Models for a self-managed Internet

By Frank Kelly

Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road, Cambridge CB3 0WB, U.K.

This paper uses a variety of mathematical models to explore some of the consequences of rapidly growing communications capacity for the evolution of the Internet. It argues that queueing delays may become small in comparison with propagation delays, and that differentiation between traffic classes within the network may become redundant. Instead, a simple packet network may be able to support an arbitrarily differentiated and constantly evolving set of services, by conveying information on incipient congestion to intelligent end-nodes which themselves determine what should be their demands on the packet network.

Keywords: stability, shadow price, fairness, explicit congestion notification, queues, short transfers.
This is a draft of a paper later published in the
Philosophical Transactions of the Royal Society A358 (2000) 2335-2348.
Pointers to related information are available at:
http://www.statslab.cam.ac.uk/~frank/smi.html .

1. Introduction

When a communication network becomes congested, a tradeoff must be made between traffic that is carried and traffic that is not. In telephone networks the traditional mechanism has been a call admission control, which blocks a newly arriving call if any of the resources that would be needed by the call are congested. The implementation of this mechanism has generally required a parallel signalling network, to monitor resource loads and make admission decisions. In contrast, congestion in the current Internet causes traffic through a congested resource to suffer packet loss, and this in turn causes at least some end-systems to reduce their load on that resource.

In a traditional telephone network, the load of a call is well-defined, and users tolerate occasional blocking in return for a guaranteed bandwidth for accepted calls. The statistical properties of packet streams carrying voice, together with the limit on load achieved by the call admission control, allow queues within the network to be kept short and quality of service for accepted calls to be assured.

In the Internet, the load imposed by end-systems is much less predictable, and the bandwidth achieved by a user fluctuates as a consequence of the behaviour of other users. Buffers, important since the early days of store-and-forward communication networks (Kleinrock 1964), are used to smooth statistical fluctuations in demand for scarce transmission capacity.

Might it be possible to design a network so that it can carry existing loads on telephone networks and the Internet, as well as new forms of traffic? A problem is that some current Internet traffic needs large buffers, causing unacceptable queueing delays for real-time applications such as telephony, a problem compounded by the current use of packet loss and delay as the default mechanism for deterring demand. Attempts to solve the problem are generally based on a small number of service classes with well defined quality and prices, delivered by a network using queueing mechanisms that treat differently packets belonging to different classes. By giving higher priority to some packets, and using large buffers for others, it may be possible to carry real-time applications over the same transmission capacity as other less delay-sensitive traffic. But this approach has drawbacks. In particular, the development of ATM traffic classes (ITU 1996) has illustrated some of the difficulties of defining service categories and incentive-compatible pricing schemes (Songhurst 1999) before the applications that might use the categories have been invented or have become widespread.

A body of work is now emerging that takes a radically different approach to differentiated services (see, for example, Crowcroft & Oechslin 1998; Gibbens & Kelly 1999a; Key & McAuley 1999a). Its premise is that a simple packet network may be able to support an arbitrarily differentiated set of services by conveying information on congestion from the network to intelligent end-nodes, which themselves determine what should be their demands on the packet network. There would then be no need for large buffers or priority queues within the network, or for connection acceptance control at the border of the network. This paper is intended to provide a brief overview of some of the mathematical models that have been found useful in the exploration of this approach.

The organization of the paper is as follows. In Section 2 we explore the impact of rapidly growing communications capacity upon queueing delays. Under various scaling regimes we argue that queueing delays become small in comparison with propagation delays. This point is hardly controversial, and yet its consequences are potentially profound. One consequence is that attempts to differentiate between service classes within the network using discriminatory queueing disciplines may become redundant. Another consequence is that a very simple network mechanism, the setting of just a single bit to mark some packets (Ramakrishnan & Jain 1990; Floyd 1994), may be enough to implement a *smart market* (MacKie-Mason & Varian 1994) for the efficient allocation of resources; since, as end-systems see more packets per round-trip time, any randomness associated with whether or not a particular packet is marked becomes less relevant than the proportion of packets marked.

If market mechanisms are able to broadly align supply and demand for communications capacity, then rather simple models of queueing behaviour may be enough to understand the dynamics of various load control algorithms. In Section 3 we review a tractable mathematical model of a network carrying adaptive traffic, from end-systems able to adjust their rates to available bandwidth. Stability, at least on a time-scale of seconds, is established by the explicit construction of a Lyapunov function. Stability on shorter time-scales, comparable with round-trip times in the network, is considered in Section 4. Simple queueing models are used to establish sufficient conditions for local stability, in terms of relationships between the gain parameters of rate or window control algorithms and the marking strategies of resources. A striking observation is that the lower the buffer level at which marking occurs, or the higher the statistical variability of traffic at the packet level, the *fewer* the possibilities for lag-induced oscillatory behaviour.

Gibbens & Kelly (1999b), Turányi & Westberg (1999) and Kelly, Key & Zachary (1999) have described how packet marking at congested resources may be used as the basis for distributed admission control of non-adaptive applications such as traditional telephony, without the need for a parallel signalling network. In Section 5 we explore the behaviour of competing aggregates of adaptive traffic and non-adaptive traffic subject to distributed admission control. In Section 6 we consider short transfers, where a file may have completed its transmission within a round-trip time. There is no possibility for adaptive control on such a short time-scale, and the effect of short transfers is to place a random background load upon the network. A simple queueing model indicates that this background load may well *improve* the stability of the adaptive traffic with which resources are shared.

The marks discussed in this paper give shadow prices at the finest possible granularity; many choices are possible about the level of aggregation at which the marks are reflected as costs or prices to economic agents. For example, the marks could allow the dispersal of charging operations such as metering, accounting and billing to customer systems; Briscoe *et al.* (2000) argue that such an architecture gives simplicity and scalability, without sacrificing commercial flexibility or security. Or the marks could be charges to third-party software running on customer systems, software which undertakes the risks associated with uncertain network loads and user behaviour (Semret & Lazar 1999; Key 1999), and presents an economic agent with pricing choices tailored for that agent.

2. Scalings for queueing delay

In this Section we study queueing delay under various scaling regimes where traffic and capacity grow in line with one another. We shall find that, in two of the three regimes considered, queueing delays decrease.

Let $X[s,t]$ be the amount of work that arrives at a queue in the time interval $[s,t]$. If the queue has a service rate C and no work is lost, then the amount of work buffered in the queue at time 0 is

$$Q = \sup_{u \geq 0} \left\{ X[-u,0] - Cu \right\}. \tag{2.1}$$

Assume that the service rate C is adequate, so that Q is defined as a proper random variable.

Now suppose that the input to the queue in the time interval $[s,t]$ becomes

$$\sum_{i=1}^{c} a\, X_i[bs, bt], \tag{2.2}$$

where the random processes $X_i[s,t]$, $i = 1, 2, \ldots, c$ are independent and each distributed as $X[s,t]$. Let $Q(a,b,c)$ label the random variable defined by the formula (2.1) when $X[s,t]$ is replaced by expression (2.2) and C is replaced by $abcC$. Thus $Q(a,b,c)$ describes the queue length in a system where three forms of scaling have been applied: the original stream $X[s,t]$ has been increased in volume by a factor a, speeded up by a factor b, and c streams have been multiplexed. The mean amount of work arriving at the queue has been increased by a factor abc, as has the service rate at the queue. Let

$$\tau(a,b,c) = Q(a,b,c)/(abcC),$$

the queueing delay under the first-in-first-out queueing discipline.

The impact of the first two forms of scaling on queueing delays is straightforward:

$$\tau(a,b,c) = \tau(1,b,c), \quad \tau(a,b,c) = b^{-1}\tau(a,1,c). \tag{2.3}$$

Certainly $\tau(a,b,c) \leq \tau(a,b,1)$, but the precise impact of the multiplexing factor c depends on the statistical properties of the streams. An illuminating example, illustrating the impact of traffic variability on several scales, is provided by a self-similar traffic model (Willinger *et al.* 1996). If the increments of $X_i[s,t]$ are stationary, with $X_i[0,t] \sim N(\lambda t, \sigma^2 t^{2H})$, corresponding to fractional Gaussian input with Hurst parameter $H \in (0,1)$, then

$$\tau(a,b,c) = c^{-1/(2-2H)}\tau(a,b,1) \tag{2.4}$$

(Norros 1994). Observe that short-range order, the case $H = 0.5$, gives a reduction of queueing time by a factor c; larger values of H, corresponding to increasing degrees of long-range order, give even larger reductions.

We have defined $\tau(a,b,c)$ to be the queueing delay under the first-in-first-out discipline, but the relationships (2.3)–(2.4) hold similarly for other important time periods that can be expressed in terms of the queue length process, such as the busy period of the queue. Detailed investigations of queueing time-scales, using real traffic traces, are described in Courcoubetis *et al.* (1999) and Gibbens & Teh (1999), and some of the implications for queueing networks are developed in Wischik (1999a).

We conclude from relations (2.3)–(2.4) that the volume scaling parameter a does not impact on queueing delay, but the speed and multiplexing parameters b and c both cause substantial reductions in queueing delay. In contrast, propagation delays depend on the speed of light and are unaffected by these scalings. These observations motivate our later models, in which queueing delays and busy periods are assumed small in comparison with propagation delays, and in which packets are served anonymously by a single queue at each resource.

The scalings of this Section assume that capacity is adequate: in an overloaded network it is certainly possible to differentiate between traffic classes using queueing mechanisms. To assure adequate capacity requires suitable capacity provisioning, to deal with growth in demand measured over days, weeks or longer (*cf.* Gibbens *et al.* 2000): we do not consider this topic here. It also requires load control, to deal with fluctuations over seconds or less. In the following Sections we consider load control for various types of traffic.

3. Rate control of adaptive traffic

In this Section we outline a tractable mathematical model of a network carrying rate adaptive traffic, following the development of Kelly, Maulloo & Tan (1998); related approaches are described in Golestani & Bhattacharyya (1998) and Low

& Lapsley (1999). The algorithm described is closely related to the window-based Congestion Avoidance algorithm of Jacobson (1988) for file transfers, but is also intended to model rate adaptive real-time applications.

Consider a network with a set J of *resources*. Let a *route* r identify a non-empty subset of J, and write $j \in r$ to indicate that route r passes through resource j. Let R be the set of possible routes, and suppose that route r carries a flow of rate x_r for each $r \in R$. Suppose that as a resource becomes more heavily loaded it generates feedback signals intended to indicate congestion to the end-systems or users responsible for routes passing through that resource.

How might the end-systems react? Consider the system of differential equations

$$\frac{d}{dt} x_r(t) = \kappa_r \left(w_r - x_r(t) \sum_{j \in r} \mu_j(t) \right) \tag{3.1}$$

for $r \in R$, where

$$\mu_j(t) = p_j \left(\sum_{r: j \in r} x_r(t) \right) \tag{3.2}$$

for $j \in J$. We interpret the relations (3.1)–(3.2) as follows. We suppose that resource j marks a proportion $p_j(y)$ of packets with a feedback signal when the total flow through resource j is y; and that user r views each feedback signal as a congestion indication requiring some reduction in the flow x_r. Then equation (3.1) corresponds to a rate control algorithm for user r that comprises two components: a steady increase at rate proportional to w_r, and a steady decrease at rate proportional to the stream of congestion indication signals received.

It is shown in Kelly *et al.* (1998) that

$$\mathcal{U}(x) = \sum_{r \in R} w_r \log x_r - \sum_{j \in J} \int_0^{\sum_{r: j \in r} x_r} p_j(z) dz$$

is a Lyapunov function for the system of differential equations (3.1)–(3.2), and it is deduced that the unique value x maximizing $\mathcal{U}(x)$ is a stable point of the system, to which all trajectories converge. The variable $\mu_j(t)$ has an interpretation as the implied *shadow price* per unit of flow through resource j at time t, and at the stable point

$$x_r = \frac{w_r}{\sum_{j \in r} \mu_j}. \tag{3.3}$$

The weights $(w_r, r \in R)$ determine the share of scarce resources obtained by different flows, and the rates x given by equation (3.3) have an interpretation in terms of a *weighted proportional fairness* criterion (Kelly 1997, Crowcroft & Oechslin 1998).

In the current Internet, the rate at which a source sends packets is often controlled by TCP, the transmission control protocol of the Internet, implemented as software on end-systems (Jacobson 1988). When a resource within the network becomes overloaded, one or more packets are lost; loss of a packet is taken as an indication of congestion, the destination informs the source, and the source slows down. The source then gradually increases its sending rate until it again receives an indication of congestion. In the future, resources may also have the ability to indicate congestion by marking packets, using an Explicit Congestion Notification bit (Floyd 1994), and current questions concern how packets might be marked and how TCP might be adapted to react to marked packets. Equation (3.1) describes a form of linear increase and multiplicative decrease similar† to that used in Jacobson's (1988) Congestion Avoidance algorithm, but designed to react less severely to congestion indication signals.

Observe that if several flows $r(1), r(2), \ldots, r(n)$ use an identical set of resources, r, say, and share the same gain parameter $\kappa_{r(1)} = \cdots \kappa_{r(n)} = \kappa_r$, then the behaviour of the aggregate

$$x_r(t) = \sum_{i=1}^{n} x_{r(i)}(t) \tag{3.4}$$

may be studied by simply removing the labels $r(1), r(2), \ldots, r(n)$ from the set R, replacing them by the aggregate label r, and using the aggregate weight

$$w_r = \sum_{i=1}^{n} w_{r(i)}. \tag{3.5}$$

One consequence of this important scaling property is that the differential equations (3.1)–(3.2) may be used to study the behaviour of the network at various levels of aggregation. In the next Section we consider forms of packet level behaviour that lead to these equations, with each route r corresponding to an individual flow, while in Section 4 we use the equations to study the behaviour of large aggregates in competition with other forms of traffic.

4. Packet scale behaviour

In this Section we describe how the differential equations (3.1)–(3.2) might arise naturally from the detailed packet level behaviour of end-system and resources, and consider stability over time-scales comparable with round-trip times in the network. The impact of time-lags on the stability of equations (3.1)–(3.2) has been considered in the network context by Kelly, Maulloo & Tan (1998), Tan (1999) and Johari & Tan (2000). Our emphasis here is on how the packet-level behaviour of marking strategies determines the functions $p_j, j \in J$, appearing in these treatments.

We begin by noting the important *self-clocking* feature of Jacobson's (1988) algorithm: the sender uses an acknowledgement from the receiver to prompt a step forward, and this produces a key dependence on the round-trip time T of the connection. In more detail, TCP maintains a window of transmitted but not yet acknowledged packets; the rate x and the window size `cwnd` satisfy the approximate relation $\texttt{cwnd} = xT$. Each positive acknowledgement increases the window size `cwnd` by $1/\texttt{cwnd}$; each congestion indication halves the window size.

Consider a variant constructed by incrementing `cwnd` by

$$\bar{\kappa}\left(\frac{\bar{w}}{\texttt{cwnd}} - f\right)$$

† Differences are discussed in Key *et al.* (1999b), Kelly (2000).

per acknowledgement, where $f = 1$ or 0 according as the packet acknowledged was marked or not. Since the time between update steps is about $T/\texttt{cwnd}$, the expected change in the rate x per unit time is approximately

$$\frac{\bar{\kappa}\left(\frac{\bar{w}}{\texttt{cwnd}} - p\right)/T}{T/\texttt{cwnd}} = \kappa\,(w - xp)\,,$$

where $\kappa = \bar{\kappa}/T$, $w = \bar{w}/T$ and p is the probability of a mark. This expression corresponds with the form of linear increase and multiplicative decrease described by equation (3.1), where the probability a packet is marked somewhere along its route is approximated by the sum of the marking probabilities at the separate resources along that route.

The model (3.1)–(3.2) ignored round-trip times within the network, in order to explore broader aspects of dynamical behaviour. Next we consider the impact of round-trip times on stability, for a very simple example.

Consider a collection of streams all using a single scarce resource. Let the round-trip time be T for each connection, and suppose connections share the same gain parameter κ. Then equations (3.1)–(3.2) become, aggregating all connections into a total flow x and taking the time-lag into account,

$$\frac{d}{dt}\,x(t) = \kappa\,(w - x(t-T)p(x(t-T)))\,. \tag{4.1}$$

The unique equilibrium point of this system does not depend on the round-trip time T, but its stability does. To explore this issue, we first recall some facts about the linear retarded equation

$$\frac{d}{dt}\,u(t) = -\alpha u(t-T), \tag{4.2}$$

where $\alpha > 0$. Solutions to equation (4.2) converge to zero as t increases if $\alpha < \pi/2T$, and the convergence is non-oscillatory if $\alpha < 1/eT$ (see Hale 1977, and Johari & Tan 2000 for a discussion of the multi-dimensional network generalization).

Let x be the equilibrium point of the system (4.1), let $x(t) = x + u(t)$, and write p, p' for the values of the functions $p(.)$, $p'(.)$ at x. Observe that x, p are related by $xp(x) = w$; we assume that the capacity of the resource C is adequate, i.e. $w < C$. Then, linearizing the system (4.1) about x, we obtain equation (4.2) with $\alpha = \kappa(p + xp')$. Hence the equilibrium point of the differential equation (4.1) is stable, and the local convergence is non-oscillatory, if

$$\kappa T\,(p + xp') < e^{-1}; \tag{4.3}$$

stability alone is assured if condition (4.3) is satisfied with e^{-1} replaced by $\pi/2$.

Next we explore some simple forms for the function $p(x)$. Suppose that the workload arriving at the resource over a time period τ is Gaussian, with mean $x\tau$ and variance $x\tau\sigma^2$, and that a packet is marked if when it arrives the workload already present in the queue is larger than a threshold level B. Then, from the stationary distribution for a reflected Brownian motion (Harrison 1985),

$$p(x) = \exp\left\{\frac{-2B(C-x)}{x\sigma^2}\right\}, \tag{4.4}$$

and the condition (4.3) for a non-oscillatory stable equilibrium becomes

$$\kappa T\left(1+\frac{2BC}{x\sigma^2}\right)p(x)<\frac{1}{e}. \tag{4.5}$$

The left-hand side of relation (4.5) is increasing in w $(= xp(x))$, and so the relation is satisfied for any $w < C$ if

$$\kappa T\left(1+\frac{2B}{\sigma^2}\right)<\frac{1}{e}. \tag{4.6}$$

Note that as the threshold level B increases, or as the variability of traffic at the packet level σ^2 decreases, the greater the possibilities for lag-induced oscillatory behaviour. The reason is straightforward: increasing B or decreasing σ^2 causes p' to increase. This increased sensitivity of the resource's load response may compromise stability, unless there is a corresponding decrease in κT, the sensitivity of response of end-systems to marks. (Recall that for the self-clocking window control algorithm described earlier in this section, $\kappa T = \bar{\kappa}$ is the window size decrement made by an end-system in response to a marked packet.) The magnitudes of κ, p' also affect the variance about the equilibrium point in the presence of noise, and speed of convergence (Kelly, Maulloo & Tan 1998): broadly, smaller values of κ or larger values of p' lessen the random fluctuations of rates at equilibrium, while larger values of κ or larger values of p' increase the speed with which changes in parameters such as w may be tracked.

Many variants of the above packet marking algorithm can be analyzed. For example, suppose we mark a packet with probability $1 - exp(-qW)$ if it arrives to find a workload of W already present (the Random Early Marking proposal of Lapsley & Low 1999, a variant of the Random Early Detection proposal of Floyd & Jacobson 1993; Floyd 1994). Then the probability a packet is marked is readily deduced to be

$$p(x)=\frac{qx\sigma^2}{qx\sigma^2+2(C-x)} \tag{4.7}$$

and a simple calculation shows that condition (4.3) is satisfied for any $w < C$ if

$$\kappa T\left(1+\frac{2}{q\sigma^2}\right)<\frac{1}{e}. \tag{4.8}$$

Observe the correspondence between q^{-1} in relation (4.8) and the threshold level B in the earlier relation (4.6): indeed the second algorithm corresponds to randomly resetting the threshold level upon each arrival, according to an exponential distribution with mean q^{-1}.

A suggestion of Gibbens & Kelly (1999a) is that a virtual buffer be maintained of finite size B, and that from the time of a virtual buffer overflow to the end of the virtual buffer's busy period, all packets leaving the real queue be marked. Thus the virtual buffer's contents evolve as if overflow is lost, while the real buffer may or may not have loss. For our Gaussian traffic model, the rate at which workload overflows the virtual buffer is

$$L(x,C)=(C-x)\left(\exp\left\{\frac{2B(C-x)}{x\sigma^2}\right\}-1\right)^{-1}$$

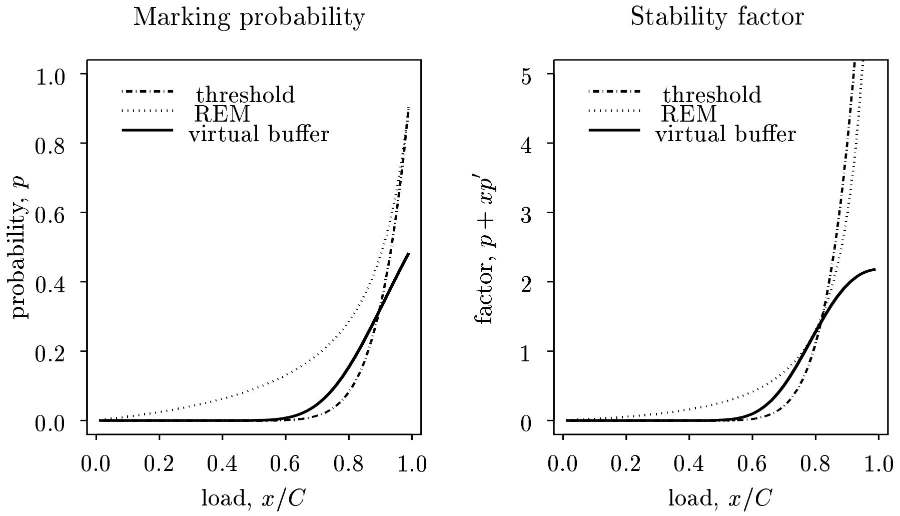

Figure 1. Marking probabilities and stability factors for the threshold, Random Early Marking and virtual buffer algorithms, for $B/\sigma^2 = 5, q = B^{-1}$.

(Harrison 1985) and the proportion of workload marked is given by

$$p(x) = -\frac{d}{dC}\, L(x, C).$$

Virtual loss from the virtual buffer causes the function $p(x)$ to increase more slowly as x approaches C, and this reduces the maximum value of the stability factor $p + xp'$ appearing in relation (4.3) - see Figure 1.

Motivations for different marking strategies, and hence for different functions $p(x)$, are many and varied (in addition to the references above and the seminal paper of Ramakrishnan & Jain 1990, the interested reader should see Wischik 1999b). But, in general, the higher the buffer level at which marking occurs, the greater the possibility of lag-induced oscillatory behaviour. An additional effect becomes apparent if the buffer level at which marking occurs is so high that the time taken by the queue to build up becomes comparable with round-trip times. In this case it becomes necessary to model instantaneous queue size as well as source rates in any dynamical representation. Bolot & Shankar (1990), Fendick *et al.* (1992) and Bonomi *et al.* (1995) have considered deterministic models of this form, establishing the local instability of the equilibrium point whatever gain parameters are used. Fendick *et al.* (1992) provide a careful analysis of the limit cycle behaviour, showing that the queue necessarily hits zero in each cycle.

The approach described in this paper, appropriate when marking occurs at much lower buffer levels, treats the stochastic effects present on the queueing time-scale as essential features of system behaviour, which are averaged out over round-trip times. The specific functions $p(x)$ used in this Section help us understand the impact of threshold levels and traffic variability at the packet level, but their precise forms are not essential for the results of the previous Section, where it is enough that $p(x)$

be smoothly increasing. Provided packet marking algorithms do not destabilize the feedback loops created by adaptive applications, the differential equations of the last section should represent the dynamical behaviour of the network on time-scales longer than a few round-trip times.

5. Distributed admission control

Gibbens & Kelly (1999b), Turányi & Westberg (1999) and Kelly, Key & Zachary (1999) have described how packet marking at congested resources allows resource allocation decisions to be distributed to the edges of networks or to end-systems, and have developed various models for the resulting distributed admission control. In this Section we consider a network carrying rate adaptive real-time traffic, of the form discussed in Section 2, and non-adaptive traffic which is subject to distributed admission control. Our aim is to explore the behaviour of competing aggregates of adaptive and non-adaptive traffic.

Suppose that routes $s \in S$ carry non-adaptive traffic, described as follows. Calls wishing to use route s arrive at the network in a stream of rate ν_s. When a call arrives a set of m_s probe packets are transmitted along route s, and the call is rejected and lost if any of these probe packets is marked. Otherwise the call is accepted and produces a flow of unit mean rate for a call holding period of unit mean. Let $y_s(t)$ represent the aggregate flow along route s. Consider the equations

$$\frac{d}{dt} y_s(t) = \nu_s \left(1 - m_s \sum_{j \in s} \mu_j(t) \right) - y_s(t) \tag{5.1}$$

for $s \in S$, where

$$\mu_j(t) = p_j \left(\sum_{r : j \in r} x_r(t) + \sum_{s : j \in s} y_s(t) \right) \tag{5.2}$$

for $j \in J$, together with equation (3.1) for $r \in R$. Equation (5.2) gives the proportion of packets marked by resource j, and this depends on the total of both adaptive and non-adaptive traffic. As before equation (3.1) describes adaptive traffic on route r, although now x_r, w_r correspond to an aggregate of flows, as described by equations (3.4)–(3.5). Equation (5.1), describing the aggregate non-adaptive traffic on route s, corresponds to an assumption that marking probabilities are small and approximately independent for each of the m_s probe packets of a call attempting route s. The equations treat the average flow along a route as evolving smoothly on a time scale comparable with a call holding time, and correspond to a scaling regime where the number of calls carried on routes is large, a regime studied in detail by Zachary (2000).

A straightforward exercise in partial differentiation yields that the strictly concave function

$$\mathcal{U}(x, y) = \sum_{r \in R} w_r \log x_r + \sum_{s \in S} \frac{1}{m_s} \left(y_s - \frac{y_s^2}{2\nu_s} \right) - \sum_{j \in J} \int_0^{\sum_{r:j \in r} x_r + \sum_{s:j \in s} y_s} p_j(z) dz \tag{5.3}$$

is a Lyapunov function for the system of differential equations (3.1), (5.1)–(5.2); hence the unique value (x, y) maximizing $\mathcal{U}(x, y)$ is a stable point of the system, to which all trajectories converge. If calls have flow rates and holding times that depend on s then the natural generalization of equation (5.1) leads to a Lyapunov function of the form (5.3), but where the parameter ν_s is the product of the arrival rate, the flow rate and the mean call holding time for calls of type s, *i.e.* the offered load on route s measured in Erlangs.

Thus the system implicitly trades off the benefits of the two types of traffic, as represented by the first two terms of expression (5.3). For example, if the load w_r of adaptive traffic on route r gradually increases, then the stable point (x, y) will gradually shift, and in particular x_r will gradually increase, as expression (5.3) places more weight on the term $w_r \log x_r$. Or, if the offered load ν_s of non-adaptive traffic on route s gradually increases, then y_s will gradually increase.

In this Section the equations (3.1), (5.1)–(5.2) are used to represent aggregates of traffic. At a finer level of detail, the number of calls on route s will fluctuate randomly, with y_s corresponding to a mean value (Kelly, Key & Zachary 1999). The impact of non-adaptive traffic on the form of the functions $p_j, j \in J$, will be treated in the next Section.

For file transfers, rather than rate adaptive real-time traffic, different models are appropriate, since the transfer time and the average rate will be inversely proportional, rather than, as in the model above, independent. Gibbens & Kelly (1999a) discuss algorithms for file transfers similar to Jacobson's (1988) Slow Start algorithm, which rapidly increase the transfer rate in the absence of congestion, but back-off otherwise. Key and Massoulié (1999) develop a network model for mixtures of file transfers and rate adaptive real-time traffic in which file transfers are either done at maximal speed or not at all.

6. Short transfers

At the start of a file transfer the congestion window maintained by TCP increases exponentially, doubling every round-trip time, until congestion is detected (the Slow-Start algorithm of Jacobson 1988), whereupon TCP switches to the Congestion Avoidance behaviour sketched in Section 4. Short files may well have completed their transfer before congestion is detected. The performance of a short transfer could clearly be improved by increasing the initial congestion window size, and the rate of exponential growth could also be varied, either up or down. Crowcroft and Oechslin (1998) discuss how a different rate of growth could be implemented; Key and Massoulié (1999) develop a theoretical framework for the choice of the rate of growth, where the choice is made by a risk-averse end-system possessing a prior distribution for the marking probability. Similarly a choice of initial window might reasonably depend upon an end-system's prior distribution, based on past experience (over preceding hours or days) of transfers from a given location.

Feedback from the network, in the form of marked packets, would then affect the choice of initial window and rate of growth of a short transfers, but only on a long time-scale corresponding to the accumulation of information from prior experience of the network. On shorter time-scales, comparable with round-trip times, the effect of short transfers will be to place an uncontrolled and random background load upon the network. If this background load on a resource over a time period τ is Gaussian, with mean $u\tau$ and variance $u\tau v^2$, then equation (4.4) for threshold marking should be replaced by

$$p(x) = \exp\left\{\frac{-2B(C-x-u)}{x\sigma^2 + uv^2}\right\}, \tag{6.1}$$

and condition (4.3) for a non-oscillatory stable equilibrium is satisfied for any $w < C - u$ if

$$\kappa T\left(1 + \frac{2B(C-u)}{(C-u)\sigma^2 + uv^2}\right) < \frac{1}{e}.$$

Alternatively, for the Random Early Marking algorithm, equation (4.7) becomes

$$p(x) = \frac{q(x\sigma^2 + uv^2)}{q(x\sigma^2 + uv^2) + 2(C-x-u)}$$

and condition (4.8) becomes

$$\kappa T\left(1 + \frac{2(C-u)}{q((C-u)\sigma^2 + uv^2)}\right) < \frac{1}{e}.$$

Thus we observe two major effects of the background load: the condition $w < C$ for adequate capacity becomes the more onerous condition $u + w < C$; but, provided this condition is met, the further condition for the equilibrium to be stable and non-oscillatory is more easily met the *larger* the infinitesimal variance v^2 of the background load. Additional variability caused by short transfers lessens the sensitivity of the resource's load response, and this improves the stability of the feedback loops created by adaptive traffic.

7. Conclusion

A consequence of rapidly growing communications capacity may be that it becomes feasible to align supply and demand, and that queueing delays become small in comparison with propagation delays. In such circumstances, a simple packet network may be able to support an arbitrarily differentiated and constantly evolving set of services, by conveying information on incipient congestion to intelligent end-nodes which themselves determine what should be their demands on the packet network. This paper has outlined work on the stability of such a self-managed network, with illustrative results for the cases of adaptive and non-adaptive traffic, and short transfers.

I am grateful to the organizers of the meeting, and to the EPSRC for its support with computing facilities under grant number GR/M09551.

References

J.-C. Bolot and A.U. Shankar (1990). Dynamic behavior of rate-based flow control mechanisms. *ACM Computer Commun. Review* **20**, 35–49.

F. Bonomi, D. Mitra and J.B. Seery (1995). Adaptive algorithms for feedback-based flow control in high-speed wide-area networks. *IEEE J. Selected Areas Communications* **13**, 1267–1283.

B. Briscoe, M. Rizzo, J. Tassel and K. Damianakis (2000) Lightweight policing and charging for packet networks. Third IEEE Conference on Open Architectures and Network Programming. www.labs.bt.com/people/briscorj/

C. Courcoubetis, V. A. Siris, and G. D. Stamoulis (1999) Application of the many sources asymptotic and effective bandwidths to traffic engineering. *Telecommunication Systems.*

J. Crowcroft and P. Oechslin (1998) Differentiated end-to-end Internet services using a weighted proportionally fair sharing TCP. *ACM Computer Communications Review* **28**, 53–67.

K.W. Fendick, M.A. Rodrigues and A. Weiss (1992). Analysis of a rate-based feedback control strategy for long-haul data transport. *Performance Evaluation* **16**, 67–84.

S. Floyd (1994) TCP and Explicit Congestion Notification, *ACM Computer Communication Review* **24**, 10–23. www.aciri.org/floyd/ecn.html

S. Floyd and V. Jacobson (1993) Random Early Detection gateways for congestion avoidance. *IEEE/ACM Transactions on Networking* **1**, 397–413. ftp://ftp.ee.lbl.gov/papers/early.pdf.

R.J. Gibbens and F.P. Kelly (1999a) Resource pricing and the evolution of congestion control, *Automatica* **35**, 1969–1985. www.statslab.cam.ac.uk/~frank/evol.html

R.J. Gibbens and F.P. Kelly (1999b) Distributed connection acceptance control for a connectionless network. *Proceedings of 16th International Teletraffic Congress* editors P.B. Key and D.G. Smith, Elsevier, 941–952. www.statslab.cam.ac.uk/~frank/dcac.html

R. J. Gibbens, S. K. Sargood, F. P. Kelly, H. Azmoodeh, R. Macfadyen and N. Macfadyen (2000) An approach to service level agreements for IP networks with differentiated services. *This volume.*

R.J. Gibbens and Y.C. Teh (1999) Critical time and space scales for statistical multiplexing in multiservice networks. *Proceedings of 16th International Teletraffic Congress* editors P.B. Key and D.G. Smith, Elsevier, 87–96.

S.J. Golestani and S. Bhattacharyya (1998) A class of end-to-end congestion control algorithms for the Internet. In *Proc. Sixth International Conference on Network Protocols.* www.bell-labs.com/user/golestani/

J. Hale (1977) *Theory of Functional Differential Equations.* Springer-Verlag, New York.

Harrison, J.M. (1985). *Brownian Motion and Stochastic Flow Systems.* Krieger.

International Telecommunications Union (1996). Recommendation I371: Traffic control and congestion control in B-ISDN. Geneva.

V. Jacobson (1988) Congestion avoidance and control. *Proc. ACM SIGCOMM '88*, 314–329. ftp://ftp.ee.lbl.gov/papers/congavoid.ps.Z.

R. Johari and D.K.H. Tan (2000) End-to-end congestion control for the Internet: delays and stability.

F.P. Kelly (1997). Charging and rate control for elastic traffic. *European Transactions on Telecommunications* **8**, 33–37. www.statslab.cam.ac.uk/~frank

F.P. Kelly (2000). Mathematical modelling of the Internet. *Proceedings of Fourth International Congress on Industrial and Applied Mathematics.*

F. P. Kelly, P. B. Key, and S. Zachary (1999) Distributed admission control. *Submitted for publication.*

F. P. Kelly, A. K. Maulloo, and D. K. H. Tan (1998) Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society* **49**, 237–252.

P. Key (1999) Service differentiation: congestion pricing, brokers and bandwidth futures. NOSSDAV 1999. www.nossdav.org

P. Key and D. McAuley (1999a) Differential QoS and pricing in networks: where flow control meets game theory. *IEE Proc Software* **146**.

P. Key, D. McAuley, P. Barham, and K. Laevens (1999b) Congestion pricing for congestion avoidance. Microsoft Research report MSR-TR-99-15. http://research.microsoft.com/pubs/

P. Key and L. Massoulié (1999) User policies in a network implementing congestion pricing. Workshop on Internet Service Quality Economics, MIT 1999. http://research.microsoft.com/research/network/disgame.htm

L. Kleinrock (1964) *Communication Nets: Stochastic Message Flow and Delay.* McGraw-Hill, New York.

D. E. Lapsley and S. H. Low (1999). Random Early Marking: an optimisation approach to Internet congestion control. Proceedings of IEEE ICON'99, Brisbane, Australia.

S. H. Low and D. E. Lapsley (1999). Optimization flow control, I: basic algorithm and convergence. *IEEE/ACM Transactions on Networking.* www.ee.mu.oz.au/staff/slow/

J. K. MacKie-Mason and H. R. Varian (1994) Pricing the Internet. In B. Kahin and J. Keller, editors, *Public Access to the Internet.* Prentice-Hall, Englewood Cliffs, New Jersey.

I. Norros (1994). A storage model with self-similar input. *Queueing Systems* **16**, 387–396.

K.K. Ramakrishnan and R.Jain (1990) A binary feedback scheme for congestion avoidance in computer networks. *ACM Transactions on Computer Systems* **8**, 158–181.

N. Semret and A.A. Lazar (1999) Spot and derivative markets in admission control. *Proceedings of 16th International Teletraffic Congress* editors P.B. Key and D.G. Smith, Elsevier, 757–766.

D.J. Songhurst (ed) (1999) *Charging Communication Networks: from Theory to Practice.* Elsevier, Amsterdam.

D.K.H. Tan (1999) *Mathematical models of rate control for communication networks.* PhD dissertation, University of Cambridge. www.statslab.cam.ac.uk/~dkht2/phd.html

Z.R. Turányi and L. Westberg (1999) Load control: lightweight provisioning of Internet resources. www.ericsson.co.hu/ethzrt/

W. Willinger, M.S. Taqqu and A. Erramilli (1996) A bibliographical guide to self-similar traffic and performance modeling for modern high-speed networks. In *Stochastic Networks: Theory and Applications.* eds F.P. Kelly, S. Zachary and I. Ziedins. Royal Statistical Society Lecture Notes Series **4**. Oxford University Press, 339–366.

D. Wischik (1999a) The output of a switch, or, effective bandwidths for networks. *Queueing Systems* **32**, 383–396. www.statslab.cam.ac.uk/~djw1005

D. Wischik (1999b) How to mark fairly. Workshop on Internet Service Quality Economics, MIT 1999.

S. Zachary (2000) Dynamics of large uncontrolled loss networks. *Journal of Applied Probability* **37**. www.ma.hw.ac.uk/~stan/papers