# QoS routing for P2P networking

Janne Linnolahti
Helsinki University of Technology
Department of Computer Science
janlintu@luukku.com

## Abstract

Growth of dynamic P2P (peer-to-peer) systems is almost unique in the history of Internet. The massive use of such systems could almost be comparised to the growth of Internet itself.

This survey reviews the field of Quality-of-Service-routing. First traditional IP-routing concept and it's QoS-routing factors are introduced. Then some IP-routing QoS-algorithms are examined, as well as some related techniques like resource reservation protocol (RSVP), differential services (DiffServ) and multiprotocol label switching (MPLS), traffic engineering, and admission control.

Then we dive into the peer-to-peer networks - where P2P features, performance issues and algorithms will be researched. Some P2P systems will be looked closely. Then some P2P searching and routing techniques are researched deeper. Overlay routing and efficient information retrieval for example. Then we look the QoS challenges for P2P-systems including interoperability problems, dynamic QoS interconnection.

It is shown that P2P is not a solution for all the problems, instead it has problems itself, security and reliability for example. When operating anonymous in the Internet, the security issues are always a problem.

## 1 Introduction

"Peer-to-peer" concept is the core of today's Internet but it will become even greater factor in the future Internet. P2P-networks grow on their own by flooding and updating the information using peer-to-peer concept. The peers (storage and computing power) are used to make the routing and data-searching faster.

This paper focuses to QoS routing and to P2P networking. The basis of QoS-based IP-routing is introduced in section 2. Techniques like resource reservation (RSVP), differential services (DiffServ) and multiprotocol label switching (MPLS) are briefly introduced.

Section 3 introduces Peer-to-Peer (P2P) scheme. The QoS features in existing IP-routing algorithms are not enough for P2P model, so we need to study some p2p-routing algorithms and -techniques. Overlay routing and auxiliary networks are created using various techniques. Efficient information retrieval is important part of the QoS-routing in efficient P2P-networks.

## 2 QoS for traditional IP-routing

### 2.1 IP-Routing Basis

IP-routing in today's Internet is focused on connectivity and supports only one type of datagram service called "best effort". It tries its best to forward user traffic, but can provide no guarantees regarding loss rate, bandwidth, delay, delay jitter, etc. Packets could get jammed in congestion and dropped due that.

"Best effort" service is not enough for today's real-time applications. Services like Video on-Demand require higher bandwidth and QoS.

#### 2.1.1 QoS Background Work

IETF has set up a QoS Routing Working Group (QOSRWG) to guide the research on QoS routing techniques. It defined a framework of QoS routing in the Internet [1]. Nowadays this QOSRWG is divided to different routing subcategories, for example diffserv, intserv, mpls, rap and rsvp. Each of these has their own charter-pages under IETF web site to collect their work together. There are linked all the internet-drafts and request-for-comments performed by these workgroups.

#### 2.1.2 QoS Definition and Concept

Quality-of-Service is a set of service requirements to be met by the network while transporting a flow [1]. QoS is service delivered to network users, which can be calculated in some measurable metrics, i.e. bandwidth, delay, jitter, cost or loss probability. Flow requirements need to be expressed using these corresponding metrics. QoS routing is needed to maximize the network utilization and improve the total throughput of the network.

#### 2.1.3 Important QoS Routing Factors

QoS routing is supposed to solve or avoid the problems mentioned before. The main objectives of QoS routing are:

- To meet the end-users QoS requirements.
- To optimize the network resource usage.
- Degrade network performance gracefully when things like congestion happen. When network is in heavy load, QoS routing is expected to give better performance (i.e. better throughput) than best-effort routing, which can degrade the performance dramatically.

Several issues related to QoS Routing: [10]

- How is QoS capability of each link determined, and resources reserved.
- What are the routing metrics used (hop count, available bandwidth delay)
- What is the granularity of routing decision (source/destination address or per flow).

**Information exchange**

One important issue is how often the routing information is exchanged between the routers. QoS routing needs to exchange more information than "best-effort" routing. QoS information such as available bandwidth needs to be exchanged. Second, the metrics used by QoS routing could be changing very quickly. If the routing information is exchanged every time the values of metrics change, it will cause a great burden for the network links and routers — consuming network bandwidth and routers' CPU cycles.

**Methods for choosing the path**

One way is to set a threshold to distinguish significant changes from minor changes. The information is exchanged only when a significant change occurs. By doing so, it can also bring stability of the QoS routes. Another method is to consider only the available resources after reservation, instead of the actual available resources based on bandwidth calculations.

**Information Maintenance**

A related problem is how to maintain the information collected. Two solutions are given:

- Keep only the routing table for best-effort traffic, and compute the paths for QoS flows on demand.
- Aggregate the flows and maintain only the information about aggregated flows.

**Administrative Control**

There are also some control issues regarding QoS.

- Flow priorities and preemption

  Different flows in the network have different QoS requirements, and should have different priorities. Critical flows should have higher priority than other flows. When voice or video flow resources are not enough, these critical flows can take the resources from flows with lower priority.

- Resource control

  In the network that has multiple different flows (i.e. DiffServ), the resources should be allocated fairly among all the flows. The starvation of lower priority flows must be avoided.

- These control schemes should be included in QoS routing. Chen [3] proposed a routing algorithm to fairly share the resources.

## 2.2 Routing Algorithms

The current Internet routing protocols are based on two routing algorithms — Distance vector algorithm and Link-State algorithm. In Distance vector algorithm, neighboring routers exchange routing information periodically. So every router can learn the routing information from others. Based on that information, the shortest path to every destination can be computed. [1]

In Link-State algorithm, every router advertises its link state information to the whole network, so every router can receive the link-state information. Such information is maintained in a local database of every router, from which the routing table is calculated using Dijkstra's shortest path algorithm. The advertising is periodically triggered by events. [1]

### 2.2.1 QoS Routing Algorithms

QoS routing algorithms should be efficient and scalable to large networks, but not too complex.

QoS-routing algorithms can be divided into 3 categories:

- hop-by-hop routing
- source-based routing
- hierarchical routing algorithm

They are classified according to the way how the state information is maintained and how the search of feasible paths is carried out. [3]

In hop-by-hop routing, each router forwards packets to the next-hop router, towards the destination.

In source routing, every router has global state information about the network, and the path is locally selected based on the state information. After the path is determined, the source router notifies the other router along that path how to forward the traffic flow. Then the flow will be routed to the destination accordingly.

Hierarchical routing is most suitable for large network. Routers are organized into logical groups in multiple levels. Routing information is at the border nodes of each group. Every node contains the detailed information about its group and integrated information about other groups. PNNI is a typical example of hierarchical routing.

### 2.2.2 QoS Multicast Routing

Multicast routing is suitable for multimedia, real-time applications, such as video-conferencing and video-on-demand. QoS routing is important factor of multicast routing.

The goals of QoS-based multicast routing are as follows:

- Scalability to large groups with dynamic membership
- Robustness in the presence of topological changes
- Support for different receiver-initiated reservations
- Support for shared reservation styles, and
- Support for "global" admission control, i.e., administrative control of resource consumption by the multicast flow.

[1]

There exist two different multicast routing solutions, source- and receiver-oriented. In multicast source routing the multicast trees are computed by the first-hop router from the source, based on sender traffic advertisements. The advantage of this scheme is that it is compatible with the current RSVP signaling model. When receiver reservations are homogeneous, this scheme works well. The disadvantages of this scheme is that it is difficult to support heterogeneous reservations.

In receiver-oriented multicast routing the sender first multicast its advertisements over a best-effort tree which can be different from the QoS tree for sending data. Then, each receiver gets the advertisements and independently computes a QoS path from the source, based on the receiver reservation. In other words, multicast path computation is broken up into multiple concurrent unicast path computations. This scheme supports heterogeneous reservations well. [1]

Chen [3] has comparised various multicast routing algorithms.

## 2.3 QoS Routing Techniques

QoS routing has to work together with many techniques. QoS-related techniques are i.e. Resource Reservation Protocols (RSVP), DiffServ and MPLS. These techniques are briefly introduced. These QoS-related techniques are researched in IETFWG's [11] routing working groups.

### 2.3.1 Resource Reservation

QoS routing and resource reservation protocols complement each other. Resource reservation provide a method for requesting and reserving network resources, while QoS routing allows the determination of a path that can accommodate the requested QoS. Combining a resource reservation protocol with QoS routing allows fine control over the route and resources at the cost of additional state and setup time. [1]

### 2.3.2 DiffServ

Differentiated Services [4] is a classification based mechanism to provide QoS for the Internet, and solve the scalability problem that exist in IntServ.

DiffServ was designed to provide a simpler, coarser approach to establishing differentiated classes of service for Internet data. DiffServ pushes complexity to the edge of the network (edge routers) where traffic classification and conditioning would take place. [10] QoS routing route the flows to the paths which have the capacity to accept the flows, or simply reject a flow if there is not bandwidth for it.

DiffServ is useful also in the P2P-routing, in fact the modification of it has been used in Kazaa where the participation levels (introduced later) are a similar technique.

### 2.3.3 Multiprotocol Label Switching Architecture

MPLS [9] has emerged as an all encompassing solution which provides QoS and Traffic Engineering for the Internet. MPLS began as a low cost routing solution, but has gradually grown in scale to include other aspects like constraint based routing and traffic engineering. [10]

MPLS is developed for reducing the complexity of forwarding in IP networks. It is particularly an approach for achieving the simplified connection-oriented forwarding characteristics of layer 2 switching technologies while retaining the equally desirable flexibility and scalability of layer 3 routing. MPLS introduces a new forwarding paradigm for IP networks in a very scalable and cost effective way. [10]

For example, packets arriving from different ports or -routers can be labeled differently.

### 2.3.4 Admission Control

Flow Admission Control (FAC) determines whether a link or a node has sufficient resources to satisfy the QoS required for a flow. FAC is typically applied by each node in the path of a flow during flow set-up to check local resource availability. [1]

Higher-level admission control determines whether or not a flow set-up should proceed, based on estimates and policy requirements of the overall resource usage by the flow. Higher-level admission control may result in the failure of a flow set-up even when FAC at each node along the flow path indicates resource availability. [1]

Higher level admission control is needed to control that the cost of the networking doesn't rise too high. The fairness might be a problem too, if smaller flows steal all the resources. So a certain level of acceptance rate need to be to be guaranteed to the larger flows.

# 3 Peer-to-Peer Networking

## 3.1 Definitions of P2P

The term "peer-to-peer" refers to a class of systems and applications that employ distributed resources to perform a function in a decentralized manner. The Intel P2P working group defines P2P as "the sharing of computer resources and services by direct exchange between systems". P2P is about sharing. A peer gives some resources and obtains other resources in return. [12]

P2P model is a communication model in which each party may act as a client and a server. Both parties have the same capabilities and both can initiate connection. It is the opposite of client/server model in the sense that there is no central entity that the other parties contact but every single entity is able to initiate connection directly with all other entities. [2]

## 3.2 Overview of P2P

Peer-to-peer systems are gaining popularity quickly due to their scalability, fault-tolerance, and self-organizing nature. Progress in P2P has been made in applications such as storage [30], DNS, media streaming [31], collaborative Web server [32], distributed content-based search [33], and even distributed firewalls [34]. Nowadays P2P-applications are more used than traditional desktop and client-server systems. P2P-system can act like a distributed computing system, it can serve the collaboration network and the most known applications are the file sharing applications. The P2P community share files using P2P-applications (Kazaa, DC++, Gnutella, Napster).

P2P model is more efficient than the traditional client/server approach. P2P allows nodes to function as both a client and a server. Bandwidth is used more wisely when machines communicate directly with each other using P2P applications. P2P reserves less bandwidth, since the flows are from peer to peer, not from the peer to server. P2P reduces the cost of owning the systems and the content.

Security is one of the biggest challenges in P2P systems. Viruses could be shared by naming them after some interesting software and when user downloads and launches the file, the virus is activated. Reliability is problem, due to anonymous nature of P2P networks.

The user action can be modeled as follows: User A connects to the central machine (Napster) or to a supernode-peer (Kazaa) and request something from it. A gets answer and decided to download thing X from computer B. Then the connection is established direct from A to B, until the download is complete.

## 3.3 Features Of P2P Systems

- Enables peer-to-peer communication and collaboration.
- Shared CPU processing improves performance in power calculations.
- Scalability is defined by how many systems and users can be supported.
- Reliability is related to systems and network failure, disconnection, availability of resources, etc. Large scalability improves reliability.
- Shared disk space (files sharing systems) gives better probability for the data to be discovered.
- Sharing the costs between all the users of community - lower the machine-, the energy- and the bandwidth costs of the system providers.
- Increases privacy and autonomy by letting the users stand anonymous behind their machines.
- P2P systems typically do not rely on established infrastructure - they build their own.

## 3.4 Performance Issues

P2P-routing updates incurred a lot of overhead compared to the physical network routing (IP-routing) in terms of network hops. But self-organization and adaptation is required to handle the changes caused by peers connecting and disconnecting from the P2P systems. This increase the scalability, fault-resilience, and the cost of ownership. Three key approaches to optimize performance [12]:

- Replication puts copies of objects/ files closer to the requesting peers. Changes to data objects have to be propagated to all the object replicas.
- Caching reduces the path length required to fetch a file and due that the amount of messages exchanged between the peers.

  Replication and the caching works quite same way. Both store files to a peers to increase the data availability.

- Intelligent routing and network organization:

  To fully realize the potential of P2P networks, it is important to understand and explore the social interactions between the peers. Ramanathan [13] determine good peers based on interests, and dynamically manipulate the connections between peers to guarantee that peers with a high degree of similar interests are connected closely. Establishing a good set of peers reduces the number of messages broadcast in the network.

  In combination with intelligent routing, replication helps to minimize the distance delay by sending requests to closely located peers.

## 3.5 P2P-Routing Models and Algorithms

Peer group management includes discovery of other peers in the community and location and routing between those peers. Discovery of peers can be highly centralized such as in Napster, highly distributed such as in Gnutella, or somewhere in between.

### 3.5.1 Centralized Directory Model

Peers connect to a central server and publish information about the content they offer for sharing. Upon request from a peer, the central index will match the request with the best peer in its directory that matches the request. The best peer could be the one that is cheapest, fastest, or the most available, depending on the user needs. Then user prefers some QoS metric, usually the bandwidth and then a file exchange occurs directly between the two peers. [12]

Examples for centralized directory model are Napster and SETI@home. The problem in such scheme is that it still has some juridical challenges. At central server must not have any information about where copyrighted data could be downloaded.

When charging policies and standards are included in the Internet routing, the QoS-service levels could be implemented in p2p-systems. Then if user want fast service they pay more. If user prefers cheaper QoS they wait longer (or forever).

People should be responsible what they do in Internet. So the anonymity should be taken away by forcing the people to authenticate when using p2p-systems. Still there is problems if some people steal others authentication id and use it in wrong way. So security is a key topic when the charging policies are standardized.

### 3.5.2 Flooded Requests Model

The flooding model is pure P2P model in which no advertisement of shared resources occurs. Instead, each request from a peer is flooded (broadcast) to directly connected peers, which themselves flood to their peers, until the request is answered or a maximum number of flooding steps (typically 5 to 9) occur. This requires a lot of network bandwidth.

If the goal is to reach all the peers in the network, then this model does not prove to be very scalable, but it is efficient in limited communities.

Some companies (eg. Kazaa) have developed "super-peer" client software (Kazaa-concept later in this paper), that concentrates lots of the requests. This leads to much lower network bandwidth requirement, at the expense of high CPU consumption. Caching of recent search requests is used to improve scalability. [12]

The QoS-priorisation could be handled using participation levels. The more interesting your data is to others the better service you get. If others are not interesting to download the data you share, it doesn't matter how much you offer.

#### 3.5.3 Document Routing Model

The document routing model is the most recent approach (used by FreeNet for example). When a document is published (shared) on such a system, an ID is assigned to the document based on a hash of the document's contents and its name. When a peer requests the document from the P2P system, the request will go to the peer with the ID most similar to the document ID. This process is repeated until a copy of the document is found. Then the document is transferred back to the request originator, while each peer participating the routing will keep a local copy.

Document routing model faces the problem that the document IDs must be known before posting a request for a given document. Hence it is more difficult to implement a search than in the flooded requests model. Also, network partitioning can lead to an islanding problem, where the community splits into independent subcommunities, which don't have links to each other. [12]

Four main algorithms have implemented the document routing model: Chord [24], CAN [19] (Content-Addressable Network), Tapestry [18], and Pastry [17]. The goals of each algorithm are similar: to reduce the number of P2P hops to locate a document of interest and to reduce the amount of routing state that must be kept at each peer.

Each of the four algorithms either guarantees logarithmic bounds with respect to the size of the peer community, or argue that logarithmic bounds can be achieved with high probability.

### 3.6 The Concept of Kazaa

What was illegal in Napster was the central database having the indices of the users shared data. Today's most used P2P-system Kazaa Media Desktop (KMD) has more sophisticated working model.

#### 3.6.1 Supernodes

Users with the fastest Internet connections and the most powerful computers are automatically designated as Supernodes. A Supernode contains a list of some of the files made available by other KMD users and where they are located. When you perform a search, your KMD first searches the nearest Supernode to you, and then sends you immediate results. This first Supernode then refers your search to other Supernodes and so on. [29]

This process is designed to make searching as fast as possible and means that your computer will never search through all the files made available by KMD users, only through the files that have been indexed by the Supernodes to which you are connected.

If you can't find a file you are looking for, try again tomorrow. You will most likely connect to a new Supernode, listing a different set of files made available by other users. [29]

#### 3.6.2 Participation Level

The key to peer-to-peer is participation and sharing. The more each individual participates and shares files, the better the experience is for everyone.

Kazaa Media Desktop includes a Participation Level. This is a title or index assigned to each user based on the way in which they use the software. Basically, the more integrity rated files you share, and are uploaded from you, the better your downloading performance will be. The aim is to reward people who 'integrity rate' their files and share content they have created so it can be uploaded - which makes the peer-to-peer world work. This can be a little hard to control since it is up to other users to choose to download files from you, but if you create and share interesting files you will give yourself a good chance of improving your level. When you request a file from another user and someone else has already requested it, the user with the highest Participation Level will receive priority. If you are the only person attempting to download a file your Participation Level does not matter. [29]

### 3.7 Important P2P-Routing Aspects

#### 3.7.1 Overlay Routing

Distributed hash table (DHT)-based overlay networks, represented by Pastry, CAN, and Chord, offer an administration-free and fault-tolerant application-level overlay network. Still they have some disadvantages. Relying on application-level routing may cause network delays and bandwidth consumption. And they typically construct a homogeneously structured overlay even though nodes in these networks usually have varying physical connectivity and packet-forwarding capacities.

Brocade [21] removes some of the constraints in systems such as Pastry [17] and Tapestry [18] and Topologically-aware CAN [20] by constructing a secondary overlay network of supernodes that are situated near the network access points such as routers. Though Brocade improves performance, it still uses logical routing in the secondary network, which incurs several physical hops for every logical hop.

Xu [23] constructs auxiliary network called expressway. By using physical proximity, forwarding capacity, node availability and node connectivity they can increase routing performance. They use the Autonomous System (AS)-level topology extracted from BGP (border gateway protocol) reports, and a novel landmark numbering strategy that can deal with changing network conditions and enable proximity neighbor selection.

Nodes that are aggregated are logically close to each other rather than physically close. Route advertisement uses a variant of the distance vector algorithm [22]. Route summarization technique is used to reduce routing state while trading off routing performance.

Simulation results show that taking underlying physical network characteristics into consideration for routing in the overlay network improves performance. The approach achieves close to optimal, shortest-path routing performance. The previous approaches Brocade and eCAN [16] stay within 2 to 6 times the performance of optimal routing. Brocade-like systems can approach optimal routing performance only when aggressive caching is used. [23]

Expressway routing outperform systems like Brocade in terms of multicast performance because expressway removes the constraints imposed by the logical structure of the overlay, and better approximates the underlying physical network. In fact, constructing an overlay that closely approximates the physical network makes it possible to deliver routing performance that is better than default IP routing. This has been shown by RON [27] and Detour [28]. [23]

#### 3.7.2 Efficient Information Retrieval

Efficient information retrieval is important from the QoS point of view. An efficient P2P information retrieval system called PeerSearch [25] has been proposed in a research. PeerSearch avoids the scalability problem of existing systems that employ centralized indexing, index flooding, or query flooding. It also avoids the non-determinism that exhibited by heuristic-based approaches. PeerSearch achieves both efficiency and determinism through an elegant combination of index placement and query routing. Given a query, PeerSearch only needs to search a small number of nodes to identify matching documents. Using IR algorithms such as vector space model (VSM) [26] and latent semantic indexing (LSI) [26], PeerSearch represents documents and queries as vectors and measure the similarity between a query and a document as the cosine of the angle between their vector representations. PeerSearch stores a document index in Content-Addressable Network (CAN) using its vector representation as the coordinates, resulting in those indices stored close to each other are also close in semantics. This unifies the problem of content- or semantic-based search with routing in an overlay network. [25]

Several features distinguish PeerSearch from other IR systems.

- PeerSearch works in a completely decentralized manner. There is no single point of failure and no complex hierarchy.
- PeerSearch supports content and semantic searches expressed in natural language, as opposed to simple keyword match.
- PeerSearch is scalable, efficient, and effective. Both indexing flooding and query flooding are avoided. The CAN routing is augmented with expressways [23], an optimization to overlay networks such as CAN. PeerSearch's effectiveness stems from the state-of-the-art IR algorithms.
- The use of vector-space transformation improves the DHT routing efficiency.
- The use of semantic-based indexing to solve various problems, including deep search in Grid, resource discovery, etc.

PeerSearch is among the first systems to allow decentralized, deterministic, and non-flooding P2P information retrieval based on content and semantics, and is also the first to apply semantic knowledge to solve various network/system problems at such a deep level. [25]

### 3.8 QoS challenges for P2P-systems

#### 3.8.1 Dynamic QoS Interconnection

Current technical and business models are not ready for supporting the next generation Internet with QoS interconnection. Only the connectivity is taken into account. This leads to a need to create policies for connections that support QoS-routing needs. Pricing policies need to be specified too.

The next generation Internet requires more control over QoS selection and inter-domain routing. Hwang et al. proposed inter-domain market-managed interconnection management mechanism, which fulfills these requirements. It supports QoS control, traffic engineering, charging, and QoS service co-ordination. The solution requires two technologies: BMP (Bandwidth Management Point) and a modified version of BGP (Border Gateway Protocol). The modified version of BGP provides additional routing exchange information such as price and QoS level specifications. Using the price and QoS attributes in the BGP UPDATE, the exchanged information can be used by the BMPs to manage and control routing policy for many different types of interconnection settlements. [15]

The analysis show that market-managed interconnection and support of inter-domain QoS routing with charging capability can increase the overall welfare of users and service providers.

#### 3.8.2 Interoperability

Although many P2P systems already exist there is still no support to enable these P2P systems to interoperate. Some of the requirements for interoperability include:

- How do systems determine that they can interoperate
- How do systems communicate, e.g., what protocol should be used, such as sockets, messages, or HTTP
- How do systems exchange requests and data, and execute tasks at the higher level, e.g., do they exchange files or search for data
- How do systems determine if they are compatible at the higher protocol levels, e.g., can one system rely on another to properly search for a piece of information

- How do systems advertise and maintain the same level of security, QoS, and reliability

The P2P Working Group [14] is an attempt to gather the community of P2P developers together and establish common ground by writing reports and white papers that would enable common understanding among P2P developers.

# 4 Conclusions

P2P is an important technology but not a solution to every problem in the future of computing. Alternatives to P2P are traditional technologies, such as centralized systems and the client-server model. Systems can participate in different degrees in the centralized/ client-server/ P2P paradigms. P2P is a strong alternative for scalability, anonymity, and fault resilience requirements. From the market perspective, cost of ownership may be the driving factor for P2P.

One big aspect including the topic is the music industry which tries to fight against copyright breakers - without good results because Napster has taught present P2P-companies how to be irresponsible of the actions what users do with their P2P-systems. While the copyright- and charging standards are still far away from reality, the users of P2P-systems are able to be like partners in crime without fear of being accused.

As the Internet grows toward a network of QoS networks, the inter-domain QoS interconnection will become one of the critical factors in the next generation Internet. The standards at this area are still incomplete and there are many open questions. By solving the interoperability problems among various P2P-networks the QoS will be radically improved. Due to the complexity of the problem, QoS-routing technologies for P2P-networks are still in the evolving research stage.

# References

[1] E. Crawley, R. Nair, B. Rajagopalan, H. Sandick, "A Framework for QoS Routing in the Internet", RFC 2386, Aug. 1998, 37 pages, http://www.ietf.org/rfc/rfc2676.txt ww.ietf.org/rfc/rfc2386.txt

[2] ORAM, A. Peer-to-Peer: Harnessing the Power of Disruptive Technologies. 2001.

[3] Shigang Chen, "Routing Support for Providing Guaranteed End-to-End Quality-of-Service", Ph.D. thesis, UIUC, May 1999, 207 pages, http://cairo.cs.uiuc.edu/papers/SCthesis.ps

[4] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, An Architecture for Differentiated Services , RFC 2475, December 1998.

[5] R. Braden, D. Clark, S. Shenker, "Integrated Services in the Internet Architecture: an Overview", RFC 1633, Jun. 1994, 33 pages, http://www.ietf.org/rfc/rfc1633.txt

[6] Braden, R., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) – Version 1 Functional Spec", RFC 2205, September 1997.

[7] "Quality of Service - Glossary of Terms", May 1999, 26 pages, http://www.qosforum.com/white-papers/qos-glossary-v4.pdf

[8] G. R. Ash, J. S. Chen, A. E. Frey and B. D. Huang, "RealTime Network Routing in a Dynamic Class-of-Service Network", Proceedings of ITC 13, 1992

[9] E. Rosen, A. Viswanathan, and R. Callon, Multiprotocol Label Switching Architecture, RFC 3031, January 2001.

[10] A Singh, G Mittal, QoS and Traffic Engineering: MPLS, DiffServ and Constraint Based Routing

[11] Active IETF Working Groups, http://www.ietf.org/html.charters/wg-dir.html

[12] D. Milojicic, V. Kalogeraki, R. Lukose, K. Nagaraja, J. Pruyne, B. Richard, S. Rollins, Z Xu, HP Laboratories Palo Alto HPL-2002-57 (http://www.hpl.hp.com/techreports/2002/HPL-2002-57.pdf)

[13] Ramanathan, M. K., Kalogeraki, V. Pruyne. Finding Good Peers in the Peer-to-Peer Networks. International Parallel and Distributed Computing Symposium, Fort Lauderdale, Florida, April 2002.

[14] Peer-to-peer Working Group. www.p2pwg.org, 2001

[15] J Hwang , J Altmann, H Oliver, A Suárez. Enabling Dynamic Market-Managed QoS Interconnection in the Next Generation Internet by a Modified BGP Mechanism, October 2002

[16] Z. Xu and Z. Zhang, "Building Low-maintenance Expressways for P2P Systems," Hewlett-Packard Labs, Palo Alto, CA, Tech. Rep. HPL-2002-41.

[17] A. 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 (Middle-ware), November 2001, pp. 329-350.

[18] B. Y. Zhao, J. D. Kubiatowicz, and A. D. Joseph, "Tapestry: An infrastructure for fault-tolerant wide-area location and routing," UC Berkeley, Tech. Rep. UCB/CSD-01-1141, April 2001.

[19] Ratnasamy, S., Francis, P., Handley, M., Karp, R., Shenker, S. 2001. A Scalable Content-Addressable Network. Proceedings of the SIGCOMM, pp 161-172.

[20] S. Ratnasamy, M. Handley, R. Karp, and S. Shenker, "Topologically-aware overlay construction and server selection," in IEEE INFOCOM 2002, June 2002.

[21] B. Y. Zhao, Y. Duan, L. Huang, A. Joseph, and J. Kubiatowicz, "Brocade: Landmark Routing on Overlay Networks," in 1st International Workshop on Peer-to-Peer Systems (IPTPS'02), 2002.

[22] M. E. Steenstrup, Ed., Routing in Communications Networks. Englewood Cliffs, NJ: Prentice Hall, 1995, ch. Distance Vector Routing, ISBN: 0130107522.

[23] Z. Xu, M. Mahalingam, M. Karlsson, Turning Heterogeneity into an Advantage in Overlay Routing, Hewlett-Packard Laboratories, Palo Alto, CA, Tech. Rep. HPL-2002-126.

[24] I. Stoica, R. Morris, D. Karger, M.F. Kaashoek, and H. Balakrishnan, "Chord: A scalable peer-to-peer lookup service for Internet applications", ACM SIGCOMM ’01, San Diego, CA, USA, 2001.

[25] PeerSearch: Efficient Information Retrieval in Peer-to-Peer Networks C. Tang, Z. Xu, M. Mahalingam, Internet Systems and Storage Laboratory HP Laboratories Palo Alto, HPL-2002-198, July 2002.

[26] M. Berry, Z. Drmac, and E. Jessup. Matrices, vector spaces, and information retrieval. SIAM Review, 41(2):335-362, 1999.

[27] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris, "Resilient Overlay Networks," in 18th ACM Symposium on Operating Systems Principles (SOSP), October 2001, pp. 131-145.

[28] S. Savage, T. Anderson, A. Aggarwal, D. Becker, N. Cardwell, A. Collins, E. Hoffman, J. Snell, A. Vahdat, G. Voelker, and J. Zahorjan, "Detour: a Case for Informed Internet Routing and Transport," IEEE Micro, vol. 19, no. 1, pp. 50-59, January 1999.

[29] http://www.kazaa.com/

[30] F. Dabek, M. F. Kaashoek, D. Karger, R. Morris, and I. Stoica, "Wide-area cooperative storage with CFS," in 18th ACM Symposium on Operating Systems Principles (SOSP ’01), October 2001, pp. 202-215.

[31] V. N. Padmanabhan, H. J. Wang, P. A. Chou, and K. Sripanidkulchai, "Distributing streaming media content using cooperative networking," in 12th International Workshop on Network and Operating Systems Support for Digital Audio and Video, May 2002.

[32] V. N. Padmanabhan and K. Sripanidkulchai, "The case for cooperative networking," in 1st International Workshop on Peer-to-Peer Systems (IPTPS), March 2002.

[33] C. Tang, Z. Xu, and M. Mahalingam, "pSearch: Information retrieval in structured overlays," in 1st Workshop on Hot Topics on Networks (HotNets-I), October 2002.

[34] A. Keromytis, V. Misra, and D. Rubenstein, "Sos: Secure overlay services," in 2002 ACM SIGCOMM Conference, August 2002, pp. 61- 72.