IP traffic and QoS control

T. Bonald, S. Oueslati-Boulahia, J. Roberts
France Telecom R&D
38-40, rue du Général Leclerc
92794 Issy-les-Moulineaux Cedex 9
France
Phone: +33 1 45 29 58 82, Fax: +33 1 45 29 60 69
{Thomas.Bonald,Sara.Boulahia,James.Roberts}@francetelecom.com

Abstract

We critically examine currently proposed QoS architectures in the light of our understanding of the statistical nature of Internet traffic. We conclude that there is a need for an alternative approach based on adequate provisioning and the application of flow-level admission control.

Keywords: Quality of service, IP traffic, admission control, transparency, accessibility.

1 Introduction

Despite intense research, there is still considerable controversy concerning the definition of a QoS architecture allowing the development of a solid business model for the commercial Internet. The successively proposed Intserv and Diffserv architectures do not take adequate account of the statistical nature of traffic and consequently fail to constitute an effective solution. We deduce the need for an alternative approach based on adequate provisioning and admission control at flow level. Our arguments are presented as a series of "claims". The objective is clearly not to provide a complete proof for each of these claims, but to present some key arguments based essentially on traffic modelling and performance analysis.

It is essential to clearly distinguish "transparency" and "accessibility" QoS requirements. Transparency refers to the network's ability to be virtually invisible with respect to quality degradation experienced by end-to-end applications. Accessibility qualifies the network's availability, i.e., the long-term fraction of time during which the network is sufficiently transparent to provide a useful service. Accessibility requirements are largely orthogonal to transparency requirements. A key conclusion of our analysis is that service differentiation with respect to accessibility is both more appropriate and more easily realizable than service differentiation with respect to the degree of transparency. The envisaged QoS architecture based on the application of flow-level admission control is capable of providing different accessibility guarantees while ensuring transparency for all admitted flows.

2 Traffic characteristics

To meet transparency and accessibility QoS requirements necessitates an understanding of the statistical nature of IP traffic.

Claim 1 *Traffic demand is predictable.*

Figure 1 shows the evolution of traffic carried on an over-provisioned high capacity transatlantic link during one week and one day. The variations in rate reflect systematic changes in the activity of the user population throughout the day. The daily average busy period load is subject to statistical variations as well as an underlying growth trend. These factors must be taken into account in determining a set of representative traffic demands to be used for provisioning.

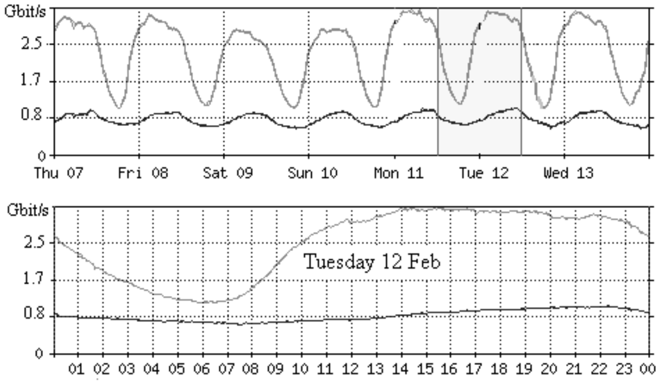

Figure 1: Weekly and daily demand profiles on an OC192 link

In addition to the systematic variations, traffic exhibits stochastic fluctuations about the average. It is these fluctuations which determine the amount of capacity necessary to meet transparency requirements, given an assumed representative demand. The stochastic variations can be modelled as a hierarchical process of sessions, flows and packets. Sessions arrive as a Poisson process [6] and consist in a succession of flows separated by “think times”. The flows relate to an instance of some application and are composed of a stream of packets. It is not feasible or necessary to describe in detail the different types of session or their composition beyond the following broad classification.

Claim 2 *Flows are basically produced by just two types of application: streaming and elastic. This broad classification remains valid as new applications emerge.*

Streaming applications require the real time transmission of audio or video signals. They produce flows having an intrinsic and generally variable data rate which must be preserved by the network. Transparency requirements are expressed in terms of packet loss and delay.

Elastic applications require the transfer over the network of some form of digital document (e.g., an e-mail, a Web page or an MP3 track). The applications are called elastic since the data rate can be varied (e.g., under the control of TCP) without detriment to transparency measured in terms of the overall transfer time.

Note that elastic flows have an intrinsic volume but a variable duration, while streaming flows have an intrinsic duration but a variable volume[1]. It is significant that, with respect to this broad classification, changes in network traffic due to evolving usage is reflected simply in different volumes of streaming and elastic traffic, respectively.

---

[1] The volume varies due to loss or to deliberate rate adaptation.

3 Over-provisioning

The current "best-effort" Internet service model has proved to be extremely successful over the last two decades, despite the high growth rate of demand. This success has been built on a few key principles and a simple pragmatic traffic engineering rule: over-provisioning.

Claim 3 *Over-provisioning solves (almost) all problems of QoS.*

Backbone IP networks are notoriously over-dimensioned. Providers guarantee network transparency by constantly maintaining low link utilization, below 50%, say. Packet delay and loss are then negligible, ensuring transparency for both streaming and elastic flows. Accessibility is frequently guaranteed for all traffic by "1 to 1" or "1 to n" link protection.

Claim 4 *Over-provisioning is not a viable solution.*

The cost of over-provisioning is not negligible. Assuming a constant demand growth rate of 100% per year, an increase in link utilization of 50% represents a delay in investment of six months. This is significant in a competitive environment when costs must be kept to a minimum. It may thus be argued that over-provisioning will eventually have to be replaced by "optimal", or "adequate", provisioning. This requires a better understanding of the relationship between demand, capacity and realized transparency.

4 Provisioning for transparency

The predictability of traffic (Claim 1) implies that it is generally possible to adequately provision the network to avoid congestion. The challenge is to determine the mechanisms and engineering rules necessary to guarantee transparency for both streaming and elastic flows (Claim 2) with a maximum utilization of network links.

Claim 5 *Reliable statistical transparency guarantees can be realized given only a broad characterization of demand in streaming and elastic traffic.*

Streaming traffic can be handled efficiently using bufferless statistical multiplexing [7, Chapter 16]. Assume there is a maximum flow peak rate so that any flow may be viewed as a succession of silent and active periods at this rate[2]. Under the assumption of Poisson session arrivals, the number of active flows has a Poisson distribution and it is straightforward to calculate the amount of capacity compatible with any required data loss rate. There is no queuing delay in this ideal model. Results presented in [3] suggest that delay remains negligible in successive queues as long as the flows are shaped to the maximum peak rate at the network ingress. It is necessary just to implement a scheduling mechanism (e.g., priority queuing) to effectively isolate streaming flows from the impact of bursty elastic traffic.

Elastic traffic shares link bandwidth dynamically under the control of TCP. Assuming sessions arrive as a Poisson process and sharing is perfectly fair, it is again straightforward to calculate the amount of capacity necessary to realize a target expected throughput [2]. This provisioning target would typically be slightly less than the maximum achievable throughput given external rate limitations such as modem speed or access line rate. We refer to this maximum achievable throughput as the elastic flow peak rate.

---

[2] e.g., an isolated packet of L bits consitutes an active period of L/R seconds at rate R bit/s.

Claim 6 *Appropriate transparency targets depend on what is economically feasible.*

The targets in question are the loss rate for streaming flows together with their maximum peak rate, and expected throughput for elastic flows. Table 1 shows the maximum utilization that can be achieved for different peak rates and link capacities, assuming for simplicity that the traffic is exclusively streaming or elastic. Transparency requirements are a loss rate of 10^{-6} and a mean throughput equal to 95% of the peak rate, respectively. With a peak to link rate ratio of 1/100, streaming traffic can represent up to 60% of link capacity while the same ratio allows a utilization greater than 90% for elastic traffic.

In an integrated network where the majority of traffic is elastic (as in the present Internet), the utilization limit for elastic traffic prevails. It is necessary, however, to strictly enforce the maximum peak rate of streaming flows to limit their impact on elastic traffic and thus ensure a high utilization of network links. This limitation on streaming flow peak rate is hard. However, any audio and video applications requiring a greater rate can still be handled as elastic flows (with the necessary rate adaptation and non-negligible packet delay and loss rate).

	Link capacity	100 Mbit/s	1 Gbit/s	10 Gbit/s
Streaming traffic	Peak rate = 1 Mbit/s	60%	86%	98%
	10 Mbit/s	15%	60%	86%
Elastic traffic	Peak rate = 1 Mbit/s	93%	99%	99%
	10 Mbit/s	67%	93%	99%

Table 1: Maximum utilization for different values of flow peak rate

Claim 7 *The scope for offering different degrees of transparency is limited.*

The above multiplexing scheme does not distinguish the particular transparency requirements of different applications. In fact, we strongly question the usefulness of such a distinction. It is not practically feasible to control the performance of buffered statistical multiplexing and thus provide differential delay guarantees[3]. Queuing behaviour depends significantly on detailed traffic characteristics (such as the self-similarity behaviour of a variable rate flow) that are impossible to predict or to monitor in real time.

A frequently advocated approach is to interpret a rough description (e.g., a token bucket) as a traffic envelope and to multiplex flows relying on worst case assumptions. The resulting inefficiency is illustrated in the results of Fig. 2 showing the maximal achievable utilization when employing different multiplexing methods. The figure reproduces results from [4, Fig. 9] where the advantage of relaxing the strict worst case assumptions of network calculus is demonstrated (the effective envelope method assumes flows are statistically independent). We have added results corresponding to bufferless multiplexing assuming the actual flow mean rate (and not just the token bucket rate) is known. The mean rate is equal to one third of the token bucket rate (this is compatible with a 10^{-6} probability of non-conformance for the traffic parameters of [4] with exponentially distributed on and off periods and a mean on period of 3 ms[4]). The difference in achievable utilization is particularly significant for a link to peak rate ratio greater than 100.

---

[3] Bufferless multiplexing ensures delays are negligible for all.

[4] The inefficiency of buffered multiplexing would have been greater had we chosen flows with more variable traffic (e.g., with heavy-tailed on or off periods).

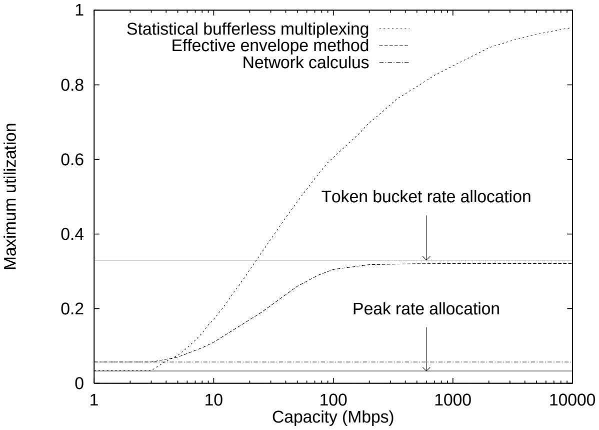

Figure 2: Bufferless vs. buffered multiplexing

Note finally, that an integrated network with a majority of elastic traffic would be even more efficient. There is then hardly any scope for improvement, even if certain applications could tolerate a higher loss rate than the assumed target of 10^{-6}.

5 Controlling accessibility

We have shown that transparency can be ensured relatively simply by means of adequate provisioning. It is not possible for a network to be transparent all of the time, however. Equipment failures, inaccurate forecasts and unpredictable traffic surges can all result in available capacity being insufficient to meet demand. In the absence of any specific control mechanism, the network is not transparent and elastic flow throughput tends to zero as the number of sessions in progress continues to increase. Network resources are then wasted as partially complete applications are aborted [2].

Claim 8 *Resource reservation does not provide an efficient accessibility control.*

In most existing data networks (e.g., ATM or frame relay) and in the proposed Intserv and Diffserv architectures, accessibility is ensured through resource reservation. A key problem is that the traffic descriptor used to allocate resources is less a precise estimation of demand than a rough bound on the data rate the customer is prepared to pay for. Experience shows that users frequently overestimate their actual bandwidth requirement by a large margin. The net result is that nominally reserved bandwidth for a group of traffic contracts sharing a given link is generally many times greater than the actual amount of traffic to be handled. The problem is compounded when it is not possible to precisely locate the paths to be followed by traffic aggregates, as in the so-called “hose model” for virtual private networks [5].

To improve utilization, the network operator has two options: reserve resources in accordance with the traffic contract but allow lower priority best effort traffic to use the residual capacity; overbook the link capacity by admitting several times more aggregates than can strictly be accommodated. The first solution is valid only if there is a sufficient amount of best-effort traffic. It constitutes a form of service differentiation and is discussed under the next claim. The second appears to contradict the very notion of assured accessibility since there is no longer any protection against overload.

Claim 9 *Class of service differentiation provides a crude, inefficient accessibility control.*

Discriminatory mechanisms such as priority queuing or class-based weighted fair queueing are effective in restricting the impact of overload to certain low priority classes. The disadvantage is that all flows in the low priority classes suffer and the resources to which they have access are wasted due to the previously mentioned phenomenon of premature abandons.

Claim 10 *Admission control is the key to controlled QoS.*

The ideal QoS architecture would perform overload control to eliminate excess traffic while maintaining close to 100% goodput: the network would ensure high utilization with all admitted traffic having guaranteed transparency. We suggest that this goal can only be achieved by means of admission control acting at a level corresponding as closely as possible to the user applications: new requests must be denied if their traffic would compromise the transparency of on-going applications. A flow-level QoS architecture realizing admission control "on the fly" without explicit signalling is described and evaluated by Benameur *et al.* [1]. A significant advantage of this architecture is the use of selective admission control thresholds to distinguish different service classes. Some users (e.g., the emergency services) have a higher accessibility guarantee than others. However, any admitted flow has guaranteed transparency whatever its service class.

6 Conclusion

Consideration of the statistical nature of Internet traffic leads us to question the effectiveness of currently envisaged QoS architectures. We propose an alternative approach based on the application of flow-level admission control. The lightweight measurement-based implementation envisaged does not pose insurmountable problems of scalability [1]. It largely remains for equipment manufacturers to meet the challenge of realizing the implied mechanisms.

References

[1] N. Benameur, S. Ben Fredj, F. Delcoigne, S. Oueslati-Boulahia and J.W. Roberts, Integrated admission control for streaming and elastic traffic, Proc. of QofIS 2001, COST 263, 69–81.

[2] S. Ben Fredj, T. Bonald, A. Proutière, G. Régnié and J.W. Roberts, Statistical bandwidth sharing: A study of congestion at flow level, Proc. of ACM SIGCOMM 2001, 111–122.

[3] T. Bonald, A. Proutière, J. Roberts, Statistical performance guarantees for streaming flows using expedited forwarding, Proc. of IEEE INFOCOM 2001, 1104–1112.

[4] R. R. Boorstyn, A. Burchard, J. Liebeherr and C. Oottamakorn, Satistical service assurance for traffic scheduling algorithms, J. Sel. Areas of Comm. 18 (12) 2651–2664, December 2000.

[5] N.G. Duffield, P. Goyal, A. Greenberg, P. Mishra, K.K. Ramakrishnan and J. E. van der Merwe, A flexible model for resource management in virtual private networks, Proc. of ACM SIGCOMM 1999, 95–108.

[6] V. Paxson and S. Floyd, Difficulties in Simulating the Internet, IEEE/ACM Transactions on Networking 9 (4) 392–403, August 2001.

[7] J. Roberts, U. Mocci and J. Virtamo, *Broadband Network Teletraffic*, Final Report of Action COST 242, Spinger, 1996.