

Traffic Theory and the Internet

Jim W. Roberts, France Telecom R&D

ABSTRACT

We argue that traffic theory, an essential component in the design of traditional telecommunications networks, should be increasingly applied in the development of the multiservice Internet. We discuss the statistical characteristics of Internet traffic at different time scales. Modeling is facilitated on identifying the notion of flow and distinguishing the categories of streaming and elastic traffic. We review mathematical modeling approaches useful for predicting the relationship between demand, capacity and performance for both streaming and elastic flows. Derived results indicate the limitations of service differentiation as a means for guaranteeing QoS and highlight the importance of traditional traffic engineering approaches in ensuring that the network has sufficient capacity to handle offered demand.

INTRODUCTION

In this article we argue that traffic theory should be increasingly used to guide the design of the future multiservice Internet. By traffic theory we mean the application of mathematical modeling to explain the *traffic-performance relation* linking network capacity, traffic demand and realized performance. Since demand is statistical in nature, performance must be expressed in terms of probabilities and the appropriate modeling tools derive from the theory of stochastic processes.

Traffic theory is fundamental to the design of the telephone network. The traffic-performance relation here is typified by the Erlang loss formula which gives the probability of call blocking, E , when a certain volume of traffic, a , is offered to a given number of circuits, n :

$$E(a, n) = \frac{a^n / n!}{\sum_{i=0}^n a^i / i!}.$$

The formula relies essentially only on the reasonable assumption that telephone calls arrive as a stationary Poisson process. It demonstrates the remarkable fact that, given this assumption, performance essentially depends only on a simple measure of the offered traffic, a , equal to the

product of the call arrival rate and the average call duration. Blocking probabilities are *insensitive* to additional details about the nature of traffic such as the distribution of call durations.

In this article we suggest that it is possible to derive similar traffic-performance relations for the Internet, even if these cannot always be expressed as concisely as the Erlang formula. Deriving such relations allows us to understand what kinds of performance guarantees are feasible and what kinds of traffic control are necessary. It is also of considerable importance to identify which traffic characteristics are essential and which can be ignored by identifying insensitivities in performance measures. The objective of traffic theory is ultimately to define simple network engineering procedures like applying the Erlang formula in the telephone network. Derivation of these procedures and proof of their general validity may, however, require somewhat sophisticated mathematical modeling.

Traffic theory currently plays a very minor role in the design of the Internet. Network provisioning is generally based on simple rules of thumb while considerable effort is spent on the design of a variety quality of service (QoS) mechanisms. The role of the latter is to allow an unspecified number of privileged users to escape the effects of congestion by giving them priority treatment. It still remains necessary to apply the appropriate traffic-performance relation if the objective is to ensure that QoS meets specific design targets for a given population of such users. An assurance that a user of premium class service would have had worse quality if it had instead used a lower priority service class is of limited value when the actual quality level can fluctuate widely depending on the network path used and its current traffic level.

It has been suggested that Internet traffic is far too complicated to be modeled using the techniques developed for the telephone network or for computer systems [1]. While we must agree that the modeling tools cannot ignore real traffic characteristics and that new traffic theory does need to be developed, we seek to show in this article that traditional techniques and classical results do have their application and can shed light on the impact of possible networking evolutions.

In the next section we discuss the nature of Internet traffic distinguishing the essential categories of elastic and streaming flows. We then introduce elements of traffic theory appropriate for these two categories before discussing the respective roles of QoS mechanisms and (over)provisioning.

STATISTICAL CHARACTERISATION OF TRAFFIC

Traffic in the Internet results from the uncoordinated actions of a very large population of users and must be described in statistical terms. It is important to be able describe this traffic succinctly in a manner which is useful for network engineering.

TRAFFIC VARIATIONS AND THE NOTION OF STATIONARITY

Observations of traffic on network links typically reveal intensity levels (in bits/sec) averaged over periods of 5 to 10 minutes which are relatively predictable from day to day (Fig. 1). Systematic intensity variations occur within the day reflecting user activity. It is possible to detect a busy period (usually in the afternoon between 2 and 5 pm) during which the traffic intensity is roughly constant. This constancy suggests that Internet traffic, like telephone traffic, can be modeled as a stationary stochastic process where statistical variations occur about an underlying constant intensity. Busy period performance is then estimated by the long term average behavior derived for the stationary process.

The precise characteristics of this stationary process depend on the composition of Internet traffic. Currently, some 90 to 95% of Internet packets use TCP and correspond to the transfer of digital documents of one form or another (Web pages, data files, MP3 tracks, ...). The congestion avoidance algorithms of TCP cause throughput to vary elastically in reaction to random changes in the set of transfers in progress. A small but growing proportion of traffic relates to inelastic streaming audio and video transmission for both interactive and playback applications.

TRAFFIC OBJECTS

The traffic process can be described in terms of the characteristics of a number of objects, including packets, bursts, flows, sessions and connections, depending on the time scale of relevant statistical variations. The preferred choice for modeling purposes depends on the object to which traffic controls are applied. Conversely, in designing traffic controls it is necessary to bear in mind the facility of characterizing the implied traffic object. This consideration is particularly important in the design of the future Internet where only the datagram and the broad destination-based aggregate used in routing are currently recognized. Traffic characterization proves most convenient at an intermediate *flow* level.

A flow is defined for present purposes as the unidirectional succession of packets relat-

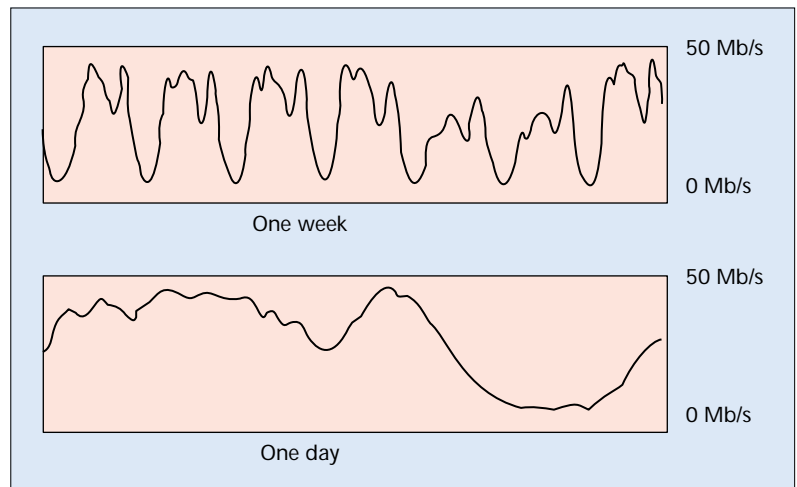


Figure 1. Traffic on an OC3 backbone link.

ing to one instance of an application (sometimes referred to as a microflow). For practical purposes, the packets belonging to a given flow have the same identifier (e.g., source and destination addresses and port numbers) and occur with a maximum separation of a few seconds. It is useful to distinguish elastic flows, where the packets in question constitute a document being transferred, and streaming flows, where the packets represent an audio or video signal. Packet level characteristics of elastic flows are mainly induced by the transport protocol and its interactions with the network. Streaming flows, on the other hand, have intrinsic (generally variable) rate characteristics that must be preserved as the flow traverses the network.

Flows are frequently emitted successively and in parallel in what are loosely termed “sessions.” A session corresponds to a continuous period of activity during which a user generates a set of elastic or streaming flows. For dial-up customers, the session can be defined to correspond to the modem connection time but, in general, a session is not materialized by any specific network control functions.

Some network service models define the notion of “connection” and control resource allocation by means of explicit signalling exchanges. The connection might be set up for a particular flow or used over a long period for an aggregation of flows between given network end points. A significant difficulty resides in defining parsimonious traffic descriptors representing the impact the connection is likely to have on network performance. Signalling overhead may also prove excessive, particularly when each connection relates to an individual flow.

ARRIVAL PROCESSES AND SERVICE REQUIREMENTS

It is well known that the arrival process of IP packets can exhibit extreme rate variations at multiple time scales (“spikes ride on ripples that ride on still longer term swells...” [2]). First reports of this behavior more than ten years ago have given rise to a large amount of

An overload control more effective than relying on user impatience would be to implement some form of admission control for elastic flows: a new flow would be rejected whenever the bandwidth it would receive falls below a certain threshold.

research aiming to explain the so-called self-similarity phenomenon and to predict its impact on network performance (see [3] for a comprehensive treatment of this phenomenon). The main reason for rate fluctuations at time scales greater than a few hundred milliseconds turns out to be extreme variability in the size of the flows making up the observed packet process. Yet more extreme variability (so-called multi-fractal behavior) occurs at smaller time scales due to the burstiness induced by TCP. It proves much simpler to describe traffic in terms of flows.

The arrival process of flows in a backbone link typically results from the superposition of a large number of independent sessions. Observations confirm the predictable property that session arrivals can be assimilated to a Poisson process. This means simply that the probability of a new arrival in a short interval of length dt is equal to λdt , where λ is the arrival intensity, and is independent of all past activity. A Poisson process results naturally when traffic is due to the independent activity of a very large population of users, each individually having a very small intensity.

As a first approximation, it is not unreasonable to assume that individual flows also occur as a Poisson process. To ignore the correlation between flow arrivals within the same session is not necessarily significant when the number of sessions is large. It is also true that results derived under the simple Poisson assumption are also often true under more general assumptions.

The size of elastic flows (i.e., the size of the documents transferred) is extremely variable and has a so-called heavy-tailed distribution: most documents are small (a few kilobytes) but the number which are very long tend to contribute the majority of traffic. The precise nature of the size distribution is important in certain circumstances, such as describing the resulting self-similar packet arrival process, and can have a significant impact on performance in some multiplexing schemes. However, it proves very difficult to describe the distribution reliably in a suitably parsimonious fashion. It is therefore highly desirable to implement controls such that performance is largely insensitive to the precise document size characteristics.

The duration of streaming flows also generally has a heavy-tailed distribution. Furthermore, the packet arrival process within a variable rate streaming flow is often self-similar [3]. As for elastic flows, it proves very difficult to precisely measure and specify these characteristics. It is thus again important to design traffic controls which make performance largely insensitive to them.

TRAFFIC DESCRIPTORS

When traffic control is performed with respect to connections, it is necessary to describe the connection in terms allowing the network to deduce its likely impact on performance. It was recognized early in ATM standards that traffic descriptors should satisfy three requirements. They should be:

- Useful for resource allocation
- Understandable by the user
- Verifiable at the network ingress

Experience has shown that it is practically impossible to reconcile these requirements. We note, in particular, that the widely used token bucket satisfies the last one (verifiability) but is hardly useful for resource allocation and is not a satisfactory descriptor for any traffic stream with random rate variations.

TRAFFIC THEORY FOR ELASTIC TRAFFIC

Exploiting the tolerance of document transfers to rate variations implies the use of closed-loop control. In this section we assume closed-loop control is applied end-to-end on a flow-by-flow basis using TCP.

PACKET SCALE PERFORMANCE

TCP realizes closed loop control by implementing an additive increase, multiplicative decrease congestion avoidance algorithm: the rate increases linearly in the absence of packet loss but is halved whenever loss occurs. This behavior causes each flow to adjust its average sending rate to a value depending on the capacity and the current set of competing flows on the links of its path. Available bandwidth is shared in roughly fair proportions between all flows in progress.

A simple model of TCP results in the following well-known relationship between flow throughput B and packet loss rate p :

$$B(p) = \frac{\text{constant}}{RTT\sqrt{p}},$$

where RTT is the flow round trip time (see [4] for a more accurate formula). The formula can also be interpreted as relating p to the realized throughput B . Since B actually depends on the set of flows in progress (each receiving a certain share of available bandwidth), we deduce that packet scale performance is mainly determined by flow level traffic dynamics. It can, in particular, deteriorate rapidly as the number of flows sharing a link increases.

PERFORMANCE AT FLOW SCALE

Consider an isolated bottleneck link and assume flows arrive according to a Poisson process. Assume further that all flows using the link receive an equal share of bandwidth. The number of flows in progress is then a random variable which behaves like the number of customers in a so-called processor sharing queue [5]. An interesting feature of this system is that, for any distribution of document size, average flow throughput is simply equal to the difference between link capacity and expected demand (measured by the product of arrival rate and mean document size).

While this model is too simple to predict performance in real networks, it usefully illustrates two interesting points which turn out to be more generally true. First, performance depends primarily on expected traffic demand (in bits/second) and only marginally on parameters describing the distribution of document sizes. Second, performance tends to be excellent as long as expected demand is less than available

capacity. Note that the second observation suggests there is limited scope for service differentiation when the network is adequately provisioned to handle the traffic of all service classes.

In overload, when expected demand exceeds link capacity, the processor sharing queue is unstable: the number of flows in progress increases indefinitely as flows take longer and longer to complete while new flows continue to arrive. This is a plausible explanation for the severe congestion events sometimes experienced in the Internet. In practice, instability is controlled by users abandoning transfers, interrupting sessions or simply choosing not to use the network at all in busy periods. However, the network is inefficiently used due to partially completed flows and sessions, and even patient users experience very poor response times.

ADMISSION CONTROL FOR ELASTIC TRAFFIC

An overload control more effective than relying on user impatience would be to implement some form of admission control for elastic flows: a new flow would be rejected whenever the bandwidth it would receive falls below a certain threshold. We have in mind a measurement-based scheme capable of roughly estimating the bandwidth available to a new flow. The admittance threshold would be chosen small enough to avoid flow rejection in normal load but large enough to ensure satisfactory throughput for admitted flows in overload. The need to perform admission control at very high speed precludes reliance on explicit signalling and resource reservation. We envisage an *implicit* admission control realized simply by discarding the first packets of flows which are not admitted [6].

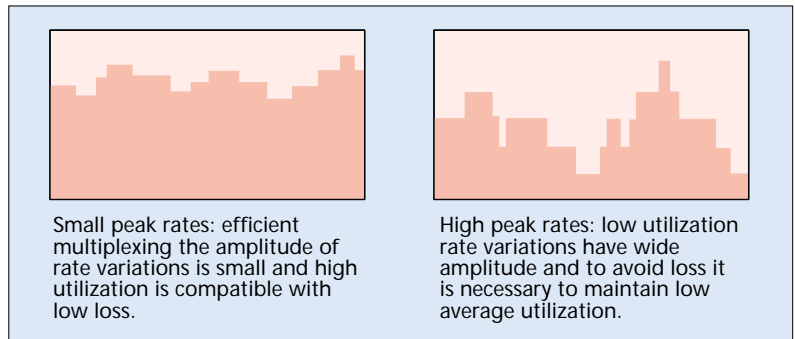
TRAFFIC THEORY FOR STREAMING TRAFFIC

We assume here that streaming traffic is subject to open-loop control: an arriving flow is assumed to have certain traffic characteristics; the network performs admission control, only accepting the flow if quality of service can be maintained; admitted flows are policed to ensure their traffic characteristics are indeed as assumed.

FLUID FLOW MODELS

The effectiveness of open-loop control depends on how accurately performance can be predicted given the characteristics of audio and video flows. To discuss multiplexing options we first make the simplifying assumption that flows have unambiguously defined rates like fluids. It is useful then to distinguish two forms of statistical multiplexing: bufferless multiplexing and buffered multiplexing.

In the fluid model, statistical multiplexing is possible without buffering if the combined input rate is maintained below link capacity. As all excess traffic is lost, the overall loss rate is simply the ratio of expected excess traffic to expected offered traffic (formally, loss rate = $E[(\Lambda_t - c)^+] / E\Lambda_t$ where Λ_t is the input rate process and c is the link capacity). It is important to



■ Figure 2. Efficiency of bufferless multiplexing.

notice that this loss rate only depends on the stationary distribution of the combined input rate Λ_t but not on its time dependent properties, including self-similarity.

The level of link utilization compatible with a given loss rate can be increased by providing a buffer to absorb some of the input rate excess. However, the loss rate realized with a given buffer size and link capacity then depends in a complicated way on the nature of the offered traffic. In particular, loss and delay performance turn out to be very difficult to predict when the input process is self-similar (see [7] p. 91).

Bufferless multiplexing thus has clear advantages with respect to the facility with which quality of service can be controlled. It is also efficient when the peak rate of an individual flow is small compared to the link rate because high utilization is compatible with negligible loss: a large number of flows can be multiplexed together and their combined rate variation is of relatively low amplitude (Fig. 2). Service differentiation with respect to loss tolerance is then of limited utility since multiplexing efficiency is high even for low loss rates.


PACKET SCALE PERFORMANCE

Packet queuing occurs even with so-called bufferless multiplexing due to the coincidence of arrivals from independent inputs. While we assumed above that rates were well defined, it is necessary in practice to account for the fact that packets in any flow are not perfectly spaced and packets of different flows arrive asynchronously. To correctly size buffers and to predict end-to-end delays, it is necessary to account for the impact of jitter which alters the instantaneous flow rate. Our ongoing research suggests jitter can be controlled quite simply as long as flows are correctly spaced at the network ingress. This is in extension of the notion of negligible jitter first developed in the context of ATM (see [7], pp 120-122).

INTEGRATING STREAMING AND ELASTIC TRAFFIC

Though we have discussed traffic theory for elastic and streaming traffic separately, integration of both types of flow on the same links has considerable advantages. By giving priority to streaming flows, they effectively see a link with very low utilization yielding extremely low packet loss and delay. Elastic flows naturally benefit from the bandwidth which would be unused if dedicated bandwidth were reserved for stream-

Bandwidth reservation provides protection to a flow or an aggregate of flows by isolating its traffic from that of other users. A practical difficulty arises, however, in matching the amount of reserved capacity to the actual rate requirement of a connection which generally varies in time.



■ **Figure 3.** *Integrated streaming and elastic traffic.*

ing traffic and thus gain greater throughput (Fig. 3). Performance of an integrated system in traffic overload requires special attention since the priority afforded to streaming flows could cause unacceptable degradation to elastic flow throughput. One possibility we have investigated is to use admission control for both elastic and streaming flows [6].

ADMISSION CONTROL

It is generally accepted that admission control must be employed for streaming flows to guarantee their low packet loss and delay requirements. Among the large number of schemes which have been proposed in the literature, our preference is clearly for a form of measurement-based control where the only traffic descriptor is the flow peak rate and the available rate is estimated in real time ([7], pp. 137–142). On high bandwidth links the same implicit admission control procedure discussed above for elastic traffic could be applied to both streaming and elastic flows with either type of flow being rejected if the estimated available bandwidth were less than a threshold. Note that the inherent tolerance of elastic flows and the fact that streaming flows are protected by priority scheduling means that the available bandwidth estimate does not need to be very precise.

QoS OR OVERPROVISIONING?

Overprovisioning is frequently opposed to the implementation of specific mechanisms as a means of ensuring satisfactory quality of service. In this section we discuss these two options developing the point of view that adequate provisioning is the key to quality of service but that mechanisms must be in place to ensure efficiency and resilience in case of overload.

QoS MECHANISMS

Currently envisaged network models provide quality guarantees using explicit resource reservation or service differentiation or a combination of both. The previous discussion reveals a number of difficulties associated with both approaches.

Bandwidth reservation provides protection to a flow or an aggregate of flows by isolating its traffic from that of other users. A practical difficulty arises, however, in matching the amount of reserved capacity to the actual rate requirement of a connection which generally varies in time. In practice, in Frame Relay and ATM networks, declared traffic descriptors are usually taken as only a rough guide. Most providers systematical-

ly practice overbooking, “allocating” available bandwidth several times over. Although this clearly violates the notion of service protection, perceived quality of service is satisfactory in most cases. It is tempting to deduce from this that reservation is not really necessary as long as admission control is employed to protect the QoS of existing traffic.

Service differentiation consists in handling different classes of traffic using particular per-hop behaviors [8]. We recognize the advantage of this when such classes have qualitatively different QoS requirements, as with streaming and elastic traffic. It is much less obvious that it is useful to differentiate traffic classes of the same type according to the degree of quality (e.g., lower packet loss or higher throughput). This is primarily because it is practically impossible to consistently achieve QoS targets which are intermediate between very good and very bad. Service differentiation is mainly effective in overload and can then be viewed as a means of protecting higher priority flows from the effects of congestion. Unfortunately, the resources allocated to the low priority classes are then inefficiently used due notably to flows being aborted.

The above discussion, based on traffic theoretic considerations, leads us to seek alternative mechanisms to ensure quality of service. A highly desirable component would be admission control at flow level. This would ensure efficient use of network resources. It is also compatible with a simple tariff structure. Service differentiation could be realized by applying different admission criteria to different traffic classes.

INFERRING THE TRAFFIC MATRIX

Adequate provisioning relies on accurate demand forecasts based on a sound knowledge of existing traffic. The traffic models discussed in the previous sections suggest that performance depends above all on expected demand in bits/sec (the product of the flow arrival rate and the average flow size). The essential data for network engineering are then the elements of the traffic matrix giving the expected demand between all boundary routers in a considered network domain, say.

It proves very difficult in practice to infer this matrix from routine traffic measurements in large IP networks where it is necessary to couple measured bit rates on network links with detailed information on the routes calculated by intra- and inter-domain routing protocols [9]. It is also necessary to account for the exceptional growth rate of Internet traffic and significant modifications in structure and locality brought by new applications and new ways of dealing with content distribution. This uncertain knowledge of the traffic matrix leads us to further suggest that the Internet employ some form of traffic-oriented adaptive routing.

TRAFFIC ROUTING

If the traffic matrix were known accurately, it would be possible to provide sufficient bandwidth on the routes of all origin-destination traffic relations to handle their demand with excellent QoS. Similarly, for a network of given

capacity it would be possible to specify a set of routes ensuring the best possible quality of service. These operations are possible to some extent with current IP routing protocols and will be considerably facilitated with the introduction of MPLS. However, static routing has obvious limitations when knowledge of the traffic matrix is uncertain and cannot easily deal with unexpected overloads or outages.

Static routing in the telephone network was replaced by adaptive routing many years ago [10]. While the original arguments which justified the introduction of adaptive routing were economic (the network required less capacity for the same quality of service), its main advantages have since been recognized to be increased network resilience. Adaptive routing can be performed on the basis of different traffic objects. The solution we prefer is to apply adaptive routing at flow level. This ties in with our advocacy of flow admission control, the decision to reject a new flow being just one possible (extreme) routing decision.

PRICING AND PROVISIONING

Much research has been performed lately on the use of pricing as a QoS mechanism (see [11], for example). Charges increase with the level of demand ensuring that only those users willing to pay the most actually use the network. These schemes tend to neglect the main function of charging which is to reimburse the costs incurred by the provider in installing and operating the network infrastructure. The use of admission control as the primary QoS mechanism, as outlined above, is intended to allow simple cost-based charging. Users would be charged in relation to the volume of traffic emitted. There is no obvious reason why one would apply different charges to streaming and elastic traffic. The only criterion to be satisfied is that the combination of flat rate and per-byte charges cover the provider's costs, providing the necessary incentive to expand the network in response to traffic growth.

CONCLUSIONS

IP traffic resulting from the activity of a large population of users can be represented for performance prediction purposes as a stationary stochastic process. This process can be modeled most conveniently at flow level distinguishing the two main categories of elastic and streaming traffic.

Our discussion of traffic theory for elastic traffic shows that performance is mainly determined by the way bandwidth is shared between contending flows. A simple processor sharing model indicates that average throughput performance is largely insensitive to detailed traffic characteristics such as the flow size distribution. Traffic theory for open-loop controlled streaming traffic is simplest when the network performs bufferless multiplexing. The packet loss probability is then insensitive to any self-similarity in the rate variations of individual flows. Integration of streaming and elastic traffic with priority service to streaming packets is advantageous to both kinds of traffic.

Taking account of the statistical nature of traffic leads us to conclude that there is limited potential for meaningful service differentiation except in overload conditions. When overload occurs, we suggest it is preferable to perform admission control, rejecting some newly arriving flows, than to allow congestion to develop even if packet loss and throughput degradation are confined to the lowest priority class. Admission control would be applied equally to elastic and streaming traffic by discarding the packets of new flows whenever estimated available bandwidth falls below a certain threshold.

The main means to ensure high quality of service is adequate provisioning which basically means avoiding overload by ensuring that the capacity of all links is greater than demand. Experience with the telephone network suggests this objective can be realized most easily if the network employs adaptive routing. We suggest the development of adaptive routing schemes at flow level using the same available bandwidth estimation necessary for admission control.

The above conclusions are somewhat at odds with current thinking about the likely evolution of the Internet and it would certainly require more space than is available here to convince most people that they are valid. We do hope, however, that we have been able to adequately demonstrate that the traffic theory dimension, accounting for the statistical nature of traffic and understanding the traffic-performance relation, is vitally important in building a network capable of meeting quality of service requirements.

REFERENCES

- [1] W. Willinger, and V. Paxson, "Where Mathematics meets the Internet," *Notices of the American Mathematical Society*, vol. 45, no. 8, Aug. 1998, pp. 961-70.
- [2] W. Leland *et al.*, "On the Self-Similar Nature of Ethernet Traffic," *IEEE Trans. Net.*, vol 2, no. 1, Feb 1994, pp. 1-15.
- [3] K. Park and W. Willinger, *Self-Similar Network Traffic and Performance Evaluation*, Wiley, 2000.
- [4] J. Padhye *et al.*, "Modeling TCP Throughput: A Simple Model and Its Empirical Validation," *Proc. SIGCOMM'98, ACM*, 1998.
- [5] L. Kleinrock, *Queueing Theory - Vol. 2*, Wiley, 1976.
- [6] J. Roberts and S. Oueslati-Boulaiah, "Quality of Service by Flow-Aware Networking," *Phil Trans. R. Soc. Lond. A*, 2000, pp. 2197-207.
- [7] J. Roberts, U. Mocci, and J. Virtamo (Eds), "Broadband Network Teletraffic," Springer Verlag, LNCS, vol. 1155, 1996.
- [8] S. Blake *et al.*, "An Architecture for Differentiated Services," RFC 2475, IETF, Dec. 1998.
- [9] A. Feldmann *et al.*, "Deriving Traffic Demands for Operational IP Networks: Methodology and Experience," *Proc. SIGCOMM 2000, ACM*, 2000.
- [10] G. Ash, *Dynamic Routing in Telecommunications Networks*, McGraw-Hill, 1998.
- [11] F. P. Kelly, "Models for A Self-Managed Internet," *Phil Trans. R. Soc. Lond. A*, 2000, pp. 2335-48.

BIOGRAPHIES

JIM ROBERTS (james.roberts@francetelecom.fr) has a B.Sc. in mathematics from the University of Surrey, UK and a Ph.D. in computer science from the University of Paris. His career has been in the area of traffic engineering research, participating in the evolution of telecommunications from POTS to the Internet, via narrowband and broadband ISDNs. He currently lead a team at France Telecom R&D working on multiservice network traffic modeling and performance evaluation.

When overload occurs, we suggest it is preferable to perform admission control, rejecting some newly arriving flows, than to allow congestion to develop even if packet loss and throughput degradation are confined to the lowest priority class.