Building An Accounting Infrastructure for the Internet

Wenjia Fang {wfang@cs.princeton.edu}

Princeton University

Abstract: This paper explores the issues involved in providing of an accounting infrastructure for the Internet. It first argues that accounting, absent in the current Internet architecture, is crucial to the success of the Internet in the future. Then it focuses mainly on the design issues involved in such provision. Finally, it provides a possible scheme in implementing such an accounting infrastructure.

1.0 Introduction

Today's Internet has evolved far beyond the expectation and scope of a government research project: it has turned into a commercial, global and integrated service network. Conceived and funded as a research project, the Internet was designed with defense and interoperability of heterogeneous networks in mind. As a result, some characteristics of a commercial network, for example, elaborate resource management and an accounting infrastructure, are absent.

By accounting, we refer to the measurement of traffic profiles of the Internet and the attribution of such profiles to the corresponding users. It is not to be confused with billing, nor pricing. Billing refers to the process of compiling the accounting information and charging the users for such usage, whereas pricing refers to the formation of prices for different types of services, usually through elaborate economic models. An accounting infrastructure is a necessary but not sufficient condition for usage-based billing and pricing. Rather, billing and pricing are two higher level artifacts on top of an accounting infrastructure. That is, usage-based billing and pricing do not necessarily follow usage-based accounting information. For example, most american households pay flat-fees for local telephone services even though the local telephone company has precise accounting information about the local phone calls.

The purpose of an accounting infrastructure, often misunderstood, goes beyond usage-based billing and pricing: it provides valuable information on traffic profiles of the network, which serves a significant role in further network designs, expansion and reengineering. In this paper, we mainly concern ourselves with the design issues of an accounting infrastructure for the Internet. More precisely, the scalability, granularity and mechanisms involved in the addition of such infrastructure to an existing distributed packet network. We are not concerned with billing and pricing.

The paper is organized as follows: Section 2 presents our motivations. Section 3 defines taxonomy and discusses the design issues. Section 4 describes an proposed usage-based accounting infrastructure and further research directions. Section 5 lists related work and section 6 offers concluding remarks.

2.0 Motivations

2.1 Design goals of the Internet

The Internet was conceived in late 1960s as the ARAPNET, a defense research project funded by the Advanced Research Project Agency (ARPA). The primary design goal of the ARPANet was "survivability in the face of failure" [6]. Owing to its military nature, the Internet gave lower priority to an accounting infrastructure. For the most part, accurate usage-based accounting is absent from the Internet, with the exception of a few specific research access points to the network [5].

As a result, most pricing and billing on the current Internet is access-based, that is, network users are charged according to their type of access to the Internet. They are primarily of the following two forms:

- By access pipe

  Most Internet Service Providers, or ISPs, provide dedicated lines to enterprises and charge according to the bandwidth the dedicated lines can support. The fees usually include a fixed start-up fee and a flat monthly fee for connection cost.

- By access time

  This is the predominant form of individual Internet access provided by ISPs. This can be seen as a variation on the access pipe, as the bandwidth is limited by the speed of the modem from 2400bps to 28.8K bps. The form of payment is also based on a flat fee. The users pay a flat monthly fee for unlimited access or a fixed amount for a given number of hours and additional for each extra hour.

Clearly, access-based accounting and pricing fail to accurately reflect the network traffic profile generated by the users.

2.2 Current trends

The past five years have seen significant changes on the Internet. First, NSF relinquished its control and funding of the NSFNet in 1993, thereby lifting access restrictions. As a result, an ever growing user base regularly accesses the Internet has increased congestions. Second, in 1996, MCI upgraded the original NSFNET backbone, the primary cross-country links of the Internet, from T3 lines (45Mbps) to a vBNS (very high Backbone Network Service) running at 622 Mbps. Moreover, many multimedia, real-time applications have been developed for the Internet. In contrast to their text-based counterparts, real-time multimedia applications demand much higher network bandwidth, usually on the order of megabits per second.

In short, the nature of the Internet has changed: it is no longer a network funded by a government agency, dedicated to research purposes. It has turned into a commercial network composed of a wide range of constituents, with a complex topology and a diverse user base. The telephony and the computer networking communities have reached a consensus that the future Internet will be an Integrated Service Packet Network (ISPN), which offers a sophisticated service model supporting a wide variety of applications for a global user base [1, 13].

2.3 Usage-based accounting

In contrast to access-based accounting, a usage-based accounting infrastructure holds end users accountable for the cumulative network resources they used when generating traffic. Such infrastructure is necessary in the Internet due to two reasons: it provides a more economically efficient way of recovering cost of providing network services, and it provides feedback to the users so that they can be more conscientious of the scarce resources. Unfortunately, neither is provided by the current access-based accounting and pricing.

The incredibly large carrying capacity of an optical fiber leads some to conclude that in the future bandwidth will be abundant and therefore there will be no need for accounting and billing for using bandwidth. We contend that this is not the case. Bandwidth in the future will be much higher by today's standards, but so will be the demand. The question is whether the increase of bandwidth will be able to keep up with the demand generated from real-time multimedia applications as well as from a growing user base; we believe this is unlikely.

The bandwidth required by any real-time multimedia applications is significantly higher than their text-based counterparts. A full-motion video, even with the most advanced lossy compression, requires bandwidth in the order of 10Mbps[1]. For the sake of arguments, the vBNS backbone networks offer 622Mbps, which can support roughly 60 simultaneous streams. The Internet is already subject to frequent congestions, and congestions are likely to be a serious problem in the future.

Access-based accounting does not reflect actual network traffic cost, or traffic burden generated by each user. When network bandwidth is abundant, this is not a problem, as all packets are delivered. However, when network bandwidth is limited and congestion occurs, the high cost of delivering a packet is not reflected in an access-based accounting system. In other words, an access-based accounting, coupled with a first-come-first-serve service model, fails to reflect the actual supply and demand of limited resources and does not deter frivolous usage.

A more profound ramification of the access-based accounting is "the tragedy of the commons": when public goods are accessed without any regulations, the public risks losing access to the goods completely. The current phenomenon of the Internet Telephone exemplifies the argument against access-based accounting and flat-rate pricing. The Internet Telephone takes advantage of the flat-rate service charge of ISPs, whereas regular telephone service is largely usage-based and heavily regulated. The advantage of a flat-rate Internet telephone service, however, will not be long-lasting: it will be offset by degraded service caused by increased congestions. The Internet Telephone will no longer be as attractive once a usage-based accounting infrastructure is in place and the internet services are properly priced to reflect the actual cost incurred.

3.0 Design space

The packet-switched nature of the Internet poses challenges for a usage-based accounting infrastructure. Unlike circuit-switched networks, where usage can be monitored by tracking the number of connections and the duration of those connections, packet-switched networks do not lend themselves to such abstraction. In a packet-switched network, state information, such as number of packets transmitted or the number of packets acknowledged, is contained in each individual packet, rather than at each switch, providing resilience against node failures. Therefore, many of the mechanisms we have learned from circuit-switched network do not transfer readily to packet-switched networks.

This section offers a discussion of several design issues, which should be addressed in any proposed accounting infrastructure. We adopt the traditional terminology for the network bandwidth providers: *Internet Service Providers, or ISPs.* We divide network bandwidth consumers into two categories, *Administrative Domains, or ADs* [9] and *End-users*. An Administrative Domain refers to a collection of hosts, users, and internal networks under a single authority, and it may choose to appear to an ISP as an single entity. An end-user is an individual who is directly requesting his network service from an ISP via a computer. For practical accounting purposes, an AD and an end-user are indistinguishable to an ISP, they only differ in how they choose to manage their resources internally.

3.1 Granularity

The most important design issue to be considered is the granularity of accounting, which can be further divided into two types:

- **Account granularity**: the end point for which traffic should be accounted.

  By each end-user.

  This is the finest granularity and the most appealing one as it holds each individual user accountable for his network traffic. However, it is not entirely feasible, as there is no unique universal identification for each end user on the Internet.

  By each IP address.

  The most natural unit of accounting, since each host on the network is uniquely identified with an IP address. This is analogous to a telephone number in the telephone network.

  By each Administrative Domain (AD).

  A possibly more desirable granularity for administrative purposes than IP addresses. As most machines within an AD share the same network address, they could be collectively identified by the tuple <IP_address, subnet_mask>. This method offers an easy way to collapse the accounting of a group of hosts within administrative domains.

- **Traffic granularity**: the unit of data transmission that should be accounted for.

  By bytes or by packets.

  The finest granularity of all and also the most appealing choice since packets on the Internet are measured in bytes.

  By flow.

  A flow is an abstraction in the packet network. It is an uninterrupted stream of packets from an application that are likely to preserve spatial and temporal locality while traversing in the network. For example, a flow in video can be all the packets transmitted between a push of the start button and a push of the stop button. Clearly, flow is a concept that is application-specific. The attributes of a flow includes its duration time (T), its required bandwidth (B), and the expected delay (D).

  The key characteristic of a flow is the physical and temporal locality of its packets. That means all packets in a flow are likely to receive similar service in terms of bandwidth and delay, and therefore, cost, along the route. This characteristic helps to consolidate the details of packet-granularity accounting and billing without losing accuracy.

  By Session.

  The concept of a session is also applications-specific. For example, in video-on-demand, a session can be any data transmission from when a video is requested, until the user wants to quit,

1. This is based on 1280x1024 pixels, 30 frames per second, 24 bits for RGB and 45:1 compression ratio.

including all the starts, stops, rewinds and forwards within a session.

A session also captures the dynamics and non-predictability of network states. It is a cruder granularity for traffic accounting.

3.2 Point of accounting and complexity

At which point of the network, shall we do accounting? Do all the routers in the Internet have to be involved or only a subset of them? Some careful observations lead us to believe that the primary purpose of accounting can be served by routers at borders of constituent networks, and the complexity of accounting tables of one ISP is only correlated to the number of its adjacent ISPs.

Figure 1 illustrates the point. ISP1 through ISP3 are separate networks, each represented by a network cloud. R1 through R7 are routers; all are border routers except R2. Let us assume for the moment that an ISP does not care which route a packet traverses within its constituency — the scope of this ISP network — then accounting is necessary when two constituencies exchange packets. In other words, we only need to account when packets leave one constituency and enter another, therefore, only border routers need to be involved in accounting. In figure 1, only R1, R3 - R7 need to be involved in accounting.

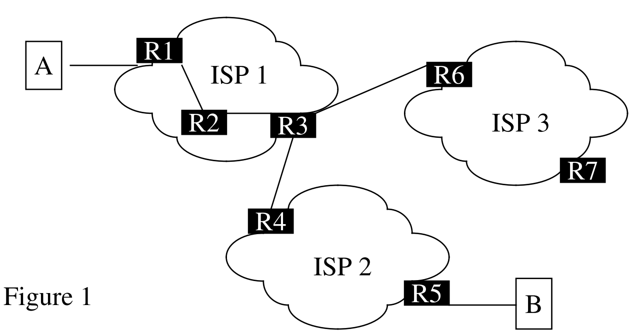

Figure 1

Another important observation is that each ISP needs to keep track of only the packets it exchanges with its immediate neighboring ISPs, without needing to know the source of packets. Consider a packet travels from host A via ISP1 and ISP2 to host B, incurring some cost. When the packet is within the constituency of ISP2, ISP2 needs not to be aware of the fact that the packet is in fact from A, it is sufficient for ISP2 to know the packet came from ISP1 for accounting purposes. At the end of the accounting cycle, ISP2 can bill ISP1 for the cost of delivery of the packet within ISP2. In turn, ISP1, which had dutifully kept track of such packet from A on R1, can charge A for both the cost incurred within ISP1 as well as ISP2.

The relaying of cost can be accomplished by the hierarchical structure of the accounting infrastructure, hence, tremendously simplifies the accounting table maintained at border routers.

3.3 Feedback Schemes

Feedback schemes are the mechanisms invoked within networks to help regulate scarce network resources. This can be done in a number of ways.

- Performance Feedback.

  This can be done by either slowing down the rate at which packets are forwarded, or dropping the packets altogether. In the extreme case of congestion, admission control will reject requests for new flows or sessions. The current Internet uses performance feedback only.

- Pricing Feedback.

  This scheme is to use pricing schemes to regulate the requests for different types of service. There are many kind of pricing feedbacks, including priority queues with different prices, real-time queues and smart-market bidding for different services in real time.

3.4 Flexible payment schemes

This refers to the fact that the true beneficiary of a transmission could be either the sender, or the receivers, or both. The decision should be left in the hands of higher level applications and users in stead of being confined and engineered into the accounting infrastructure. For example, the true beneficiary of a video-on-demand request is usually the receiver, the reasonable account to be charged for email is the sender, and in the case of an video-conferencing, the participants should be able to choose a scheme they see as appropriate, since they may choose to split the cost.

3.5 Authentications

Authentication is not a problem in accounting as long as the account granularity can be identified by network addresses. For example, an IP address or a <IP_address, subnet_mask> can uniquely identify an entity on the Internet. When an ISP provides a physical connection to an AD, it can attribute all the traffic generated from that physical link to that particular AD. Similarly, each ISP, connecting to other ISPs via some switches, can attribute the traffic exchange between itself and its neighbors unmistakably. This is analogous to the physical telephone cable connected to each household, and authentication is not necessary before a phone call can be made.

On the other hand, authentication does present a problem when the account granularity is by user account. There is no unique identification for users on the Internet currently, and therefore, it is difficult to attribute traffic generated from hosts to each user. In this case, each user needs to authenticate to hosts before they could use the Internet services. This is analogous to making a phone call using calling card from any public telephone. The telephone number does not serve to identify the user, hence, a calling ID is needed.

3.6 Quality of Service (QoS)

Quality of Service refers to the mechanism used to provide bounded delay and certain bandwidth to real-time applications, as they are much more sensitive to network conditions. It ought to be realized that QoS is not a property intrinsic to particular applications, but rather a service guarantee that can be applied to all applications in general. Decisions of whether certain performance guarantees are desirable for a particular application should be left to the users, instead of being assumed by the architecture. For example, email is not intrinsically a low-priority application, nor is video inherently a high-priority application. A user can choose an low-delay email service at the expense of delaying a video application.

3.7 Conclusions

A good accounting infrastructure is distributed, scalable and low-overhead. The challenge of building an accounting infrastructure lies more in adding it to an existing packet network than in the pure technical challenge of providing it. The success of one such accounting infrastructure is proportional to the increased utilities of a usage-based packet network over the cost of providing such accounting infrastructure. Keeping this in mind, we will never run into the unfortunate situation of providing a monolithic operating system which consumes too much computer resources itself.

4.0 MetroCard

In this section, we describe MetroCard, a system which can facilitate Internet accounting. In essence, MetroCard is similar to a toll-road system in which the tolls are dynamically adjusted to reflect the actual network congestion and delay. The MetroCard differs from a toll-road system in that it does not allow actual transfer of money (of any sort) from the user packet to the routers, therefore avoiding the overhead of authentication and verification. MetroCard only serves to record the cost incurred when a stream of packets traverses the network.

The MetroCard assumes three functions, currently absent in the Internet architecture, to be operational:

- Changes in routing protocols to have more than one possible route from the same IP address [15], as well as additional information in the routing table to include the delay and congestion situation on this route.
- Changes in IP header to incoporate an “Accounting” field.
- A new protocol in border routers to maintain and manage accounting tables and recognize “bill” packets.

4.1 MetroCard components

There are three components in the MetroCard system. The first component is an accounting field in IP header, which is a quadruple of *<Service_Level, Accu_Tolls, Budget, Max_Charge>*.

<Service_Level> is the service level the packet desires. This is not changed throughout the transmission.

<Accu_Tolls> is the tolls incurred within a constituent, and cleared each time the packet travels into a new constituent.

<Budget> is the budget of the packet. It is decremented as the packet travels. Therefore, its value is the budget remained for the rest of the travel.

<Max_Charge> is the maximum the packet is willing to pay, a upper limit of how much the user is willing to pay the tolls. It is held constant throughout the transmission. *<Budget>* is initially set to be *<Max_Charge>*.

The second component is the current tolls posted by the routers on different routes, which are dynamically adjusted to reflect the congestions and delays on this particular route. In order to implement MetroCard, the routers need to adopt more sophisticated routing and enlarge routing tables to accommodate both multiple routes for the same IP destination and its delay. Forwarding a packet would involve matching the *<Service_Level>* to the best possible route available. Of course, when there is no congestions and the delay is the minimum, there is no reasons for tolls to go beyond zero.

The third component includes a new protocol and accounting tables installed on the border routers of constituents. These tables are necessary to record the tolls incurred for a packet within the constituents. The third component also includes the bill packets, which are generated by the destination host upon receiving of a regular data packet. The bill packet is sent back to the originating host, thus conveying the final tolls incurred throughout the transmission of the packet, which is necessary for the complete accounting cycle.

4.2 An example

We illustrate the complete MetroCard system with an example. Figure 2 illustrates six steps by which an accounting cycle completes. ISP1 and ISP2 are two constituents, which are connected via two border routers R2 and R3. Host A, the originating host, is directly connected to ISP1 Host B, the destination host, is directly connected to ISP2. Steps are listed in the figure by sequenced numbers in small circles.

Step 1: Host A sends out an IP packet with IP accounting field <2, 0, 100, 100>. On a scale of 5, <*Service_Level*> 2 is considered requesting a moderate service. <Max_Budget> is set to be 100, and <Budget> is initialized to be the same. <Accu_Tolls> is 0.

Step 2: The packet traverses through ISP1 without encountering much congestion, so by the time it reaches the boundary of the constituent ISP1, the <Accu_Tolls> only shows a modest 5. The accounting field is now <2, 5, 100, 100>.

Step 3: As the packet travels into a different constituent ISP2, the border switch S resets <Accu_Tolls> field, and subtracts that from <Budget> field. In the mean time, S also records that packet owes ISP1 5 unit for the tolls. The accounting field is <2, 0, 95, 100>.

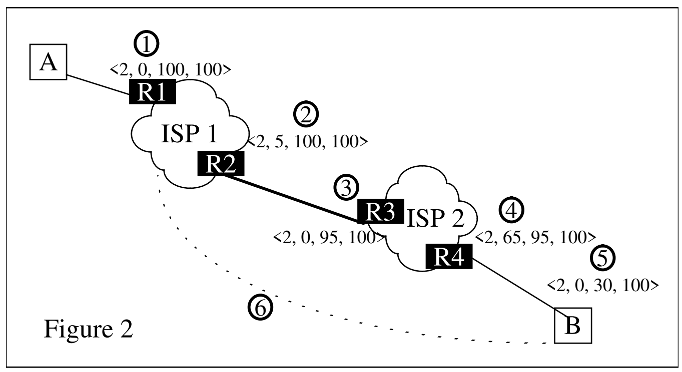

Figure 2

Step 4: Packet traverses through ISP2, encountering severe congestion, so the cost incurred is a steep 65, as indicated by <Accu_Tolls> field by the time it gets out ISP2. The accounting field is now <2, 65, 95, 100>.

Step 5: Packet arrives in the destination B. Before ISP2 hands over the packet to B, it records <Accu_Tolls> as what ISP1 owes ISP2, also subtracts that from <Budget>, leaving 30 left in the <Budget> field. The accounting field is now <2, 0, 30,100>.

Step 6: Host B receives the packet and constructs a “bill” packet which has 100 - 30 = 70, indicating that it had cost A 70 units to for this particular packet to travel from host A to host B.

Note that the bill packet is treated by the routers as a special case, the transmission of the bill packet is not charged.

Table 1 lists values in accounting tables for ISP1 and ISP2 after each above steps, respectively.

4.3 Discussion

There are several design issues needs further thinking in the MetroCard system. For example, how to deal with retransmission of data packet, how to cope with the potential loss of bill packets and how should packets without enough budget treated in the network. Also, it is not entirely trivial for the MetroCard to deal with different traffic granularity like flows or sessions.

TABLE 1. Values of accounting tables after each step in fig. 2

tables Main'd by	after step 1	after step 2	after step 3	after step 4	after step 5	after step 6
ISP 1	[A, 0]	[A, 5]	[A, 5]	[A,5]	[A,5]	[A,5]
ISP 2	[ISP1, 0]	[ISP1, 0]	[ISp1, 0]	[ISP1, 65]	[ISP1, 65]	[ISP1, 65]

5.0 Related Work

Estrin and Zhang made the first attempt to consider usage-based accounting design issues [9]. Their paper is instrumental in our ideas and we believe the issues are much more relevant today than seven years ago, when their paper was first published.

There has been some recent research in using pricing schemes to provide feedbacks to users. Cocchi [8] presented a scheme for pricing in a computer network with multiple priorities. Cocchi used computer simulations to confirm his thesis that in a network with multiple service priorities it is possible to set prices so that users of every application type are more satisfied with the combined cost and performance of a network with service-class sensitive prices than they would with flat pricing. Thus a priority pricing scheme is always achievable which will enhance total community utility.

Parris [12] *et. al* offer a scheme for real-time pricing in computer networks which can support reservation of resources. Their scheme is based on charging per real-time channel based on the resources reserved, including the type of service, time of day, and channel lifetime. The authors assume homogenous network, and reduce their analysis to a single node. It is not clear how to scale their scheme to the multiple, very heterogeneous nodes, of the Internet.

Braun and Claffy [2] realized the need for distributed and hierarchical accounting infrastructure and proposed using the precedence field in IP header to prioritize scheduling. However, without any pricing schemes associated with priorities, users are likely to specify the highest priority possible, therefore defeating the purposes of prioritizing.

MacKie-Mason and Varian [11] offer an alternative model of real-time Internet Pricing. Their model, called the "smart-market", assumes the packets will carry a bid, which reflects the price that the user is willing to pay. The prices of each route is also dynamically determined when congestion occurs. Their model is the closest to ours, but rather than looking at pricing schemes, we are more concerned with building a distributed accounting infrastructure.

Clark [6], realizing the linkage between the treatment of individual packets and the resulting overall transfer rate is not obvious from the above schemes, proposed a "service profile" scheme. Each application sets an expected profile for the packets it will send. A meter, at higher level of the protocol stack, like TCP, tags each packet as whether it is "in" or "out" of that particular profile. As the congestion situation changes, the packets from an application will collectively get more or less uniform service.

6.0 Summary

We presented our arguments for a usage-based accounting infrastructure. The increasing popularity of the Internet and the proliferation of real-time applications will render network bandwidth to be a scarce resource. The current access-based accounting infrastructure fails to attribute the usage of the network resources and reflect the actual cost of the transmission whereas a usage-based accounting infrastructure, coupled with proper pricing and billing, can regulate network bandwidth more efficiently.

We also discussed design issues of an accounting infrastructure for a packet-switched network, more specifically, the granularity, scalability and mechanisms involved. Finally, we presented the MetroCard system, a distributed accounting system which dynamically captures the network state of transmissions.

7.0 Acknowledgment

I thank many colleagues who provided valuable comments on an earlier draft of the paper: Dirk Balfanz, Yuqun Chen, Edward Felten, Mariusz H. Jakubowski, Terence P. Kelly of Princeton University and Brian Levine of UC Santa Cruz.

References:

[1] R. Braden, D. Clark, S. Shenker, "Integrated Services in the Internet Architecture: an Overview", RFC1633, http://tiger.coe.missouri.edu/~research/rfc/our-rfc-index.html.

[2] H. Braun, K. Claffy. "Network analysis issues for a public Internet", http://www.nlanr.net/Papers/hwb-menu.html.

[3] H. Braun, K. Claffy, G. Polyzos, "A framework for flow-based accounting on the Internet", http://www.nlanr.net/Papers/hwb_menu.html

[4] V. Cerf, as told to Bernard Aboba, "How the Internet came to be", http://www.zilker.net/users/internaut/internt.html.

[5] K. Claffy, G. Polyzos, H. Braun, "Traffic Characteristics of the T1 NSFNET backbone", http://www.nlanr.net/Papers/kc_menu.html

[6] D. Clark, "The Design Philosophy of the DARPA Internet Protocols", ACM SIGCOMM'88 Proceedings.

[7] D. Clark, "Adding Service Discrimination to the Internet", MIT LCS Technical Report, 1995.

[8] R. Cocchi, S. Shenker, D. Estrin, L. Zhang, "Pricing in Computer Networks: Motivation, Formulation and Example", ACM SIGCOMM'93 proceedings.

[9] D. Estrin, L. Zhang, "Design Considerations for Usage Accounting and Feedback in Internetworks", ACM Computer Communications Review, 1990 20(5), pp. 56-66

[10] D. Estrin, "Policy Requirements for inter-administrative domain routing", Technical Report, RFC1125, USC Computer Science Dept.

[11] J. MacKie-Mason & H. Varian, "Pricing the Internet", presented for conference "Public Access to the Internet", JFK School of Government, May 26027, 1993.

[12] C. Parris, S. Keshave, D. Ferrari, "A Framework for the Study of Pricing in Integrated Networks", The TENET Group technical report. ftp://tenet.berkeley.edu/pub/techreports/1992/92_018

[13] S. Shenker, "Service Models and Pricing Policies for an Integrated Services Internet", ftp://ftp.parc.xerox.com/

[14] S. Shenker, D. Clark, L. Zhang, "A Scheduling Service Model and a Scheduling Architecture for an Integrated Services Packet Network", AMC/IEEE Trans. on Networking.

[15] Shree M, J.J. Garcia-Luna-Aceves, "Congestion-oriented shortest multipath routing". Proceedings IEEE INFOCOM '96. March 1996).IEEE Computing. Soc. Press, 1996. p. 1028-36 vol.3.

[16] L. Zhang, S. Deering, D. Estrin, S. Shenker, D. Zappala, "RSVP: A new Resource ReSerVation Protocol", IEEE Network, Sept, 1993.