1
Mathematical modelling of the Internet

Frank Kelly
Statistical Laboratory, University of Cambridge,
16 Mill Lane, Cambridge CB2 1SB, UK

Abstract

Modern communication networks are able to respond to randomly fluctuating demands and failures by adapting rates, by rerouting traffic and by reallocating resources. They are able to do this so well that, in many respects, large-scale networks appear as coherent, almost intelligent, organisms. The design and control of such networks present challenges of a mathematical, engineering and economic nature. This paper outlines how mathematical models are being used to address current issues concerning the stability and fairness of rate control algorithms for the Internet.

1 Introduction

In the current Internet, the rate at which a source sends packets is controlled by TCP, the transmission control protocol of the Internet [8], implemented as software on the computers that are the source and destination of the data. The general approach is as follows. 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 TCP then gradually increases its sending rate until it again receives an indication of congestion. This cycle of increase and decrease serves to discover and utilize whatever bandwidth is available, and to share it between flows. In the future, resources may also have the ability to indicate congestion by marking packets, using an Explicit Congestion Notification mechanism [4], and current questions concern how packets might be marked and how TCP might be adapted to react to marked packets.

Just how the available bandwidth within a network *should* be shared raises interesting issues of stability and fairness. Traditionally stability has been considered an engineering issue, requiring an analysis of randomness and feedback operating on fast time-scales, while fairness has been considered an economic issue, involving static comparisons of utility. In future networks the intelligence embedded in end-systems, acting rapidly on behalf of human users, is likely to lessen this distinction.

In Section 2 and 3 we describe a tractable mathematical model of a network and use it to analyse the stability and fairness of a simple rate control algorithm, following closely the development in [10], [11]. Stability is established by showing that, with an appropriate formulation of an overall optimization problem, the network's implicit objective function provides a Lyapunov function for the dynamical system defined by the rate control algorithm. The optimum is characterized by a proportional fairness criterion. The work reviewed in Section 2 forms part of a growing body of research on the global optimization of networks, with important related approaches described by Golestani and Bhattacharyya [5] and Low and Lapsley [16].

In Section 4 we describe a dynamical system which represents more closely the TCP algorithm, in particular the generalization MulTCP proposed by Crowcroft and Oechslin [2], and discuss its stability and fairness. Several authors have described network models of TCP based on systems of differential equations, and the results of Section 4 are close variants on the work reported in [5], [6], [7], [13], [14], and [15].

The dynamical system representations of Sections 3 and 4 model only gross characteristics of the flows through a network. And yet these flows evolve as a consequence of the fine detail of software operating at the packet level. In Sections 4 and 5 we briefly describe, following [6], some of the statistical approaches used to relate these macroscopic and microscopic levels of description, and some of the insights provided by the global model into how packets should be marked at resources and how TCP might react.

There are clear analogies between the approach to large-scale systems described in this paper and that familiar from early work on statistical physics. The behaviour of a gas can be described at the microscopic level in terms of the position and velocity of each molecule. At this level of detail a molecule's velocity appears as a random process, with a stationary distribution as found by Maxwell [18][1]. Consistent with this detailed microscopic description of the system is macroscopic behaviour best described by quantities such as temperature and pressure. Similarly the behaviour of electrons in an electrical network can be described in terms of random walks, and yet this simple description at the microscopic level leads to rather sophisticated behaviour at the macroscopic level: the pattern of potentials in a network of resistors is just such that it minimizes heat dissipation for a given level of current flow (Thomson and Tait [23]). Thus the local, random behaviour of the electrons causes the network as a whole to solve a large-scale optimization problem, an analogy that is discussed further in [9].

Such physical analogies are immensely provocative, but of course the mathematics of large-scale communication networks, such as the global telephone network or the Internet, differs from that of physical models in several respects. Most notably we generally know the microscopic rules (they are coded in our software) but not their macroscopic consequences; and we can *choose* the microscopic rules, an ability that makes it all the more important to learn how to

[1] It is interesting to note the substantial influence this paper had upon Erlang, the early pioneer of research on telephone traffic [3].

predict consequences.

Finally, a note on the title of this paper. The topic discussed here is just one of many that might have been covered in a paper with this title, and no attempt is made here to survey all the important models that arise in the study of communication networks. Some indication of the huge scope for mathematical modelling may be found in the collections [19], [12] and in the recent article by Willinger and Paxson [25].

2 Rate control of elastic traffic

Consider a network with a set J of *resources*. Let a *route* r be a non-empty subset of J, and write R for the set of possible routes. Associate a route r with a user, and suppose that if a rate $x_r > 0$ is allocated to user r then this has *utility* $U_r(x_r)$ to the user. Assume that the utility $U_r(x_r)$ is an increasing, strictly concave function of x_r over the range $x_r > 0$ (following Shenker [21], we call traffic that leads to such a utility function *elastic* traffic). To simplify the statement of results, assume further that $U_r(x_r)$ is continuously differentiable, with $U_r'(x_r) \to \infty$ as $x_r \downarrow 0$ and $U_r'(x_r) \to 0$ as $x_r \uparrow \infty$. Suppose that as a resource becomes more heavily loaded the network incurs an increasing cost, perhaps expressed in terms of delay or loss, or in terms of additional resources the network must allocate to assure guarantees on delay and loss: let $C_j(y)$ is the rate at which cost is incurred at resource j when the load through it is y. Suppose that $C_j(y)$ is differentiable, with

$$\frac{d}{dy} C_j(y) = p_j(y), \tag{2.1}$$

where the function $p_j(y)$ is assumed throughout to be a positive, continuous, strictly increasing function of y bounded above by unity, for $j \in J$. Thus $C_j(y)$ is a strictly convex function.

Next consider the system of differential equations

$$\frac{d}{dt} x_r(t) = \kappa_r \left(w_r(t) - x_r(t) \sum_{j \in r} \mu_j(t) \right) \tag{2.2}$$

for $r \in R$, where

$$\mu_j(t) = p_j \left(\sum_{s: j \in s} x_s(t) \right) \tag{2.3}$$

for $j \in J$. We interpret the relations (2.2)–(2.3) as follows. Suppose that resource j generates a continuous stream of feedback signals at rate $y p_j(y)$ when the total flow through resource j is y; that resource j sends a proportion x_r/y of these feedback signals to a user r with a flow of rate x_r through resource j; and that user r views each feedback signal as a congestion indication requiring some reduction in the flow x_r. Then equation (2.2) corresponds to a rate control algorithm for user r that comprises two components: a steady increase at rate proportional to $w_r(t)$, and a steady decrease at rate proportional to the stream of congestion indication signals received.

Initially we shall suppose that the weights w are fixed. The following theorem is taken from [11].

Theorem 2.1 *If $w_r(t) = w_r > 0$ for $r \in R$ then the function*

$$\mathcal{U}(x) = \sum_{r \in R} w_r \log x_r - \sum_{j \in J} C_j \left(\sum_{s: j \in s} x_s \right) \tag{2.4}$$

is a Lyapunov function for the system of differential equations (2.1)–(2.3). The unique value x maximizing $\mathcal{U}(x)$ is a stable point of the system, to which all trajectories converge.

Proof The assumptions on $w_r, r \in R$, and $p_j, j \in J$, ensure that $\mathcal{U}(x)$ is strictly concave on the positive orthant with an interior maximum; the maximizing value of x is thus unique. Observe that

$$\frac{\partial}{\partial x_r} \mathcal{U}(x) = \frac{w_r}{x_r} - \sum_{j \in r} p_j \Big(\sum_{s: j \in s} x_s \Big);$$

setting these derivatives to zero identifies the maximum. Further

$$\begin{aligned} \frac{d}{dt} \mathcal{U}(x(t)) &= \sum_{r \in R} \frac{\partial \mathcal{U}}{\partial x_r} \cdot \frac{d}{dt} x_r(t) \\ &= \sum_{r \in R} \frac{\kappa_r}{x_r(t)} \left(w_r - x_r(t) \sum_{j \in r} p_j \Big(\sum_{s: j \in s} x_s(t) \Big) \right)^2, \end{aligned}$$

establishing that $\mathcal{U}(x(t))$ is strictly increasing with t, unless $x(t) = x$, the unique x maximizing $\mathcal{U}(x)$. The function $\mathcal{U}(x)$ is thus a Lyapunov function for the system (2.1)–(2.3), and the theorem follows. □

At the stable point

$$x_r = \frac{w_r}{\sum_{j \in r} \mu_j}. \tag{2.5}$$

This equation has a simple interpretation in terms of a charge per unit flow: the variable μ_j is the *shadow price* per unit of flow through resource j.

Next suppose that each user r is able to vary its own weight $w_r = w_r(t)$, but is required to pay the network at the rate $w_r(t)$ per unit time. The parameter $w_r(t)$ can be viewed as the *willingness to pay* of user r; Songhurst [22] provides a discussion of related charging and accounting mechanisms, and of their implementation. In particular, we suppose that user r is able to monitor its rate $x_r(t)$ continuously, and chooses to vary smoothly the parameter $w_r(t)$ so as to satisfy

$$w_r(t) = x_r(t) U_r'(x_r(t)) : \tag{2.6}$$

this corresponds to a user who deduces from observation of the network that its current charge per unit flow is $\lambda_r = w_r(t)/x_r(t)$, and then chooses $w_r = w_r(t)$ to solve the optimization problem

$$\begin{array}{ll} \text{maximize} & U_r\left(\dfrac{w_r}{\lambda_r}\right) - w_r \\ \text{over} & w_r \geq 0. \end{array}$$

In [11] a similar proof to that of Theorem 2.1 is used to establish the following result.

Theorem 2.2 *The strictly concave function*

$$\mathcal{U}(x) = \sum_{r\in R} U_r(x_r) - \sum_{j\in J} C_j\left(\sum_{s:j\in s} x_s\right). \tag{2.7}$$

is a Lyapunov function for the system of differential equations (2.1)–(2.3), (2.6), and hence the unique value x maximizing $\mathcal{U}(x)$ is a stable point of the system, to which all trajectories converge.

Thus if each user r is able to choose its own weight, or willingness to pay, w_r, and does this so as to optimize its own utility less payment, then the overall system will converge to the rate allocation x maximizing the aggregate utility (2.7). At the optimum the relation (2.5) will again hold, but where now the weights w have been chosen by users themselves.

3 Fairness

The expression for an aggregate utility (2.7) requires that utilities and costs are additive. In this Section we consider what can be done without this assumption.

The vector

$$\left(U_r(x_r), r\in R;\ -\sum_{j\in J} C_j\Big(\sum_{s:j\in s} x_s\Big)\right) \tag{3.1}$$

lists the utilities of the various users and the cost of the network. A rate allocation x is *Pareto efficient* if any alteration that strictly increases at least one of the components of the vector (3.1) must simultaneously strictly decrease at least one other component [24]. By the strict monotonicity of the components of the vector (3.1), *every* allocation x is Pareto efficient. Say that an allocation x is *achievable* by weights w if x maximizes the function (2.4), and hence is a stable point of the system (2.1)–(2.3). From the strict concavity of the components of the vector (3.1), together with the supporting hyperplane theorem, it follows that every allocation x is achievable by some choice of weights w.

Pareto efficiency becomes more interesting when routing choices are allowed. If a user's utility is a function of the sum of its flows over several routes, then it is quite possible to construct flow patterns which are not Pareto efficient. The flow patterns which emerge as stable points of the routing generalization [11] of the system (2.1)–(2.3) are, however, just those that are Pareto efficient.

Thus many, or even all, rate allocations x are Pareto efficient. Is there a rationale for choosing any particular one? This question is often couched in terms of fairness, and the following construction is familiar in contexts as diverse as political philosophy [20] and data networks [1]. Say that the allocation x is *max–min fair* if for any other vector x^* that leaves the final component of (3.1) constant and for which there exists r such that $x_r^* > x_r$, there also exists s such that $x_s^* < x_s \leq x_r$. The max–min fairness criterion gives an absolute priority to the smaller flows, in the sense that if $x_s \leq x_r$ then no increase in x_r, no matter how large, can compensate for any decrease in x_s, no matter how small. A max–min fair allocation is necessarily Pareto efficient, and thus achievable by some choice of the weights w.

It is possible to describe the stable point of the system (2.1)–(2.3) in the language of fairness. Say that the allocation x is *proportionally fair* for the weight vector w if for any other vector x^* that leaves the final component of (3.1) unaltered, the aggregate of weighted proportional changes is zero or negative:

$$\sum_{r \in R} w_r \frac{x_r^* - x_r}{x_r} \leq 0.$$

This criterion favours smaller flows, but less emphatically than max–min fairness. The allocation (2.5) is proportionally fair for the weight vector w.

4 Network models of TCP

The differential equations (2.1)–(2.3) represent a system that shares several characteristics with Jacobson's TCP algorithm [8] operating in the current Internet, but also has several differences, which we now discuss.

A flow through the Internet will receive congestion indication signals, whether from dropped or marked packets, at a rate roughly proportional to the size of the flow; and the response of Jacobson's congestion avoidance algorithm to a congestion indication signal is to halve the size of the flow. Thus there are *two* multiplicative effects: both the number of congestion indication signals received *and* the response to each signal scale with the size of the flow. A further important feature of Jacobson's algorithm [8] is that it is *self-clocking*: the sender uses an acknowledgement from the receiver to prompt a step forward and this produces an important 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 `cwnd` $= xT$. Each positive acknowledgement increases the window size `cwnd` by 1/`cwnd`; each congestion indication halves the window size. Crowcroft and Oechslin [2] have proposed that users be allowed to set a parameter m, which would *inter alia* multiply by m the rate of additive increase and make $1 - 1/2m$ the multiplicative decrease factor in Jacobson's algorithm. The resulting algorithm, MulTCP, would behave in many respects as a collection of m single TCP connections; its smoother behaviour is more plausibly modelled by a system of differential equations.

For MulTCP the expected change in the congestion window `cwnd` per update step is approximately

$$\frac{m}{\texttt{cwnd}}\ (1-p) - \frac{\texttt{cwnd}}{2m}\ p \tag{4.1}$$

where p is the probability of congestion indication at the update step. Since the time between update steps is about $T/\texttt{cwnd}$, the expected change in the rate x per unit time is approximately

$$\frac{\left(\frac{m}{\texttt{cwnd}}\ (1-p) - \frac{\texttt{cwnd}}{2m}\ p\right)/T}{T/\texttt{cwnd}} = \frac{m}{T^2}\ (1-p) - \frac{x^2}{2m}\ p.$$

Motivated by this calculation, we model MulTCP by the system of differential equations

$$\frac{d}{dt}\,x_r(t) = \frac{m_r}{T_r^2} - \left(\frac{m_r}{T_r^2} + \frac{x_r(t)^2}{2m_r}\right)\sum_{j\in r}\mu_j(t), \tag{4.2}$$

where T_r is the round-trip time for the connection of user r, and where $\mu_j(t)$ is again given by equation (2.3). We again view $p_j(y)$ as the probability a packet produces a congestion indication signal at resource j, but we do not now insist that $p_j(y)$ satisfies relation (2.1). Note that if congestion indication is provided by dropping a packet, then equation (4.2) approximates the probability of a packet drop along a route by the sum of the packet drop probabilities at each of the resources along the route. To hold $x(t)$ to the positive orthant, augment equation (4.2) with the interpretation that its left-hand side is set to zero if $x_r(t)$ is zero and its right-hand side is negative.

Theorem 4.1 *The function*

$$\mathcal{U}(x) = \sum_r \frac{\sqrt{2}m_r}{T_r}\arctan\left(\frac{x_rT_r}{\sqrt{2}m_r}\right) - \sum_{j\in J}\int_0^{\sum_{s:j\in s}x_s} p_j(y)dy \tag{4.3}$$

is a Lyapunov function for the system of differential equations (4.2), (2.3). The unique value x maximizing $\mathcal{U}(x)$ is a stable point of the system, to which all trajectories converge.

Proof The function $\mathcal{U}(x)$ is strictly concave on the positive orthant, with

$$\frac{\partial}{\partial x_r}\mathcal{U}(x) = \frac{m_r}{T_r^2}\left(\frac{m_r}{T_r^2} + \frac{x_r^2}{2m_r}\right)^{-1} - \sum_{j\in r}p_j\Big(\sum_{s:j\in s}x_s\Big).$$

Thus

$$\frac{d}{dt}\mathcal{U}\big(x(t)\big) \quad = \quad \sum_{r\in R}\frac{\partial\mathcal{U}}{\partial x_r}\cdot\frac{d}{dt}\,x_r(t)$$

$$= \sum_{r \in R} \left(\frac{m_r}{T_r^2} + \frac{x_r(t)^2}{2m_r} \right)^{-1} \left(\frac{d}{dt} x_r(t) \right)^2 ,$$

establishing that $\mathcal{U}(x(t))$ is strictly increasing with t, unless $x(t) = x$, the unique x maximizing $\mathcal{U}(x)$. The function $\mathcal{U}(x)$ is thus a Lyapunov function for the system, and the theorem follows. □

Comparing the form (4.3) with the aggregate utility function (2.7), we see that MulTCP can be viewed as acting as if the utility function of user r is

$$\frac{\sqrt{2} m_r}{T_r} \arctan \left(\frac{x_r T_r}{\sqrt{2} m_r} \right)$$

and as if the network's cost is the final term of the form (4.3).

The system (4.2) has stable point

$$x_r = \frac{m_r}{T_r} \left(\frac{2(1-p_r)}{p_r} \right)^{1/2} ,$$

where

$$p_r = \sum_{j \in r} p_j .$$

If we were to omit the factor $(1-p)$ in the expression (4.1) then the implicit utility function for user r would be

$$-\frac{2m_r^2}{T_r^2 x_r}$$

and the stable point of the system would be

$$x_r = \frac{m_r}{T_r} \sqrt{\frac{2}{p_r}}, \tag{4.4}$$

recovering the inverse dependence on round trip time and the inverse square root dependence on packet loss familiar from the literature on TCP [17][2].

We have not presumed that the function $p_j(y)$ is defined by relation (2.1); indeed $p_j(y)$ could not in general satisfy this relation if congestion can only be signalled by dropping packets. But the functions $p_j(y)$ and $C_j(y)$ are both strictly monotone functions of y, and this is enough to ensure that the stable point is Pareto efficient. The same conclusion cannot be reached when users or

[2]The constant of proportionality, $\sqrt{2}$, is sensitive to the difference between MulTCP and m single TCP connections, as well as other features of TCP implementations discussed in [17]. With m single TCP connections the connection affected by a congestion indication signal is more likely to be one with a larger congestion window. This bias towards the larger of the m connections increases the final term of the expectation (4.1) and decreases the constant of proportionality.

the network have routing choices, and choose routes generating fewer congestion indication signals. For example, suppose that

$$p_j(y) = \left(\frac{y}{N_j}\right)^{M_j}, \tag{4.5}$$

a form that would arise if resource j were modelled as an $M/M/1$ queue with service rate N_j packets per unit time at which a packet is marked with a congestion indication signal if it arrives at the queue to find at least M_j packets already present. Suppose the network suffers unit loss for each such arriving packet: then

$$\frac{d}{dy} C_j(y) = (M_j + 1)\left(\frac{y}{N_j}\right)^{M_j},$$

and so the packet marking probability understimates the shadow price by a factor $M_j + 1$. This discrepancy is enough to allow the construction of examples exhibiting Braess' paradox, where the addition of capacity to a network, followed by routing adaptation to the extra capacity, leads to a new stable point at which every user is worse off (for a discussion of Braess' paradox couched in the optimization framework of this paper see [9]). Such a stable point is certainly not Pareto efficient.

We end this section with such an example. Consider the network illustrated in the Figure, where TCP connections send packets from node a to node d, with some connections routed via node b and others via node c. Suppose congestion indication signals are generated at the various queues at rates (4.5), and returned

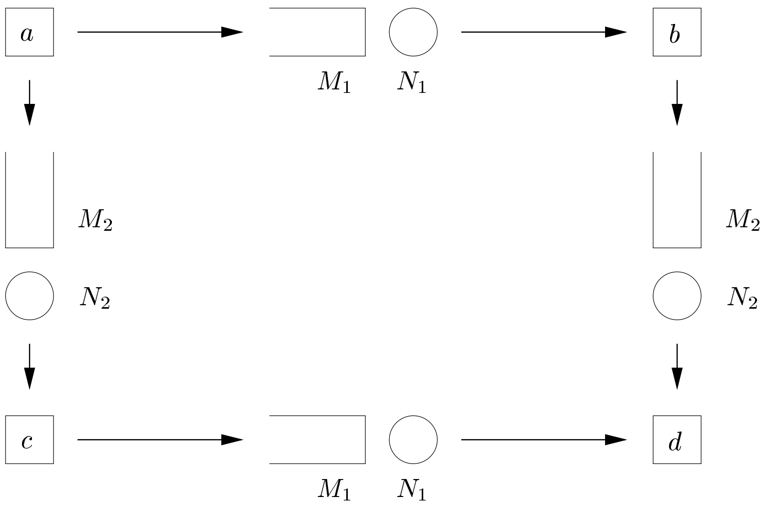

FIG. 1. Braess' paradox for a network. Adding a link between nodes b and c, followed by routing adaptation to the extra capacity, may *reduce* the throughput of all connections.

to the sources at node a via other routes not explicitly modelled. Suppose that there are 2000 connections in total, that round-trip times are 1000 time units, and that throughputs are determined by relation (4.4). Suppose that connections are routed so as to equalize the rate at which congestion indication signals are generated along different routes. Then approximately 1000 connections will be routed along each of the two possible routes, and, with the parameter choices $(N_1, M_1, N_2, M_2) = (10, 50, 20, 5)$, $p_1 = 0.006$ for the horizontal links, while $p_2 = 0.019$ for the vertical links. Next suppose a high capacity link, at which no congestion indication signals are generated, is introduced between nodes b and c. This allows some connections to be routed along a third route, from a to b to c to d. At the new equilibrium, the load on the horizontal links becomes about 1044 connections and the load on the vertical links about 956 connections, with $p_1 = p_2 = 0.013$, and the throughput of each connection *drops* by nearly 3 per cent.

5 Packet marking strategies

We have briefly indicated in Section 4 how TCP software on users' computers can perform operations, at the packet level, that implement the end-system behaviour modelled in the differential equations (4.2); similarly a relatively minor variation, discussed in [6], [13], [14], is enough to achieve the end-system behaviour modelled by the differential equations (2.2). In this Section we consider how the behaviour required of *resources* may be implemented at the packet level. We have argued in earlier sections that the rate of packet marking at resource j should signal its shadow price, defined as the derivative (2.1), and we next describe how this be done in a simple robust manner, at the packet level.

For simplicity of exposition let us suppose that packets are all of equal length. A simple model describing a queue with a finite buffer is as follows. Let Y_{t-1} be the number of packets that arrive at the resource in the interval $(t-1, t]$, and let Q_t be the queue size at time t. Then the recursion

$$Q_t = \min\{M, Q_{t-1} - I\{Q_{t-1} > 0\} + Y_{t-1}\}$$

describes a queue with a buffer capacity of M that is able to serve a single packet per unit time; the number of packets lost at time t is

$$[Q_{t-1} - I\{Q_{t-1} > 0\} + Y_{t-1} - M]^+.$$

Define a busy period to end at time t if $Q_{t-1} = 1$, and a busy period to begin at time t if $Q_{t-1} \leq 1$ and $Q_t \geq 1$.

The impact of an additional packet upon the total number of packets lost, its shadow price, is relatively easy to describe. Consider the behaviour of the queue *with* the additional packet included in the description of the queue's sample path. Then the additional packet increases the number of packets lost by one if and only if the time of arrival of the additional packet lies within a *critical congestion interval*, defined as a period between the start of a busy period and the loss, within the same busy period, of a packet; otherwise the additional packet does not affect the number of packets lost. Packets arriving during critical congestion intervals should, ideally, be marked.

It will often be difficult to determine the shadow price of a packet while the packet is passing through the queue; it will, in general, be unclear whether or not the current busy period will end before a packet is lost. Nevertheless the above ideal behaviour gives considerable insight into the form of sensible marking strategies, and suggests simple, robust strategies which result in a rate of packet marking at a resource which is equal to its shadow price [6].

The above discussion shows that the network shadow prices, the variables $\mu_j(t)$ appearing in equation (2.3), may be identified, at least statistically, on the sample path describing load at resource j. This identification provides the required linkage between microscopic, packet-level, marking strategies, often of an apparently crude and statistical nature, and sophisticated macroscopic behaviour interpretable as the global optimization of a large-scale network.

Bibliography

1. D. Bertsekas and R. Gallager (1987) *Data Networks*. Prentice-Hall.
2. 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.
3. A.K. Erlang (1925) A proof of Maxwell's law, the principal proposition in the kinetic theory of gases. In E. Brockmeyer, H.L Halstrom, H.L. and A. Jensen, *The life and works of A.K. Erlang*, Copenhagen, Academy of Technical Sciences, 1948. 222–226.
4. S. Floyd (1994) TCP and Explicit Congestion Notification, *ACM Computer Communication Review* **24**, 10-23. www.aciri.org/floyd/ecn.html
5. 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/
6. R.J. Gibbens and F.P. Kelly (1999) Resource pricing and congestion control, *Automatica* **35.** www.statslab.cam.ac.uk/~frank/evol.html
7. P. Hurley, J. Y. Le Boudec and P. Thiran (1998) A note on the fairness of additive increase and multiplicative decrease. In *Proc. 16th International Teletraffic Congress*, Edinburgh, P. Key and D. Smith (eds). Elsevier, Amsterdam. 467-478.
8. V. Jacobson (1998) Congestion avoidance and control. In *Proc. ACM SIGCOMM '88*, 314-329. A revised version, joint with M.J. Karels, is available via ftp://ftp.ee.lbl.gov/papers/congavoid.ps.Z.
9. F.P. Kelly (1991). Network routing. *Phil. Trans. R. Soc. Lond. A*,**337**, 343-367. www.statslab.cam.ac.uk/~frank/
10. F.P. Kelly (1997). Charging and rate control for elastic traf-

fic. *European Transactions on Telecommunications* **8**, 33-37. www.statslab.cam.ac.uk/~frank/elastic.html

11. 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. www.statslab.cam.ac.uk/~frank/rate.html.
12. F. P. Kelly, S. Zachary and I. Ziedins (eds) (1996) *Stochastic Networks: Theory and Applications*. Royal Statistical Society Lecture Notes Series, Vol **4**. Oxford University Press, Oxford.
13. P. Key and D. McAuley (1999) Differential QoS and pricing in networks: where flow control meets game theory. *IEE Proc Software* **146**.
14. P. Key, D. McAuley, P. Barham, and K. Laevens. Congestion pricing for congestion avoidance. Microsoft Research report MSR-TR-99-15. http://research.microsoft.com/pubs/
15. S. Kunniyar and R. Srikant. End-to-end congestion control schemes: utility functions, random losses and ECN marks. University of Illinois.
16. S. H. Low and D. E. Lapsley. Optimization flow control, I: basic algorithm and convergence. To appear, *IEEE/ACM Transactions on Networking.* www.ee.mu.oz.au/staff/slow/
17. M. Mathis, J. Semke, J. Mahdavi, and T. Ott (1997) The macroscopic behaviour of the TCP congestion avoidance algorithm. *Computer Communication Review* **27**.
18. J.C. Maxwell (1860) Illustrations of the dynamical theory of gases. *Philosophical Magazine* **20**, 21-37.
19. D. Mitra (ed) (1995) Advances in the fundamentals of networking. *IEEE J. Selected Areas in Commun.* **13**, 933-1362.
20. J. Rawls (1971) *A Theory of Justice*, Harvard University Press.
21. S. Shenker (1995) Fundamental design issues for the future Internet. *IEEE J. Selected Areas in Commun.* **13**, 1176-1188.
22. D.J. Songhurst (ed) (1999) *Charging Communication Networks: from Theory to Practice*, Elsevier, Amsterdam.
23. W. Thomson and P.G. Tait (1879) *Treatise on Natural Philosophy*, Cambridge.
24. H.R. Varian (1992). *Microeconomic Analysis,* third edition, Norton, New York.
25. W. Willinger and V. Paxson (1998) Where Mathematics meets the Internet. *Notices of the American Mathematical Society*, **45**, 961-970. www.ams.org/notices/