Quality of Service in the Internet - the issues of debate

Mika Ilvesmäki
Helsinki University of Technology
Laboratory of Telecommunications Technology
P.O. Box 3000
02015 HUT, Finland
Email:lynx@tct.hut.fi

Abstract

In this work, we discuss some of the basic concepts relating to Quality of Service in the Internet. We start by defining the concepts of Quality of Service and Class of Service. We briefly look at service classes and proceed to observe the Internet router architecture and the enhancements that need to be or have been implemented to enable different service levels. We then move to deal with two cardinal issues: the introduction of state to the core of Internet and the contradictory position between the user and the network in triggering the Quality of Service. We find the issue of Quality of Service in the Internet being versatile and subject to speculation.

1 Introduction

Bringing Quality of Service (QoS) to the Internet is an issue of controversy. Traditionally Internet has been the network for all with equal opportunity to use its resources, but the use of QoS would change this. The recent developments of software and the emerging new services with increasing commercial efforts suggest that QoS, or at least different levels of service, Class of Service (CoS), should be introduced to the Internet. The emergence of various bandwidth hungry and delay sensitive applications, such as Voice over IP and video conferencing, require, or at least benefit from QoS or some other form of prioritization guarantees. Similarly, the probable growth of new QoS sensitive applications [1] using the Internet protocol might expect some sort of QoS guarantees from the network. Several IETF[1] workgroups, such as IntServ[2] and DiffServ[3], have participated in the discussion and definition of Internet service architectures, but their work has not yet reached to any conclusive solutions. Various architectural and technological solutions have emerged and the heated debates for and against these solutions have dominated the discussion the on the future Internet. Nevertheless, what is Quality of Service, what does it mean in the Internet and what are the various aspects that need to be addressed should QoS ever take its place in the 'Net'.

In the current Internet, there are no guarantees for either relative or absolute QoS and it is debatable if we can ever expect the Internet to provide absolute end-to-end QoS [2, 3, 4, 5]. The contemporary Internet is characterized by the diversity of its networking technologies. In the core of the network ATM and Frame Relay, both able to offer QoS [6], are slowly pushing FDDI–solutions to the background. This trend would implicate that the core Internet could be able to offer some kind of service level as the penetration of QoS capable technologies reaches an adequate level. However, in the edges of the network the multitude of network solutions is overwhelming. All the different LAN technologies, some capable to offer QoS and some not, create a substantial obstacle to an absolute end–to–end Quality of Service in the Internet. Furthermore, the ever growing diversity in access technologies, such as the introduction of xDSL techniques and the strong foundation of traditional PSTN–modem solutions suggests that offering a possibility of consistent QoS in the Internet would introduce a number of problems regarding, for instance, the definition of QoS parameters.

The rest of this work is structured as follows: In Section 2, we will first define the concepts of Quality and Service and their relation in the Internet and introduce the concept of Class of Service. In Section 3 we will take a look at how the contemporary Internet key component, the Internet router, functions, how it should be developed in the future and what efforts have already been started. Then we look at some of the founding issues that need to be dealt with, before any kind of QoS can be brought to the network. Particularly, in Section 4 we will deal with the concept of state and how, and in what way, it could be intro-

[1]http://www.ietf.org/
[2]http://www.ietf.org/html.charters/intserv-charter.html
[3]http://www.ietf.org/html.charters/diffserv-charter.html

duced and used in the Internet. Following, in Section 5 we define the fundamental question of whether there should be Quality or Class of Service in the Internet and debate on who should initiate or trigger the use of a service class in the network.

2 Definition of QoS and CoS

Quality, in the Internet frame of reference, could be understood as the combination of exactly defined measures such as data loss, delay, jitter and use of network resources associated with the feeling or notion of Quality that the user of the network experiences. Major difficulty lies in defining the Quality as a function of both the measures **and** the human factor.
Service is defined as something that the users use, either explicitly (interactive Internet services like telnet, ftp or www) or implicitly (various 'invisible' applications like name services etc.). The actual definition of the service can be done in various ways:

- By protocol identifiers (in the IP header)
- By (the parts of) the addresses of the sender and/or receiver (in the IP header)
- By TCP/UDP source and/or destination port numbers (in the TCP/UDP header)
- By the ingress point of the traffic in the network
- By a combination of any of the above

Different ways of service definition serve different purposes. In some instances it could be nice to group services according to the applications that are used (port numbers) or by the user of the service (sender's IP address), and in other cases a coarser service definition might be enough. The actual policy with which packets of particular service are treated might be identified with the IP precedence bits [7]. See Figure 1 for IPv4 header structure and one option of the intended functionality of the IP precedence bits[4].
In this light, the Quality of Service in general terms and when speaking of networking, means that the user of some service receives a predefined, but not necessarily a constant amount of resources from the network guaranteeing that the user's packets are delivered to their destination within the set parameters and performance bounds.
On a related issue, Class of Service (CoS) is a closely related concept to QoS. Using CoS instead of QoS means that the traffic of one user is treated better than the traffic of another. No absolute guarantees are given, only promise to differentiate traffic. The

[4]Sometimes referred to as TOS field, as in Type of Service.

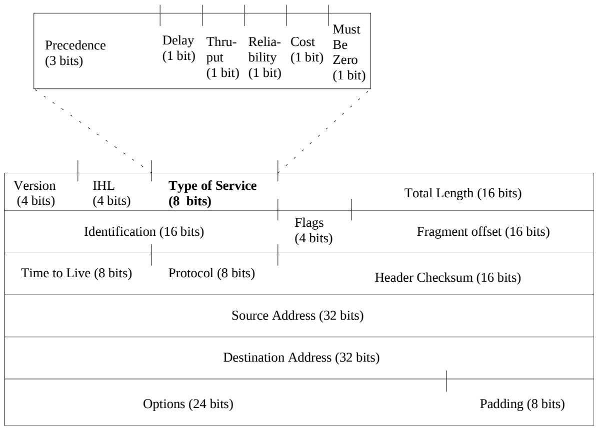

Figure 1: IPv4 header and the Precedence bits

Table 1: IP precedence values

IP precedence values	
1000	minimize delay
0100	maximize throughput
0010	maximize reliability
0001	minimize monetary cost
0000	normal service (the default value)

Class of Service solution will most probably be the concept first deployed in the Internet before the actual end-to-end absolute Quality of Service.

2.1 Service classes

To realize the QoS or CoS in the network we have to introduce service classes[8]. The service classes define the parameters with which the traffic is handled in the network. No clear consensus exists how many service classes there might be: Some suggest up to eight service classes [7] while others are content with three or four classes [9, 10]. In [11] the semantics of the IP precedence field values in the IP header are defined to be as shown in Table 1.
It could be stated that the number of different service classes is somewhat secondary to the more important question: Who or what decides which kind of service is to be given to a particular traffic flow[5][8]. We will come later back to this question in Section 5.

3 Internet router - the key component

The Internet could be said to consist, in a very basic level, of routers and connections between them.

[5]The term flow refers to traffic streams representing a particular user or application[8].

This is illustrated in Figure 2, where edge (or access) routers send their traffic to each other via the core routers. Bringing service differentiation or QoS to the Internet requires several changes in the router architectures and especially requires broadening the ways we think an Internet router should function. Extending the functionality of Internet routers is one of the essential issues in bringing QoS to the network.

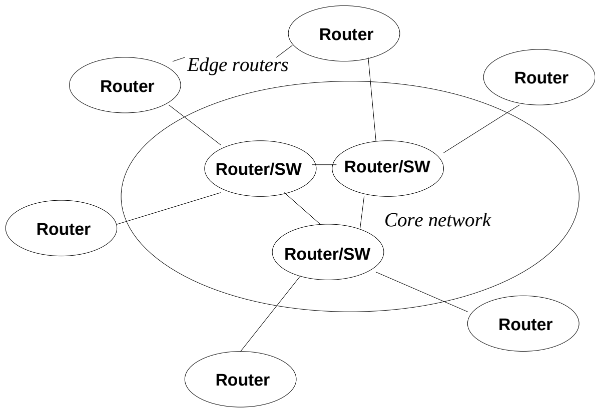

Figure 2: Basic Internet structure

3.1 Conventional routers

Conventional router solutions tend to make forwarding decisions for each packet separately even if the decisions for a series of packets would be identical. Information is queued packet by packet to the routing/forwarding processor on the IP datagram level, realizing the OSI layer three functionality. This queue is traditionally served based on first come first serve (FCFS) scheduling creating the so-called best effort service. All the links connected to the router are utilized from the common pool by the route processor on a need to use -basis. This process is illustrated in Figure 3 where the router first takes the packet, looks up the destination address using routing tables and then finally sends the packet to the appropriate network interface according to the routing information obtained from the routing tables.
Hence, a traditional Internet router inherently utilizes statistical multiplexing since its resources are shared between all users. In the case of actual queue overflow (packet drop), the resource usage is further controlled by the TCP flow control in the sending hosts. TCP flow control is designed to aid in dividing the resources as fairly as possible between all users [12].

3.2 Current trends in router design

The amount of traffic and number of users in the Internet are increasing. This increase calls for im-

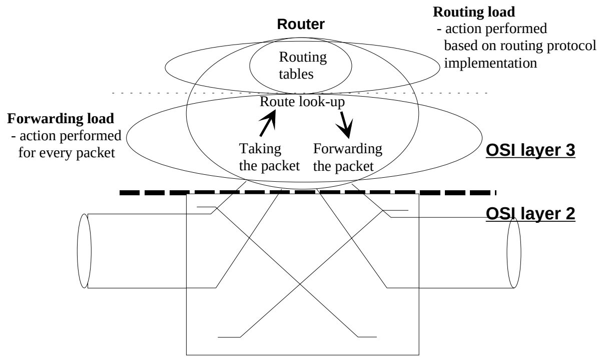

Figure 3: Conventional routing procedure

provement of performance in the network elements for IP packet forwarding and routing, the Internet routers. The first natural need is to be able to either forward the increased amount of packets faster or somehow find a way to reduce the amount of packets offered to the forwarding process. In addition to increased traffic and the growing number of users, there is also a clear need for different levels of service emerging amongst the users of the network. Whether this need arises from using applications sensitive to performance parameters (bandwidth, delay, jitter) or from the general need to receive increased and better performance than other users in the network (for instance, priority traffic in corporate networks) it is clear that some form of traffic classification needs to be done in order for the network to be able to offer the different service levels to the traffic[13].
The improvements that have been proposed in relation of router architectures and bringing Quality of Service to the Internet relate

- To enhancing router performance by implementing some or all of the functionalities of the routers in hardware instead of software[14, 15, 16, 17, 18, 19, 20, 21, 13, 22, 23]; this is the so-called Gigabit–approach.
- To speeding up the access to routing tables, where advanced search methods are used [24, 25, 26, 27, 28, 29]
- or to using sophisticated ideas of mapping a large number of packets to a single identifier[15, 30, 31, 32, 33, 34, 35, 36], usually in the OSI layer 2; a solution referred to as multi-layer switching or IP switching.

In addition, QoS based routing is being investigated so that preferred paths could be determined for prioritized traffic[37, 38].
The Internet router is the key component in the current Internet as it provides both the routing and

packet forwarding functionality in the network. It seems likely that the dominant position of the router will continue and it will also be providing mechanisms for realizing Quality of Service in the future Internet. The various components in routers that provide the tools and mechanisms for Quality of Service [12] and should therefore be implemented are:

- Traffic shaping, to provide predictable traffic behavior,
- Admission control, to limit excessive resource usage,
- and Differential congestion management, providing preferential treatment to particular traffic classes in times of congestion in the network.

4 Bringing state into the Internet

The traditional paradigm of Internet has been stateless. Routers do not keep connection state information. This is to improve the robustness of the communication system, routers are designed to be stateless, forwarding each IP packet independently of other packets. Consequently, redundant paths can be exploited to provide robust service in spite of failures of intervening routers and networks. All state information required for end-to-end flow control and reliability is implemented in the hosts, in the transport layer or in application programs. All connection control information is thus co-located with the end points of the communication, so it will be lost only if an end point fails.
Bringing QoS or consistent classes of service independent of the traffic in the network we have to introduce some kind of state, if not to all of the Internet routers, at least to some. There exist different kinds of state: local state, soft-state and hard-state. The basic function for the state is to spread around the information on the relevant QoS / CoS related decisions that have been made. The goal is to achieve consistency in the network regarding, for instance, policy and routing decisions.
The discussion circulates mainly around whether to bring state to the Internet or not. The actual protocol implementations are already many and will most probably be, even in the future, ISP dependent choices. Between Autonomous Systems(AS) the change of state information and related data will most probably be communicated using the existing routing protocols with, perhaps, slight modifications.

4.1 Local state

The basic form of state, the local state, is already used in the Internet routers. Some routing decisions or policies are cached for quick access. The single device may be speeded up by introducing local state, for instance, in the accessing of routing tables the most frequent address lookup decisions could be cached or the incoming packets could be mapped to a separate path already in the input port of the router, bypassing the forwarding stage altogether (refer to Figure 3). The neighboring devices need not to know anything about these alterations and network can be argued to function a bit more efficiently. This kind of local state is closely related to the advanced router architectures mentioned earlier. Local policy decisions could also be defined for particular service classes and types (=applications).
The downside of local state is that the neighboring node, or the rest of the network in general, does not know anything about the state of an individual node. This may lead to inconsistencies in service levels. Also, if the local decisions differ in the edge routers (refer to Figure 2) the core routers may need to maintain large policy tables to achieve consistency in the network. This introduces problems in scaling and router management.

4.2 Soft State

Soft state, simplified, means that two or more devices negotiate and agree on some property that will be remembered in both of the devices for an arbitrary length of time. Some routing protocols and recent suggestions for utilizing connection oriented link layer technologies have brought the concept of soft state to the Internet. Also many of the routing protocols utilize soft state; the routing tables are updated and link states or distance vectors are updated as timers expire. No device utilizing the soft state may function under the assumption that neighboring devices know about the state beyond the timeout value. No explicit state teardown messages exist when using soft state and the state is lost, or brought down as timeout values are met without the request the renew the state information. Bringing the state up is also, in many cases, only implicitly acknowledged. In the Integrated Services architecture the protocol used to reserve the resources (RSVP) utilizes the concept of soft state [9, 10].

4.3 Hard State

Extending the soft state timeouts to infinity and adding explicit state teardown messages we have defined the concept of hard state. Hard state is most commonly found in telephone networks and BISDN

(ATM) systems. Systems using hard state usually have complex signaling protocols exist. The complexity is necessary because hard state systems assume that the state information is present in other devices and they base their operation on this assumption. Therefore necessary mechanisms to guarantee the state need to exist.
In the Internet, the hard state means that throughout the path in the network, some service (related) parameters are set and resources are allocated for a certain traffic flow. The parameters are reset and resources freed only when they are explicitly told to do so. Hard-state is being brought to the Internet as a by-product of using ATM as link layer (OSI layer2) technology [1, 39].

4.4 End-to-end or hop-by-hop

In close relation to using state in the Internet is the question whether the state information should be spread around to the whole packet path (end-to-end) or to only parts of the paths (hops). Usually the packet paths extend over many ASs that may or may not have identical thoughts on policy. In the current Internet the number of router hops between two end hosts may rise to dozens or even higher. This may introduce additional processing delays and could possibly be avoided by using state in the routers. The more widespread the state is in the network the less we have unexpected phenomena in the course of packet forwarding and routing.
The hop-by-hop has been the traditional approach; the hop is understood to be just the hop from one router to another. At the moment it seems that it would be easier to start implementing the state in a hop-by-hop fashion and then later on extend the definition of hop to include parts of the path as one hop or even the whole path as one hop. This is the approach that is suggested in DiffServ architecture [40]. Within one Autonomous System (AS) consisting of several routers there is a similar policy for traffic flows.
In the IntServ approach, the RSVP has been suggested to function in an end-to-end manner[9, 10]. This would require updating of the whole Internet to support RSVP functionality and would likely introduce scalability problems in the core of the network.

4.5 Connectionless vs. connection oriented

As we are thinking in spreading the QoS / CoS decision info we need to also take into account the layer 2 that actually transports the data. The emergence of connection oriented technologies, specifically ATM, offer at least the possibility to QoS, though they do not require the use of QoS parameters and even provide services or means for this kind of traffic (UBR in ATM, CIR set to zero in Frame Relay). To allocate resources in connection oriented environment we need advanced Connection Admission Control functionality and modern queue management algorithms and methods. The difficulty lies in defining the parameters of IP traffic to characterize the traffic properties and this characterization affects the overlaying Internet Technologies; TCP flow control, for instance.
In the connectionless environment we can provide for Classes of Service, differentiated levels of service with the aid of priority queues and related algorithms.
In the current Internet the diversity of networking technologies also suggests that the thought of connectionless environment and the subsequent implications might be more readily applied.
The problem of providing absolute or relative guarantees of delivery depends also very much on the underlying network techniques. Even if we have just one hop of Ethernet on an otherwise ATM infested path, we can not anymore say anything definite on the Quality of Service bounds that the path has.
However, whether we are using connectionless or connection oriented technologies we have still two important questions left unanswered. Do we want to use QoS or CoS in the Internet and who do we let to make the requests and decisions on service levels.

5 User vs. Network

Finally, we come back to one of the most important questions relating to Quality of Service and the Internet. Who or what decides that a particular traffic flow is to be prioritized. One of the most important questions regarding the state in the Internet, is do we have to scale the state information with the individual user or with the applications that users are using or with some other aggregated form of traffic. Service aggregation is an issue that has not been studied much. A simplified suggestion for service aggregation is presented in Figure 4, where in a connection oriented environment (ATM) services are mapped to Virtual Paths containing Virtual Circuits for individual traffic flows.
The assignment of prioritization has traditionally been done in two ways:

1. The network issues prioritization with no, or little regard to user wishes. An example of such scheme, in a manner of speaking, might be the contemporary Internet.
2. The user requests either explicitly or implicitly resources from the network. The network then notifies the user if such requested resources exist and reserves these resources. An example of

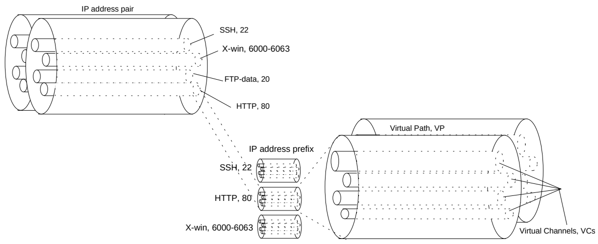

Figure 4: Aggregating services in the Internet

> implicit user request of prioritization might be the PSTN. When the user picks up the handset the nearest exchange reserves a timeslot to the connection. This timeslot could be understood as absolutely prioritized to the user.

The fundamental issue of whether we should assign QoS to an IP packet flow implicitly by network elements or explicitly by user requests is debated quite extensively in [8] where the user initiated resource allocation is slightly favored, though network initiated QoS is seen to have some advantages.
If the user triggers the QoS she, most likely, wants also to measure it. This strongly suggests that the given QoS level should be absolute and measurable. This implicitly requires that the individual user has her QoS parameters stored throughout the packet path. Furthermore, this means that every node has to store information on every flow that passes through. Many say that this per-flow state may never be practical or feasible in the Internet or that only a very limited set of connections might have their individual states in the network. Furthermore, in the case of user triggered QoS there certainly may be situations when the service request can not be fulfilled. A situation that does not exist in the current Internet as such.
Via network triggered resource allocations the QoS can be relative, CoS. Difficulty lies in building appropriate business models and service architectures. It may be hard to convince a customer that she is being serviced better than the next one. This implies that QoS should be geared towards groups of individual users, enterprises and such. One problem arises in determining the set of services that are to be classified to various service classes. Only using fixed sets of applications limits the service to those applications[8]. Some work has been done to automate the construction of service sets, Network Service Profiles[41, 42]. The aggregation of traffic flows under similar policy identifiers is also a problematic issue, since we can not for sure know what kind of traffic needs what kind of parameters[8] at different instances.

6 Conclusions

It should be obvious that even in the Internet there are no free lunches; if we give resources to a group of users or services we unavoidably take the resources away from another group. Put in another way: Bringing QoS to the network leads to a situation where the once fair network is favoring some services at the price of neglecting others. This is the case, when resources become limited and the little amounts of resources have to be divided. To control the usage of resources and limiting the number of users that wish to use the theoretically scarce resources a service architecture is needed. Examples of these kinds of efforts are the Integrated and Differentiated service architectures respectively. In addition, pricing also plays a major role when determining who wants, or is able to in the first place, to use QoS in the Internet.
It may be questionable if we are ever going to see absolute QoS on a flow by flow basis. In the future we may actually have applications (or users) which know exactly what kind of QoS they need, and with the aid of IPv6 [43] and RSVP [9] we may be able to assign explicit QoS to a session on–demand. Then, it is debatable if the network will ever have the capacity to provide for per-flow absolute QoS. More likely, we will see aggregates of services being prioritized with different weights. This way the adaptation is in balance both the network and the applications are flexing with their performance.
The essntial issue is the per-flow state scalability. The reasonable state table requirements in the edges of the network (access routers) grow up to be impossible demands in the core of the network. This effect is further enhanced with the recent trends in traffic patterns where 80 % of the traffic is forwarded outside of the LANs. This means that there can not exist single customers, or per-flow state, in the core of the network; not at least in the wide sense and the majority of Internet users will be offered only differentiated CoS. Also the question of connectionless and connection oriented techniques is answered if per-flow state is limited. Traffic is prioritized in routers and connectionless technologies suffice as the transport media.
The actual use and implementation of state, whether it be soft-state or hard-state is indifferent. Soft state might be lighter to implement and requires no network wide 'flag day'.
In addition, the end-to-end solution is not necessary in the CoS scheme but we can concentrate on prioritizing traffic on a hop-by-hop basis. This also makes interoperability between ISPs less of an issue and gives time to old networking technologies to be updated. After the technology has penetrated the definition of hop might extend from router- to-router, to AS-to-AS to the end-to-end solution.

Whatever the final solution, QoS or CoS, the pricing, business models, service architectures need to carefully accounted for and should be used to gear the use of QoS to those applications (and users) that can not exist without service guarantees.

7 Acknowledgements

This work was sponsored by the IPANA-project in the laboratory of Telecommunications Technology and the author is also sponsored by the Nokia Foundation.

References

[1] Eugenio Guarene, Paolo Fasano, and Vinicio Vercellone. IP and ATM integration perspectives. *IEEE Communications Magazine*, January 1998.

[2] M. Borden, E. Crawley, B. Davie, and S. Batsell. *Integration of Real-Time Services in an IP-ATM Network Architecture.* August 1995.

[3] S. Shenker, C. Partridge, and Guerin R. *Spesification of Guaranteed Quality of Service.* February 1997.

[4] Eric Crawley, Raj Nair, Bala Rajagopalan, and Hal Sandick. *A Framework for QoS-based Routing in the Internet.* March 1996.

[5] Bala Rajagopalan and Raj Nair. *Quality of Service (QoS) -based Routing in the Internet - Some Issues.* October 1996.

[6] *Traffic Management Specification Version 4.0.* ATM FORUM, April 1996.

[7] Roger Bohn, Hans-Werner Braun, Kimberly C. Claffy, and Stephen Wolff. Mitigating the coming internet crunch: Multiple service levels via precedence. March 1994.

[8] Scott Shenker. Fundamental design issues for the future internet. *IEEE Journal on Selected Areas in Communications*, 13(7):1176–1188, September 1995.

[9] Paul P. White. Rsvp and integrated services in the internet: A tutorial. *IEEE Communications Magazine*, 35(5):100–106, May 1997.

[10] Paul P. White and Jon Crowcroft. The integrated services in the internet: State of the art. Technical report, University College London, 1996.

[11] P. Almqvist. *Type of Service in the Internet Protocol Suite.* July 1992.

[12] Paul Ferguson and Geoff Huston. Quality of service in the internet: Fact or fiction. 1998.

[13] S. Keshav and Rosen Sharma. Issues and trends in router design. *IEEE Communications Magazine*, 36(5):144–151, May 1998.

[14] Guru Parulkar, Douglas C. Schmidt, and Jonathan S. Turner. *IP on ATM: a Strategy for Integrating IP with ATM.* Washington University.

[15] Peter Newman, Greg Minshall, Tom Lyon, and Larry Huston. IP switching and gigabit routers. *IEEE Communications Magazine*, pages 64–69, January 1997.

[16] Fahad Hoymany and Daniel Mossé. Switchborne router for high performance packet forwarding of connectionless traffic in atm networks. In *ICCCN '96*, October 1996.

[17] Fouad A. Tobagi. Fast packet switch architectures for broadband integrated services digital networks. In *Proceedings of the IEEE*, volume 78, pages 133–167. IEEE, IEEE, January 1990.

[18] Tong-Bi Pei and Charles Zukowski. Putting routing tables in silicon. *IEEE Network Magazine*, pages 42–50, January 1992.

[19] Robert J. Walsh and Cüneyt M. Özveren. The GIGAswitch control processor. *IEEE Network*, pages 36–43, January/February 1995.

[20] Ahmed Tantawy, Odysseas Koufopavlou, Martina Zitterbart, and Joseph Abler. On the design of a multigigabit IP router. *Journal of High Speed Networks*, 3(3), 1994.

[21] Pankaj Gupta, Steven Lin, and Nick McKeown. Routing lookups in hardware at memory access speeds. Technical report, Stanford University, 1998.

[22] Craig Partridge, Philip P. Carvey, Ed Burgess, Isido Castineyra, Tom Clarke, Lise Graham, Michael Hathaway, Phil Herman, Allen King, Steve Kohalmi, Tracy Ma, John Mcallen, Trevor Mendez, Walter C. Milliken, Ronald Pettyjohn, John Rokosz, Joshua Seeger, Michael Sollins, Steve Storch, Benjamin Tober, Gregory D. Troxel, David Waitzman, and Scott Winterble. A 50–gb/s IP router. *IEEE/ACM Transactions on Networking*, 6(3):237–248, June 1998.

[23] Dan Decasper, Zubin Dittia, Guru Parulkar, and Bernhard Plattner. Router plugins a software architecture for next generation routers. In *Proceedings of SIGCOMM '98*. IEEE/ACM, IEEE/ACM, 1998.

[24] Mikael Degermark, Andrej Brodnik, Svante Carlsson, and Stephen Pink. *Small Forwarding Tables for Fast Routing Lookups*. Luleå University of Technology, 1997.

[25] Jon Turner Marcel Waldwogel, George Varghese and Bernhard Plattner. Scalable high speed IP routing lookups. In *Proceedings of SIGCOMM '97*, pages 25–36. ACM/SIGCOMM, ACM/SIGCOMM, 1997.

[26] Stefan Nilsson and Gunnar Karlsson. Fast address lookup for internet routers. In Paul Kuhn and Roya Ulrich, editors, *Broadband Communications: The Future of Telecommunications*, pages 11–22. IFIP, Chapman & Hall, 1998.

[27] E. Besson and P. Brown. Performance evaluation of hierarchical caching in high-speed routers. In *IEEE 1998 Global Telecommunications Conference*. IEEE, IEEE, November 1998.

[28] Henry Hong-Yi Tzeng. Longest prefix search using compressed trees. In *IEEE 1998 Global Telecommunications Conference*. IEEE, IEEE, November 1998.

[29] E. Filippi, V. Innocenti, and V. Vercellone. Address lookup solutions for gigabit Switch/Router. In *IEEE 1998 Global Telecommunications Conference*. IEEE, IEEE, November 1998.

[30] Yasuhiro Katsube, Ken-Ichi Nagami, and Hiroshi Esaki. *Cell Switch Router - Basic Concept and Migration Scenario*. Toshiba Corporation, 1996.

[31] Hiroshi Esaki, Shigeo Matsuzawa, Akiyoshi Mogi, Ken-Ichi Nagami, Tatsuya Jinmei, Toru Kon'no, and Yasuhiro Katsube. Cell switch router (CSR) – label switching router supporting standard ATM interfaces. In Seyhan Civanlar and Indra Widjaja, editors, *Broadband Networking Technologies*, volume 3233, pages 2–10. SPIE, SPIE, November 1997.

[32] Yakov Rekhter. Tag switching architecture - overview. In Seyhan Civanlar and Indra Widjaja, editors, *Broadband Networking Technologies*, volume 3233, pages 11–19. SPIE, SPIE, November 1997.

[33] Hassan M. Ahmed, Ross Callon, Andrew G. Malis, and John Moy. IP switching for scalable IP services. In *Proceedings of the IEEE*, volume 85, pages 1984–1997. IEEE, December 1997.

[34] Bruce Davie, Paul Doolan, and Yakov Rekhter. *Switching in IP Networks – IP Switching, Tag Switching and Related Technologies*. Morgan Kaufmann Publishers, Inc., 1998.

[35] Peter Newman, Greg Minshall, and Thomas L. Lyon. IP switching – ATM under IP. *IEEE/ACM Transactions on Networking*, 6(2):117–129, April 1998.

[36] Arun Viswanathan, Nancy Feldman, Zheng Wang, and Ross Callon. Evolution of multiprotocol label switching. *IEEE Communications Magazine*, 36(5):165–173, May 1998.

[37] George Apostolopoulos, Roch Guérin, Sanjay Kamat, and Satish K. Tripathi. Quality of service based routing: A performance perspective. In *Proceedings of SIGCOMM '98*. IEEE/ACM, IEEE/ACM, 1998.

[38] Liang Guo and Ibrahim Matta. On state aggregation for scalable QoS routing. 1997.

[39] Philip Dumortier. Toward a new IP over ATM routing paradigm. *IEEE Communications Magazine*, January 1998.

[40] Florian Baumgartner, Torsten Braun, and Pascal Habegger. Differentiated services: A new approach for quality of service in the internet. In Harmen R. Van As, editor, *High Performance Networking*. Kluwer Academic Publishers, 1998. IFIP TC-6 Eighth International Conference on High Performance Networking (HPN'98) Vienna, Austria, September 21-25,1998.

[41] Mika Ilvesmäki, Marko Luoma, and Raimo Kantola. Flow classification schemes in traffic–based multilayer IP switching — comparison between conventional and neural approach. *Computer Communications*, 21(13):1184–1194, September 1998.

[42] Mika Ilvesmäki, Raimo Kantola, and Marko Luoma. Adaptive flow classification in IP switching: The measurement based approach. In Raif O. Onvural, editor, *Internet Routing and Quality of Service*, volume 3529 of *Proceedings of SPIE*. SPIE, SPIE, November 1998.

[43] William Stallings. *IPv6: The New Internet Protocol*. 01.09.1997 1997.