Provisioning for QoS in IP Networks - The Big Picture

Ulrich Fiedler

fiedler@tik.ee.ethz.ch

Computer Engineering and Networks Laboratory, ETH Zurich, Switzerland

April 15, 2002

Abstract

1 Introduction

This report tries to analyze and explain the ongoing shift of interest away from research topics related to provisioning QoS.

Provisioning in general means to provide network resources, i.e. link and forwarding capacities, to enable network services. In IP networks these network services are limited a single service class, *best-effort*. This service class transports the traffic of all applications. Best-effort means that there is no guarantee when or even if the traffic arrives at its destination. No traffic is protected against harmful cross traffic. However, some of the applications, like VoIP or streamed video, have stringent network QoS requirements.

Provisioning QoS means to provide sufficient network resources that all applications meet their QoS requirements. Provisioning QoS can be achieved either by over-provisioning the network or by means of explicit QoS support mechanisms. However, the choice of for or against explicit QoS support mechanisms is sincerely affected by how the demand of network services and technologies evolve.

We thus report trends on demand of network services and technologies and well known explicit QoS support mechanisms. Moreover, we review how the trends affect provisioning QoS. Finally, we conclude with a discussion and a summary.

2 Technological Trends

Table 1: Averaged Annualized Growth Rates (Estimation)

Technology/Demand	Annualized Growth Rates
Internet backbone link bandwidth (US)[1]	120%
Computing performance of CPU's [2, 3]	60%
Disc Storage Space[4]	110 %
Random Access Time of Dynamic RAMs [5]	10%
Internet traffic volume (US)[1]	100%
Traditional Voice Services[1]	10%
Cellular Services (US)[1]	30%-40%

In this section we look at the trends how technologies evolve and how these trend may affect research issues in the fields of network QoS support mechanisms and provisioning. Forecasting technological trends in general has a bad track record. However, predicting the pace of progress in optimizing existing technologies has worked considerably well in the past. Moore's law predicting the integration of semiconductors is probably the most prominent example. Table 1 shows some compiled figures on the progress of key technologies that affect performance of network services. For comparison we have added some figures on the progress in demand of these services.

We briefly highlight three major trends. A first trend is that bandwidth supply in the backbone out-paces demand due to progress in photo-electronics. There seems to be no such thing as the *tragedy of the commons*, i.e. new bandwidth is not suddenly saturated by new traffic (see Odlyzko [6] for further justification). This trend comes together with the fact that interconnectivity between service providers gets better. Thus there is good reason to expect that service quality improves which certainly alleviates the need for QoS support mechanisms.

A second trend is that as disk storage out-paces bandwidth. We see two implications with this trend: i) mirroring and caching content at sites close to the user, which further accelerates response and download times, is of use ii) for rich media, i.e. audio/video, a store-and-replay usage model is technologically feasible. Odlyzko [4] even argues that this model thus likely to be favored over a streaming usage model. However, feasibility of a store-and-replay usage model would considerably dilute the need for QoS support mechanisms.

A third trend is that access times for RAMs progress much slower than bandwidth supply in the backbone. Thus complex QoS support mechanisms that require access to temporal storage such as weighted fair queueing (WFQ), which requires to handle virtual queues, may run into problems. Although integrated network processors can somewhat ameliorate these problems in access networks where processing speeds is slower than in the backbone. This trend likely to hampers the deployment of explicit QoS support mechanisms in access networks. An exception to this trend are wireless links. The increasing popularity of wireless access fosters the need for explicit QoS support over lossy wireless links where bandwidth is sparse. However, this support can be done locally; thus does not require any explicit end-to-end QoS support mechanisms.

3 Explicit QoS support mechanisms

In this section we revise explicit QoS support mechanisms for end-to-end QoS support. See the book of Ferguson and Huston [7] for a more detailed description of the mechanisms revised. We do not revise any layer two technologies such as ATM and MPLS.

3.1 Integrated Services

The most simple service model to provision network for QoS is to to segregate flows with QoS requirements from any cross traffic. This model requires signaling and resource reservation along a network path to support application QoS requirements before sending the data. The Integrated Services architecture (IntServ) [8] implements such a service model [9]. Loosely spoken the architecture is kind of IP incarnation of the public switched telephone network (PSTN).

The IntServ service model proposes to implement a *guaranteed service* [10] and a *controlled-load service* [11] in addition to best-effort service. The guaranteed service is intended for applications that require hard guarantees in delay and bandwidth. The controlled-load service is intended for applications that require a sort of enhanced best-effort service. However, the integrated service model a significant drawbacks which prevented deployment: It requires every intermediate network device, including high speed backbone routers, to keep per flow states. This property is usually described as "IntServ does not scale".

3.2 Differentiated Services

The Differentiated Service (DiffServ) service model and architecture [12, 13, 14] is alternative approach to provision QoS which was proposed to overcome the problem of scalability. Diff-

Serv provides different levels of service by aggregating flows with similar QoS requirements. At the network edges, packets are classified and marked with *code-points*. Inside the network, packets are forwarded solely depending on their code-points.

The IETF proposes that services like should be built on specified per hop behaviors (PHB). They therefore specify the expedited forwarding (EF) PHB [15] and the assured forwarding (AF) PHB group [16]. EF PHB is intended to build virtual wire type of services, like e.g. *Premium Service*, to forward traffic of interactive voice and video applications. Wroclawski and Charny [17] proposed to implement guaranteed service based on EF. The AF PHB group is intended for loss differentiation among data traffic of different priority in times of network congestion. Wroclawski and Charny [17] proposed to implement controlled-load service based on AF.

To built any guaranteed or controlled-load service, we need admission control and service level agreements (SLAs) with service level specifications (SLSs) between providers along the way down any network path. See the survey paper of Xiao [18] or the Ph.D. thesis of Fankhauser [19] for details of implementation.

The problem with deployment of DiffServ to provision QoS in networks is that DiffServ still requires all-or-nothing network upgrades for providers. These upgrades trigger dramatic changes to network operations, peering arrangements, and business models [20]. Another problem is the absence of suitable means to verify service quality by users or providers. Odlyzko [6] argues that given the utilization patterns usually encountered in IP networks such as bursty traffic, short flows, low link capacity utilization, deployment is not likely. The reason is that permanent peak allocation of resources for high-prio traffic leads to astronomical costs. Alternatively, building just another over-subscribed best-effort type of service for high-prio traffic does not make much sense either.

4 Network Provisioning

No matter if with or without explicit support, provisioning QoS will only work if networks, i.e. link and buffer capacities, are adequately provisioned. However, complex procedures to provision capacity of IP networks are likely to fail in an environment where annualized growth of capacities is very large and subject to considerable fluctuations. Predicting future traffic demand is a very hard problem. The great advantage of IP, namely its great flexibility to to support new applications, aggravates this problem.

Generous over-provisioning of backbone capacities, which has become affordable due to progress in in photo-electronics, simplifies network operation. Large Providers (e.g. France Telecom [21], AT&T [6]) report that network congestion due to over-load can then be reduced to trans-oceanic links and neuralgic peering points. Moreover, congestion at peering points will be alleviated by the fact that interconnectivity between providers, i.e. autonomous systems, increases (see e.g. Vukadinovic et al. [22]). Moreover, France Telecom's Ben Fredj et al. [21] show evidence that limiting provisioning procedures to solely account for link utilization is sufficient to ensure adequate performance for TCP regulated traffic. They even suggest that the commonly used 60% utilization limit might be overly conservative. They point out that adequate provisioning procedures coupled with traffic routing strategies to avoid demand overload are "the key" to quality of service. Therefore strategies to enable load balancing and network resilience are likely to get more attention by the QoS research community than sophisticated methods that enable precise provisioning.

5 Resilient Networks

Resilient networks are another means of QoS provisioning. Resilient Networks were proposed as overlay networks in the Internet to improve robustness and availability of Internet path between participating hosts (see e.g. the RON project [23]). Participating hosts examine the condition of network paths between themselves to provide a more available and better-performing routing service than plain-vanilla IP routing. Resilient networks are able to reduce the number of observed Internet outages by a factor of three to ten, and produce considerable improvements in latency and loss [24]. Although not end-to-end principle compliant, resilient networks have caught considerable attention in the research community[24].

6 Summary and Discussion

Lately we observed a shift of interest away from provisioning QoS related problems in the QoS research community. We thus have analyzed trends how key network technologies and traffic demand evolve. Moreover, we have explained why explicit QoS mechanisms like IntServ and DiffServ are not likely to be deployed. Complex provisioning procedures to provision QoS (with our without explicit support) are likely to fail in environments like the Internet where demand is fast growing and hard to predict. Resilient networks have become a somewhat interesting alternative to explicit QoS mechanisms.

References

[1] A. M. Odlyzko, "Internet growth: Myth and reality, use and abuse," *iMP: Information Impacts Magazine*, Nov. 2000.

[2] L. Roberts, "Beyond Moore's Law: Internet Growth Trends," *IEEE Computer*, , no. 1, pp. 117–119, Jan. 2000.

[3] Open Systems Group, "Standard Performance Evaluation Corporation," http://www.spec.org/.

[4] A. M. Odlyzko, "Internet TV: Implications for the long distance network," *Television over the Internet*, 2001.

[5] B. Price, *High Performance Memories: New Architecture DRAMs and SRAMs - Evolution and Function*, John Wiley & Sons, 1999, revised edition.

[6] A. M. Odlyzko, "The Internet and other networks: Utilization rates and their implications," *Information Economics and Policy*, vol. 12, pp. 341–365, 2000.

[7] P. Ferguson and G. Huston, *Qualtiy of Service: Delivering QoS on the Internet and in Corporate Networks*, John Wiley & Sons, 1998.

[8] S. Blake et. al., "Integrated services in the internet architecture: An overview," RFC 1633, Internet Request For Comments, June 1994.

[9] S. Shenker, "Fundamental Design Issues for the Future Internet," *IEEE JSAC*, pp. 1176–1188, Sept. 1995.

[10] S. Shenker et al., "Specification of guaranteed quality of service," RFC 2212, Internet Request For Comments, Sept. 1997.

[11] J. Wroclawski, "Specification of the controlled-load network element service," RFC 2211, Internet Request For Comments, Sept. 1997.

[12] S. Blake et. al., "An architecture for differentiated services," RFC 2475, Internet Request For Comments, Dec. 1998.

[13] J.Wroclaski D. Clark, "An approach to service allocation in the internet," IETF Draft, July 1997.

[14] L. Zhang K. Nichols, V. Jacobson, "A two-bit differentiated services architecture for the internet," IETF Draft, Apr. 1999.

[15] V. Jacobson et. al., "An Expedited Forwarding PHB," RFC 2598, Internet Request For Comments, June 1999.

[16] J. Heinanen et. al., "Assured Forwarding PHB Group," RFC 2597, Internet Request For Comments, June 1999.

[17] A. Charny J. Wroclawski, "Integrated service mappings for differentiated services networks," IETF Draft, Feb. 2001.

[18] X. Xiao and L. Ni, "Internet qos: A big picture," *IEEE Network*, pp. 8–18, Mar. 1999.

[19] G. Fankhauser, "A network architecture based on market principles," ETH Ph.D. thesis No. 13750, Oct. 2000.

[20] "Qbone home page," http://qbone.internet2.edu/, Apr. 2002.

[21] S. Ben Fredj, T. Bonald, A. Proutiere, G. Regnie, J. Roberts, "Statistical Bandwidth Sharing: A Study of Congestion at Flow Level," in *proc # acm-sigcomm*, San Diego CA, USA, Aug. 2001, pp. 111–122, ACM.

[22] D. Vukadinovic, P. Huang, T. Erlebach, "On the Spectrum and Structure of Internet Topology Graphs," in *Proceedings of I2CS*, Kuehlungsborn, Germany, June 2002, ACM.

[23] "Resilient overlay networks," http://nms.lcs.mit.edu/projects/ron/.

[24] "End–to–end interest mailing list," http://www.postel.org/pipermail/end2end-interest/2001-October/, Oct. 2001.