the network.




Figure 1: The partition tree, \mathcal{P}_2 , on the left corresponds to the points on the right.

3 Partition Tree

We have n points in our d -dimensional space. The *Partition Tree*, \mathcal{P}_d , is the main data structure used in RAQ. Similar to the data structure used in [11], \mathcal{P}_d partitions the d -dimensional space, so that in the final level, each region has only one point. Each internal vertex of the tree corresponds to a region in space, and the root represents the whole space. Each pair of the sibling vertices divide their parent region into two parts, and each leaf represents an undivided region called a *cell*, each corresponding to a single point in that region. Figure 1 portrays the partitioning of \mathcal{P}_2 .

Each vertex x is assigned a *label* to specify its corresponding subspace. We define $x_{label} = ((p_1, d_1), (p_2, d_2), \dots, (p_{r(x)}, d_{r(x)}))$ where:

$r(x)$: The distance of x from the root of the tree.

p_i : The plane equation that partitions the i th region into two parts.

d_i : Determines one side of the plane p_i .

$parent(x)_{label} = ((p_1, d_1), \dots, (p_{r(x)-1}, d_{r(x)-1}))$

$sibling(x)_{label} = ((p_1, d_1), \dots, (p_{r(x)}, \bar{d}_{r(x)}))$, where \bar{d}_i is the opposite side of d_i .

$root_{label} = \lambda$, i.e. the empty string.

We treat the labels as strings. The expression $l_1 \sqsubset l_2$ means that l_1 is a prefix of l_2 , $|l|$ represents the size of l (i.e. the number of (p_i, d_i) pairs) and $+$ is the concatenation operator. Also, $[l] \stackrel{\text{def}}{=} \{x \in V \mid l \sqsubset x_{label}\}$ where l is a label and V is the vertex set of the partition tree. Obviously, for a vertex x , $|x_{label}| = r(x)$.

4 Design Principles of RAQ

RAQ is a structure on the nodes of a network. Each node maps to one point in the d dimensional search space. A partition tree $\mathcal{P}tree$ is constructed over the points and thus each node corresponds to a single cell. We say that a node owns its cell and is responsible for providing data to the queries targeting any point in that cell. Since there is a one-to-one map between nodes and the leaf points of the partition tree, we use them interchangeably. So, for example, we assume having labels for each node.

Moreover, each node has several *links* to other nodes of the network. Each link is basically the addressing information of the target node which can be its IP address. The links are established based on the partition tree information and the following rule:

Connection Rule: Consider node x and its label $x_{label} = ((p_1, d_1), (p_2, d_2), \dots, (p_k, d_k))$. The connection rule of node x implies that x is connected to exactly one node in each of these $|x_{label}|$ sets:

$$[(p_1, \bar{d}_1)], [(p_1, d_1), (p_2, \bar{d}_2)], \dots, [(p_1, d_1), (p_2, d_2), \dots, (p_k, \bar{d}_k)]$$

For example, in Figure 1, node 1 is connected to one node in each of these sets: $\{2\}, \{4\}, \{6, 7, 8\}$. We will show that the *joining* and *leaving* mechanisms guarantee the maintenance of connection rule over the network.

Lemma 1 *An arbitrarily chosen vertex has links to $O(\log n)$ other nodes in the network.*

It is important to note that the partition tree is not directly maintained by the elements of RAQ; given the coordinates and the labels of the leaves, all information of the partition tree can be uniquely obtained, and these are the only data which are maintained by the nodes of the network. In fact, the partition tree is the *virtual data structure* of RAQ.

It is obvious that $\mathcal{P}tree$ is a balanced tree with the height of $O(\log n)$ when it is first constructed. We will show that this property holds even in the dynamic environment where the nodes join and leave the network frequently.

5 Exact Match Query

In RAQ, exact-match queries are of the form $Exact\text{-}Query(target, metadata)$. The value of *target* is the coordinate of the point that we are searching for, and *meta-data* contains the data to be used after the query reaches the target. Note that the queries aim to reach the target and the responses may vary in different cases. As mentioned, the target of a query is a node whose region contains the query's target point.

We say, a point p matches a label l at level k , if k is the greatest value of i such that the subspace induced by a \mathcal{P}_d vertex x with $x_{label} = ((p_1, d_1), (p_2, d_2), \dots, (p_i, d_i))$ contains p and $x_{label} \sqsubset l$. In other words, say l represents a leaf y in $\mathcal{P}tree$, then, x is the lowest common ancestor of y and the node containing p .

Lemma 2 *Suppose node x receives a query targeting point p and p matches x_{label} at level k . If $k = |x_{label}|$ then the cell of x contains p . Otherwise, x has a link to a node y so that y_{label} matches p at a level greater than k .*

Proof: Let $x_{label} = ((p_1, d_1), \dots, (p_{r(x)}, d_{r(x)}))$. If $k = |x_{label}|$, then, obviously, x contains p . If not, from the connection rule, we know that x is linked to a node y with $((p_1, d_1), \dots, (p_{k+1}, \bar{d}_{k+1})) \sqsubset y_{label}$. Therefore, according to the definition, p matches y_{label} at a level not less than $k + 1$. \square

Now, the algorithm for exact match routing becomes clear. Once the query Q is received by a node x , if x contains the target point, then we are done. Otherwise, x sends the query via a link to a node y with a label that matches the target point at a higher level. This will continue until the query reaches the target.

Theorem 1 *The exact match query is resolved via $O(\log n)$ message passing.*

Proof: Suppose the target of query Q is the node x . From lemma 2, Q will reach to x in at most $|x_{label}|$ steps and $|x_{label}| = O(\log n)$. So, the routing operation is performed by $O(\log n)$ message passing. \square

6 Range Query

We assume that a range query Q is of the form $Range\text{-}Query(label, pivot, func, d_1, d_2, metadata)$ where $label$ implies that Q must be sent only to the nodes x so that $label \sqsubset x_{label}$, we denote the label of Q by Q_{label} . The initial value of Q_{label} is set to empty string. The value of $pivot$ is the coordinate of the point that the distances are measured from, and $func$ is the distance function.

The above range query means that Q should be sent to every node in the network with the distance of $d_1 \leq d \leq d_2$ from $pivot$. $func$ can be any distance function \mathcal{F} with the following characteristic: Given a point p , a hyper-cubic subspace S and a distance d , let $A = \{x | x \in S \text{ and } \mathcal{F}(x, p) < d\}$. The problem of whether or not A is empty must be computable. For example, \mathcal{F} can be L_p -norm function, in which case the answer space of Q is $\{x | d_1 \leq (\sum (p_i - x_i)^p)^{1/p} \leq d_2\}$.

To handle the range queries, we use *split the space, duplicate the query* method, or split-duplicate for short. Suppose that node x receives a range query Q and $x_{label} = ((p_1, d_1), \dots, (p_{r(x)}, d_{r(x)}))$. Obviously, $Q_{label} \sqsubset x_{label}$. If $Q_{label} = x_{label}$, then x itself will give the appropriate response. Otherwise, we iterate the following sequence of operations:

1. Duplicate Q and name the results Q_1 and Q_2 . Set $Q_{1label} = Q_{label} + (p_{|l|+1}, d_{|l|+1})$ and $Q_{2label} = Q_{label} + (p_{|l|+1}, \bar{d}_{|l|+1})$.
2. If the answer space of Q has intersection with the subspace induced by Q_{2label} , then send Q_2 via the link to node y where $Q_{2label} \sqsubset y_{label}$. Note that by the connection rule, y exists.
3. Iterate split-duplicate operation subsequently on Q_1 while the splitted subspace has intersection with the answer space.

Lemma 3 *If node x receives range query Q , then Q will be routed to all nodes y where $Q_{label} \sqsubset y_{label}$ and the intersection of the cell of y and answer space of Q is not empty.*

Proof: We prove this by backward induction on $|Q_{label}|$. If $|Q_{label}| > |x_{label}|$, then obviously $Q_{label} \not\sqsubset x_{label}$ thus the induction basis holds.

Suppose that x receives the query and the induction hypothesis holds for $k > |x_{label}|$. If $Q_{label} = x_{label}$ then x is the only target of the query and we are done. Otherwise, two new queries are generated by the algorithm, while the second query Q_2 is sent to its adjacent node via the appropriate link if the subspace induced by Q_{2label} has nonempty intersection with the answer space. The size of the labels of these new queries is increased by one. Thus, by the induction hypothesis, Q will be routed to all nodes z with $Q_{label} \sqsubset z_{label}$ where the cell of z has a nonempty intersection with the answer space. The union of the induced spaces of these labels covers the whole space of Q_{label} . The claim is therefore correct. \square

Theorem 2 *RAQ resolves range queries in at most $O(\log n)$ communication steps. In other words, a query will be received by a target node by crossing $O(\log n)$ intermediate hops.*

Proof: As we mentioned, the basic query enters the system by initializing its label to an empty string. By lemma 3, the range-query will be received by all nodes whose cells have nonempty intersections with the query answer space. In each communication step, the size of the query label is increased at least by 1. Thus, when a node receives a query, the distance to the source must be $O(\log n)$. \square

7 Joining and Leaving

In this section, we describe the joining and leaving mechanism and demonstrate the validity of our claim in section section 4 that the partition tree always remains balanced.

7.1 Joining Mechanism

Suppose that node x wants to join to the network. x chooses a fairly random point p in the space and then finds y , one of the active nodes in the network. Several mechanism can be adopted for the arriving node to find an active node; we assume that RAQ uses the same mechanism as in CAN [11].

Sending an exact match query by y , the node z whose cell contains p is found. z divides its cell into two parts, with one containing the corresponding point of z and the other including p . We assume that x possesses the cell containing p . This is just a simple insertion into the partition tree. This is done by updating the labels of x and z . Since we are not directly maintaining this tree, this update is sufficient.

The connections are now updated to follow the given connection rule: x chooses one random point in each of the subspaces induced by the labels specified by the connection rule. For each of these points, say r , z routes an exact match query to find the node that owns r . Consequently, x establishes a connection link to this node.

Theorem 3 *Join operation is done via $O(\log^2 n)$ message passing.*

Proof: The arriving node finds its region by an exact match query. By lemma 1, the arriving node has to create $O(\log n)$ connections. Establishing each connection is done by an exact-match lookup, via $O(\log n)$ message passing. Therefore, the whole operation is completed by $O(\log^2 n)$ message passing. \square

7.2 Leaving Mechanism

Let x be the node that wants to leave the network. After x leaves the network, its sibling in $\mathcal{P}tree$ will maintain the region once belonging to x . From $\mathcal{P}tree$ viewpoint, leaving is just a simple deletion of a leaf in a binary tree, so z_{label} and thus $\mathcal{P}tree$ will be updated easily. The difficult part is updating the connection links of the nodes that have links to x . To handle this issue, *Departure Links* or for short *dlinks* are defined below.

In RAQ, node b maintains *addressing information* of a , or a *dlink* to a , when node a establishes a connection link to node b . When b decides to leave the network, it sends a message to each of the nodes referred to by its dlinks. In the following, we refer to the number of dlinks of b as its *d-degree*.

Theorem 4 *The expected value of d-degree is $O(\log n)$.*

Proof: Here we argue the validity of our claim. From the mechanism described above to establish a connection link and from the dynamic structure of the network where the nodes frequently join and leave the network, we can fairly conclude that the probability that a node v has a link to a node u is equal to the probability that u has a link to v . We avoid discussing the unimportant details of this claim. Accordingly, $E[v_{degree}] = E[v_{d-degree}]$ for an arbitrary node v in the network. \square

Corollary 1 *By lemma 1 and theorem Theorem 4, each node of the network maintains the addressing information of $O(\log n)$ nodes of the network.*

Consider the time when x is leaving the network. x sends a departure message to all of the nodes on its dlinks. As mentioned, every connection in RAQ is a link to a subspace. Each of the nodes that receives this departure message chooses a new random point, say p , in the corresponding subspace and sends an exact match query via x to establish a new connection link to the node that possesses p . After these operations, x will peacefully leave the network and the connection rule of RAQ is maintained.

Theorem 5 *The leave operation is done via $O(\log^2 n)$ message passing.*

Proof: According to theorem Theorem 4, $O(\log n)$ links must be updated. Each update is performed by $O(\log n)$ message passing, thus the total number of passed messages is $O(\log^2 n)$. \square

According to the discussion, arriving nodes are distributed all over the space. Thus, the partition grows uniformly and remains balanced. Uniform distribution of the nodes also implies that the nodes leave the network randomly in the entire space. We can thus conclude that the claimed proposition is valid.

8 Probabilistic RAQ

In this section, we describe the idea of PRAQ, a probabilistic variation of RAQ. The deterministic and nondeterministic models are different mainly in connection rule and consequently exact match query handling. Due to lack of space, we will not go through the details.

In nondeterministic structure, there are additional links called *Cyclic Links*. Consider the leaves of $\mathcal{P}tree$ in order that they appear in Post-Order Traversal of the tree and name the sequence \mathcal{A} . Illustratively, for the tree presented in Figure 1 the \mathcal{A} sequence is: 1, 2, 4, 6, 7, 8. Let $\mathcal{A} = a_1, \dots, a_n$. $a_{i_{prev}} \stackrel{\text{def}}{=} a_{i-1}$ and $a_{i_{next}} \stackrel{\text{def}}{=} a_{i+1}$ where the additions and subtractions are done in modular n . A node x has two cyclic links to x_{prev} and x_{next} . These links form the *Universal Cycle* which includes all the network's nodes. Universal Cycle assures us that nodes are always connected. Accordingly, connection rule is modified:

Connection Rule . Consider node x with label $((p_1, d_1), (p_2, d_2), \dots, (p_k, d_k))$. The connection rule of node x implies that x is connected to one of the vertices of $[((p_1, d_1), \dots, (p_i, d_i))]$ with probability $\frac{1}{k-i+1}$ for $2 \leq i \leq k$.

Theorem 6 *The average number of links of a node in the probabilistic model is $O(\log \log n)$.*

Proof: The expected value of the number of the links of node x is $\sum_{i=2}^{r(x)} \frac{1}{r(x)-i+1} + 2$ and by the inequality below the claim is proven.

$$\sum_{i=1}^{\log n} \frac{1}{i} \leq \log \log n$$

\square

Now we show how we can emulate RAQ by PRAQ with dilation $O(\log n)$. From to the connection rule of RAQ we know that the link from a node x to a neighbor y , is in fact a link to a subspace induced by a label.

Theorem 7 *A single message passing step in RAQ can be emulated in PRAQ in $O(\log n)$ steps.*

Proof: Suppose x wants to send a message to the subspace induced by label l . If x has no link to a node in $[l]$, it sends the message to its immediate neighbor in Universal Cycle in the direction of l . This operation will be repeated until the message is received by a node which is either connected to $[l]$ or itself is in $[l]$. Each node y in the path has a link to $[l]$ with probability $\frac{1}{r(y)-l} = O(\frac{1}{\log n})$. Thus the expected number of traversed nodes before reaching $[l]$ is $O(\log n)$. \square

Corollary 2 *PRAQ resolves exact-match and range queries in $O(\log^2 n)$ communication steps.*

Departure and arrival of the nodes in the probabilistic model is similar to the deterministic one and is done by emulating message passing. Inserting a node to the Universal Cycle is as simple as inserting a node in a cyclic link list.




Figure 2: Left: The histogram of the number of peers with a specific degree, Right: The plot of the average number of the messages passed once a node inserts.




Figure 3: Left: The graph of the average number of the messages in resolving a single exact match query, Right: The distribution of the length of the routing paths.

9 Experimental Evaluation

We have performed simulation to evaluate RAQ. We have simulated networks with 2^8 to 2^{16} nodes to measure different metrics of the structure. A network is constructed by inserting nodes one by one using the methods discussed in section section 7. In Figure 2, a histogram of the number of links of each of the nodes in the network for networks of different sizes is presented. As we argued in lemma 1, the average is in $O(\log n)$. Theorem Theorem 3 claims that the join operation is done via $O(\log^2 n)$ message passing which is the same as what we get through the experiment. The next plot in Figure 2 shows the average number of the messages that are transferred in inserting a single node.

The principal metric in evaluating RAQ is the number of the messages passed to answer a single exact match query. We generate exact match queries with random source and destination. According to the simulation's result, Figure 3, the average number of the messages is approximately $\frac{\log n}{2}$. In Figure 3, the distribution of the length of routing paths is displayed. These results confirm our arguments in section section 5.




Figure 4: Left: The plot of the average number of duplicated queries generated to handle range queries, Right: The graph of the average length of the routing paths.

To evaluate the fault tolerability of RAQ, we have simulated the case that a node fails to handle a query by a 10 percent probability. In this situation, when node u sends a query to node v and v rejects the message, u will send the query randomly to one of its other neighbors. The simulation results in Figure 3 show that RAQ is still doing well in this case and routes the queries only slightly slower. Thus it is a reliable structure.

Figure 4, presents the results of simulating range queries. We generate 500 range queries for small networks and 5000 queries for the larger ones. A range query is generated by choosing a hypercubic subspace as the answer space which has about 0.05 percent length along each dimension. The left plot in Figure 4, represent the number of duplicate queries generated to solve the query. As is illustrated by the graph, the number of target nodes of the range queries and the number of messages grow linearly by the size of the network due to the constant size of the answer space. Similar to the exact match queries, the average length of routing latency is about $\frac{\log n}{2}$ which confirms our claims in section section 6.

10 Conclusion

In this paper we presented RAQ, a range-queriable distributed data structure for peer-to-peer networks to organize its multi-dimensional data, and to efficiently support exact and range queries on the data. Our structure is easy to implement and uses $O(\log n)$ memory space for each of its n nodes. The exact match query can be performed, as in other works, by $O(\log n)$ message passing. The main contribution of this paper is that the structures broadcast the range query to the target nodes within at most $O(\log n)$ link traversing steps. We showed that all properties of RAQ can be maintained when nodes join the network or leave it. Also, we have discussed a probabilistic model with fewer number of links at each node. We demonstrated our claims by results of simulation.

We are currently working on other extensions of the RAQ model to reduce the complexity of the degree of network's nodes. We also intend to validate our results through experimental evaluation using real data. Other ideas can be to design a fault tolerant model to handle different faults such as the situation that the nodes abruptly leave the network. Load balancing is another important property of the RAQ to look at. In this case, we are going to study the situations that the data points are not uniformly distributed in the search space; also, the computing power of the nodes of network can be different. Further works on these topics are underway.

Acknowledgements

The authors would like to thank Cyrus Shahbi, Farnoush Banaei-Kashani and Mohammad Toossi for their supports, useful comments and reviews.

References

- [1] J. Aspnes, y. Kirschz, A. Krishnamurthy. Load balancing and locality in range-queriable data structures. in Proc. PODC the twenty-third annual ACM symposium on Principles of distributed computing
- [2] J. Aspnes and G. Shah. Skip Graphs. In Proceedings of Symposium on Discrete Algorithms, 2003.
- [3] F. Banaei-Kashani and C. Shahabi, SWAM: Small-World Access Model. In Proc. CIKM 2004, Thirteenth Conference on Information and Knowledge Management CIKM 2004, Nov 2004.
- [4] Gnutella, <http://gnutella.wego.com>
- [5] J. Gao and P. Steenkiste, An Adaptive Protocol for Efficient Support of Range Queries in DHT-based Systems. Proc. 12th IEEE International Conference on Network Protocols (ICNP'04), Berlin, Germany, Oct. 2004.
- [6] N. J. A. Harvey, M. B. Jones, S. Saroiu, M. Theimer, and A. Wolman. SkipNet: A Scalable Overlay Network with Practical Locality Properties. In Proc. of Fourth USENIX Symposium on Internet Technologies and Systems, 2003.
- [7] D. R. Karger and M. Ruhl. Simple Efficient Load Balancing Algorithms for Peer-to-Peer Systems. In ACM Symposium on Parallelism in Algorithms and Architectures, June 2004.
- [8] J. Kleinberg. The small-world phenomenon: An algorithmic perspective. In Proc. 32nd ACM Symposium on Theory of Computing (STOC 2000), pages 163-170, 2000.
- [9] D. Malkhi, M. Naor, and D. Ratajczak. Viceroy: A scalable and dynamic emulation of the butter y. In Proc 21st ACM Symposium on Principles of Distributed Computing (PODC 2002), pages 183-192, 2002.
- [10] H. Nazerzade. Making Querical Data Networks Navigable. In Proc. ICI 2004 International Conference on Informatic, Sep 2004.
- [11] S. Ratnasamy, P. Francis, M. Handley, and R. M. Karp. A scalable content-addressable network. In Proc. ACM SIGCOMM 2001, pages 161-172, 2001.
- [12] S. Ratnasamy, J. Hellerstein, and S. Shenker. Range Queries over DHTs. Technical Report IRB-TR-03-009, Intel Research, 2003.
- [13] A. I. T. Rowstron and P. Druschel. Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems. In IFIP/ACM International Conference on Distributed Systems Platforms (Middleware 2001), pages 329-350, 2001.
- [14] R. Seidel. Exact upper bounds for the number of faces in d-dimensional Voronoi diagrams, DIMACS Series, volume 4. American Mathematical Society, 1991.
- [15] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In Proc. ACM SIGCOMM 2001, pages 149-160, 2001.