

Use of CBR for IP over ATM

J. Aracil, D. Morato, M. Izal and C. Donezar
Dept. de Automática y Computación
Universidad Pública de Navarra, Spain

ABSTRACT

Internet traffic burstiness allows for statistical multiplexing gain in the available bandwidth of an ATM link. However, a dynamic allocation bandwidth assignment (ABR) has to be performed. In this paper we evaluate the real advantages of ABR versus CBR for Internet service provisioning. We consider performance parameters such as connection setup delay and active waiting time due to flow control and show that CBR schemes can be a good alternative for Internet service provisioning over ATM networks.

Keywords: IP over ATM, Internet service provisioning, quality of service.

1. INTRODUCTION

Nowadays, the Internet is turning into a multimedia network that aims at offering real-time services such as video-conference. Along with the growing demand of such multimedia services there is also a great deal of interest in non real-time (or elastic¹) applications such as WWW or email. Usually, such elastic applications run on top of TCP.

Best-effort service is currently offered by Internet Service Providers (ISPs) but there is a considerable interest in adding service discrimination to the current network. New versions of the IP protocol provide capabilities for a better granularity in Quality of Service (QOS). Furthermore, it is foreseen that IP services run on top of ATM networks, which will offer QOS guarantees.

ATM backbones for IP traffic are currently being used. The lack of maturity of the ATM signalling standards complicates matters as far as QOS discrimination is concerned. However, as ATM standards evolve, we will face the need of accomodating IP services into ATM service classes (CBR, VBR, ABR, UBR).

From the service provider standpoint, the ATM backbone provides a dedicated bandwidth which is hired to the operator with flat rate. A number of users are connected to the ISP and they are usually billed with a flat rate as well. The problem that the service provider must solve is how to share the ATM bandwidth so that a better QOS can be provided to the users.

Dynamic bandwidth allocation schemes, such as ABR and UBR, provide a better link utilization since the available bandwidth is statistically shared among traffic sources. However, there is a complexity cost associated to such schemes. Therefore, it is important to evaluate the real QOS improvements that can be achieved.

Our analysis scenario consists of a dedicated bandwidth which is shared by a number of users. We assume that the service provider aims at offering different QOS to the users, which are billed according to the QOS perceived. Two different approaches for bandwidth allocation are considered: constant allocation (CBR) and dynamic allocation (ABR). We will show that ABR does not offer significant advantages over a CBR allocation for elastic services, at a cost of a higher complexity.

The paper is structured as follows: section 2 presents the state of the art in Internet service provisioning and pricing, section 3 describes the user traffic model. Results and discussion are presented in section 4. Finally, section 5 presents the conclusions that can be drawn from this study.

Send correspondence to J. Aracil: Email: javier.aracil@upna.es, Telephone : +34 48 16 97 33, Fax : +34 48 16 92 81

2. STATE OF ART AND CONTRIBUTIONS

A number of proposals have been made regarding Internet service provisioning and pricing. Both issues are strongly related, since QOS discrimination implies billing for the different performance guarantees offered by the network. The key issues are:

- How to describe user traffic in terms of the parameters available in the UNI signalling standards, in order to allocate resources to the different streams.
- What is the utility function for users, namely how to map QOS guarantees (bandwidth and end-to-end delay) into user satisfaction.

David Clark addresses the latter² arguing that user satisfaction grows with transactions throughput. Consider a large file downloading operation, a large instantaneous bit rate is useless unless such bit rate is maintained during the whole transaction. Otherwise, the user is kept waiting until the transaction finishes. Since it is possible to know the file sizes before the transaction takes place bandwidth allocation is fairly simple. The performance measure is total transaction delay and the sustained throughput during the transaction can be estimated easily. In order to map such throughput requirements into UNI parameters an in/out flag is used to mark packets. Throughput is being monitored during the transaction duration and packets are marked accordingly. An out flag means that the packet is more likely to be discarded. We will adopt transaction delay as the QOS metric for Internet service provisioning.

Another proposal for bandwidth allocation is the "smart market" proposal by Mackie-Maison and Varian.³ Packets are marked with a bid, which is the price that the user would pay at a point of congestion. Packets whose bid price is below a threshold are held in queue.

Edell, McKeown and Varaiya⁴ propose a billing gateway system to charge users that share a billed link. Varaiya⁵ distinguishes between fixed or access charge, usage charge and service quality charge.

The key issue in all the abovementioned proposals is to be able to allocate bandwidth dynamically and to charge users according to varying network load conditions. From the service provider standpoint, a resource allocation and pricing algorithm should meet the following conditions:

- *Complexity* should be kept at a minimum. A *complexity cost* is associated to any billing algorithm. A rather precise but complex algorithm is not practical since it implies a processing cost, which may lead to additional HW and SW investments.
- It must be *understandable* by users. Human factors are extremely important for the design of billing algorithms. The number of parameters must be kept at a minimum. Namely, it is preferable to provide users with a simple QOS scale under the generic name of "speed" rather than offering more complicated alternatives.
- For the IP over ATM scenario, service QOS parameters must map into ATM service classes QOS parameters. Simplicity is also a key point in this issue.

This paper adds to the abovementioned studies by analyzing the Internet service provisioning scenario in an IP over ATM environment. We consider feasibility of dynamic versus constant allocation schemes taking into account user QOS metrics. There is a large literature on performance evaluation of ABR and CBR schemes, however, we take a different standpoint since we consider practical service provisioning issues such as complexity cost and QOS provided to end-users measured in transaction delay.

3. INTERNET USER TRAFFIC MODEL

Recently, a number of papers have appeared regarding Internet traffic analysis and characterization.⁶⁻¹² Internet traffic is self-similar since the packet counting process in fixed length time intervals presents the scaling property in its finite dimensional distributions. Namely, no matter the time scale that we consider the process looks the same in a distributional sense.

The high variability of this traffic, which does not smooth out as we aggregate time intervals, relates to the fact that user activity periods are heavy-tailed. User thinking times also follow a heavy-tailed law.¹³ The superposition (statistical multiplexing) of such individual user streams with *on-off* heavy-tailed periods gives a Fractional Brownian Motion (FBM). A FBM is a Gaussian self-similar process which can be used to model accurately the packet birth process coming from a multiplex of Internet sources.⁸ Such process is characterized by an autocorrelation function which follows an hiperbolyc law, in contrast with short-range dependent models which present exponential decay.

Even though the abovementioned papers offer considerable insight into Internet traffic analysis they are not able to portray the Internet user behavior considered in this paper. First, all the abovementioned papers consider either academic or industry networks, in which it is not necessary to set up a connection to the Internet provider site before the data transfer (and billing process) takes place. Secondly, end users are mainly students, faculty or company employees whose network use is supposedly not related to entertainment. Residential users profile is completely different in terms of network use.

In previous studies,¹⁴ we analyze four days worth of traffic in the UCB campus FDDI rings and perform extensive analysis of the trace. We focus on a phenomenological (causal) modeling rather than adopting an abstract approach. Some of the conclusions obtained are still valid in the scenario under analysis.

First, we observe that a significant amount of bandwidth is consumed by traffic-intensive applications such as file transfers (FTP). We add service discrimination to the source level analysis performed in¹³ obtaining the following results:

- Users spawn file transfer sessions (such as FTP control sessions) during which several file transfers occur. This birth process is Poisson.
- File size and duration follow both a heavy-tailed distribution with infinite variance.
- User think times within file transfer sessions are also heavy-tailed.

This results are represented graphically in figure 1.




Figure 1. User model

We argue that it is possible to follow a fluid-flow approach using a Poisson batch arrival process with heavy tail batch sizes to model individual streams. Such approximation represents a worst performance case since we neglect the *off* periods within file tranfer sessions.

Tsybakov and Georganas show in⁶ that Poisson arriving bursts with the Pareto distribution of active periods and constant rate 1 form an asymptotically second-order self-similar process. Therefore, we are considering a self-similar traffic process to model a multiplex of Internet users.

4. RESULTS AND DISCUSSION

4.1. Analysis scenario

Figure 2 shows the scenario under analysis, which consists of an ISP providing Internet access to a large population of users.




Figure 2. Scenario under analysis

Such users are grouped into different queues (VP/VC pairs) with a different QOS assigned. Our goal is to determine what bandwidth in the ATM link should be assigned to each user group and what ATM service class offers a better trade-off between complexity and QOS optimization. We do not consider UBR service since it does not provide any QOS guarantees at all.

CBR is the simplest way to assign bandwidth to each queue. In fact, the most part of current signalling implementations allow only CBR allocation. ABR is a more sophisticated bandwidth allocation procedure but it allows statistical multiplexing gain in the available bandwidth. The use of ABR implies sending resource management cells to control flow and to guarantee fair allocation of the available bandwidth.

Service disciplines that achieve this goal are known, such as WFQ.¹⁵ However, it is not clear how to map their parameters into the ABR service class parameters. Furthermore, the more link bandwidth the more processing capabilities are need in the multiplexor to handle WFQ.

ABR is a good alternative in comparison to fixed allocation schemes *if the QOS gain in terms of transaction delay justifies the complexity cost*. We will evaluate three factors that contribute to total transaction delay:

- Connection setup overhead.
- Active waiting time due to flow control.
- Transmission time.

ABR makes transmission time decrease, since the available bandwidth is shared among users. However, we need to evaluate the abovementioned contributions to total transaction delay in order to determine the real impact of ABR on performance gain. In the next subsections, we evaluate connection setup overhead, active waiting due to flow control and transmission time.

4.2. Connection setup overhead

Before the transaction starts the TCP three-way handshake takes place. Connection setup overhead depends on roundtrip delay and server processing speed. We measure connection setup overhead for WWW services, as one of the most popular services in the Internet.

Figure 3 shows the percentage of transaction delay that is consumed in connection setup for an HTML page size of 10 KB. We perform 30 *GET* commands to a WWW server in the same LAN (roundtrip delay of 0.01 s.) and in the wide area network (roundtrip delay 0.1 s.).

The results show that even in the best case (client and server in the same LAN) a 20% of transaction delay is due to connection setup. Since we are performing 30 consecutive trials per server we are not taking into account the effect of DNS, which would make this percentage increase.




Figure 3. Percentage of transaction delay lost in WWWW connection setup

4.3. Active waiting due to flow control

Figure 4 shows the percentage of transaction delay which is due to active waiting due to TCP flow control, which is evaluated by simulation. Such active waiting time depends on queue occupancy (link utilization) since unacknowledged packets may be held in queue. Analytical calculations are difficult because a feedback effect takes place. Therefore, we perform a simulation analysis, obtaining the results of figure 4.




Figure 4. Percentage of transaction delay lost in active waiting for window sizes of 4 KB, 8 KB and 16 KB

We observe that the effect of flow control can be significant. Window sizes of 8 KB are usual for TCP clients running on PCs. For a fairly loaded network (around 50%) a 40% of total transaction delay is due to TCP flow control. Therefore, the effect of flow control is significant *even if the utilization factor is not too large*.

4.4. Transmission time

Rare events of the form $P(W > B)$ take place in an exponential fashion, as predicted by the Large Deviations Theory.¹⁶ In fact, one could approximate the buffer occupancy beyond a threshold B by:

$$P(W > B) \approx e^{-\delta B} \quad (1)$$

In order to determine the convergence rate δ assume that we have a discrete-time model where X_n customers arrive at time n , ($n > 1$) We aim at estimating $P(W > B)$ after an idle period has occurred. If c customers can be served in each time interval the buffer occupancy reaches B if the arrival rate a exceeds c during $B/(a - c)$ time units. Using the result of Cramer's theorem we have

$$P(W > B) \approx \exp \left\{ -\frac{B}{a^* - c} \Lambda^*(a^*) \right\} \quad (2)$$

where

$$\frac{\Lambda^*(a^*)}{a^* - c} = \min_{a > c} \frac{\Lambda^*(a)}{a - c} \quad (3)$$

and $\Lambda^*(a^*) = \sup_{\theta} [\theta a - \Delta(\theta)]$ being $\Lambda(\theta)$ the log-moment generating function $\Lambda(\theta) = \text{Log}E[e^{\theta x_1}]$.

Assume that the arrival process is a Poisson process of rate $\lambda = N\lambda'$ where λ' is the end user rate and N the total number of users. Batches are distributed according to a Pareto distribution

$$f_X(x) = \alpha k^\alpha x^{-\alpha-1} \quad (4)$$

where k is the minimum batch size and α is the exponent parameter of the Pareto distribution. We assume a value of α in the range $-1 \leq \alpha \leq 2$, which gives rise to a self-similar process. This implies finiteness in the first moment and infinite variance. However, since the batch size cannot take an infinite value in our scenario (file sizes are always limited), we approximate $f_X(x)$ by a truncated Pareto distribution in the range $k \leq x \leq x_{max}$, where x_{max} represents the maximum file size (50 KB).

Thus, finiteness in all moments can be assumed and, therefore, the log-moment generating function exists. Assume that the time interval duration is Δt and that Δt is small in comparison with $1/\lambda$. With such assumptions the probability of exactly one arrival is $\lambda\Delta t$ and the probability of more than one batch arrival is $o(\Delta t)$, therefore

$$\Lambda(\theta) = \text{Log}E[e^{\theta X}] = \text{Log} \int_k^{x_{max}} e^{\theta x} \lambda f_X(x) dx = \text{Log} \int_k^{x_{max}} \lambda \alpha k^\alpha x^{-\alpha-1} e^{\theta x} dx \quad (5)$$

The *effective bandwidth* $c(\delta)$ or minimum value of c for which 1 holds is given by⁵

$$c_{min} = \frac{\Delta(\delta)}{\delta} \quad (6)$$

Figure 5 shows $P(W > B)$ as a function of the capacity $c(\delta)$. λ equals 0.1 and the mean file size is 10 KB, with $\alpha = 1.2$. Similar curves are obtained with other parameters.

Figure 5. Delay versus equivalent bandwidth

We observe two different performance regions:

- If the assigned capacity is high then a bandwidth increase does not have a significant influence in the QOS perceived.

- If the assigned capacity is small then a bandwidth increase translates into a significant improvement in the perceived QOS

This results match out intuition. Since transaction sizes are Pareto distributed service time variance is very high. If utilization is high a small increase in bandwidth assignment translates into a significant improvement.

However, we observe from the previous subsections that the effect of TCP flow control and connection setup overhead is rather significant. At a utilization of around 50% a 20% of total transaction delay is lost in connection setup and a 50% is lost in active waiting due to flow control. ABR proves more useful as utilization increases but we observe that the active waiting time increases accordingly, therefore the performance increase obtained by ABR is not as important as figure 5 suggests.

As far as billing and pricing QOS is concerned, note that ABR would allow some users in the same QOS group to obtain a higher QOS due to the available bandwidth. Such available bandwidth is fairly shared by all users in a link, so a good approach would be billing according user total traffic. However, we face with the complexity of keeping track of a large number of users transmitting over the link. On the other hand, it seems unpractical for users since being charged for the file sizes is awkward and difficult to measure. A CBR approach is much simpler to implement. Besides, users have a better understanding of what they are paying for. A simple CAC algorithm would let users select between different QOS scales: slow, fair and very fast. The service provider then partitions the ATM link capacity in several constant bit rate VP/VCS according to the different QOS groups. Time-of-day billing schemes may be used, which will contribute to traffic smoothing.⁴

5. CONCLUSIONS

According to the theoretical analysis and simulation and taking into account connection setup overhead and active waiting time we can state that CBR is a good alternative for Internet service provisioning. Furthermore, the complexity associated with this solution is much smaller than that of ABR services. Billing users is also made simpler and much more understandable to them. Therefore, in order to evaluate ATM service classes for elastic Internet services, CBR allocation should also be taken into account.

REFERENCES

1. T. Nishida and K. Taniguchi, "QOS control and service models in the Internet," *IEICE Transactions in Communications* **E78-B**, pp. 1162–1175, April 1995.
2. D. D. Clark, "Adding service discrimination to the internet," tech. rep., MIT LCS, September 1992.
3. J. K. Mackie-Mason and H. R. Varian, "Pricing congestible network resources," *IEEE Journal on Selected Areas in Communications* **13**(7), pp. 1141–1149, 1995.
4. R. Edell, N. McKeown, and P. P. Varaiya, "Billing users and pricing for TCP," *IEEE Journal on Selected Areas in Communications* **13**(7), pp. 1162–1175, 1995.
5. J. Walrand and P. Varaiya, *High-Performance Communication Networks*, Morgan Kaufmann, San Francisco, CA, 1996.
6. B. Tsybakov and N. D. Georganas, "On self-similar traffic in ATM queues: Definitions, overflow probability bound and cell delay distribution," *IEEE/ACM Transactions on Networking* **5**, pp. 397–409, June 1997.
7. W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, "On the self-similar nature of Ethernet traffic," *IEEE/ACM Transactions on Networking* **2**, pp. 1–15, January 1994.
8. I. Norros, "On the use of Fractional Brownian Motion in the theory of connectionless networks," *IEEE Journal on Selected Areas in Communications* **13**, pp. 953–962, August 1995.
9. V. Paxson and S. Floyd, "Wide area traffic: The failure of Poisson modeling," *IEEE/ACM Transactions on Networking* **4**, pp. 226–244, April 1996.
10. M. Parulekar and A. Makowski, "Tail probabilities of a multiplexer with self-similar traffic," in *IEEE INFOCOM '96*, vol. 3, pp. 1452–1459, 1996.
11. B. K. Ryu and S. B. Lowen, "Point process approaches to the modeling and analysis of self-similar traffic - part I: Model construction," in *IEEE INFOCOM '96*, vol. 3, pp. 1468–1475, March 1996.

12. W. Willinger, M. S. Taqqu, W. E. Leland, and D. V. Wilson, "Self-similarity in high speed packet traffic: Analysis and modeling of ethernet traffic measurements," *Statistical Science* **10**(1), pp. 67–85, 1995.
13. W. Willinger, M. S. Tacqu, R. Sherman, and D. V. Wilson, "Self-similarity through high-variability: Statistical analysis of Ethernet LAN traffic at the source level," in *ACM SIGCOMM 95*, pp. 100–113, (Cambridge, MA), 1995.
14. J. Aracil, R. Edell, and P. Varaiya, "A phenomenological approach to Internet traffic self-similarity," in *35th Annual Allerton Conference on Communications, Control and Computing*, (Urbana-Champaign, Illinois), October 1997.
15. A. Demers, S. Keshav, and S. Shenker, "Analysis and simulation of a fair queueing algorithm," *Journal of Internetworking: Research and Experience* , pp. 3–26, 1989.
16. A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*, Jones and Bartlett, Boston, MA, 1993.