

The Clifford Paterson Lecture, 1995: Modelling Communication Networks, Present and Future

F. P. Kelly

Phil. Trans. R. Soc. Lond. A 1996 **354**, 437-463

doi: 10.1098/rsta.1996.0016

Email alerting service

Receive free email alerts when new articles cite this article - sign up in the box at the top right-hand corner of the article or click [here](#)

To subscribe to *Phil. Trans. R. Soc. Lond. A* go to:
<http://rsta.royalsocietypublishing.org/subscriptions>

The Clifford Paterson Lecture, 1995

Modelling communication networks, present and future

BY F. P. KELLY

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

Contents

	PAGE
1. Introduction	438
(a) Outline	439
2. Loss networks	440
(a) Erlang's formula	440
(b) A network model	440
(c) Examples	442
3. Dynamic routing and stability	444
(a) Instability	444
(b) Dynamic alternative routing	446
(c) Extensions	447
4. Coalitions and global routing	450
(a) Dominant coalitions	450
(b) The core	451
(c) Transit payments	452
(d) Network evolution	454
5. Broadband traffic and statistical sharing	454
(a) Effective bandwidths	455
(b) Multiplexing models	457
(c) Charging mechanisms	459
6. Conclusion	460
References	461

Modern communication networks are able to respond to randomly fluctuating demands and failures by allowing buffers to fill, 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. In this paper I describe some of the models that have proved useful in the analysis of stability, statistical sharing and pricing, in systems ranging from the telephone networks of today to the information superhighways of tomorrow.

Phil. Trans. R. Soc. Lond. A (1996) **354**, 437–463

Printed in Great Britain

437

© 1996 The Royal Society

TeX Paper

1. Introduction

It was with some trepidation that I gave the Clifford Paterson Lecture, not least because my background is that of a mathematician, and the topic on which I spoke, the mathematical modelling of communication networks, is but loosely connected with the area of the Lecture, namely electrical science and technology. Indeed I did wonder if perhaps a mistake had been made: in the Yearbook of the Society there are three Kellys, and either of the others would appear to have better credentials in the area of electrical science!

I have, however, taken some comfort from the work of Clifford Paterson himself (Ryde 1949). He was the founding Director and moving spirit of GEC's Research Laboratories in the first half of this century, and had a remarkable spread of interests. He put great efforts into promoting a closer liaison between science and industry, and into the development of national and international standards, both areas of relevance for this lecture.

Perhaps I should begin by saying what the lecture will *not* contain. I shall not discuss in any detail the way that future communication networks will look to the user. I share the common belief that advances in computing and communications will transform our economy and society, and in ways we shall find hard to predict. Few who have seen the impact of the Internet on the methods of working of research groups could disagree with this point, but it is a point that is made frequently, perhaps repetitively, in our current media.

Nor shall I attempt to outline the enormous impetus that has been given to the subject of *mathematics* by the demands of modern communication networks. Probability and information theory, logic and combinatorics, number theory and cryptography – whole areas of mathematical science now have as their primary focus of application problems that arise in computing and communication. But there is no chance of covering such diversity in an hour, or a month, even if I were able to do so.

Instead I plan to consider one general issue, the stability and control of large-scale networks, and to briefly sketch some of the mathematical models which have played a role in its understanding.

The global telephone network is the largest single machine yet constructed by the human race. It is sufficiently tightly coupled to be regarded as a single machine since, after all, its function is to connect any two telephones anywhere in the world. The reliable operation of this network requires the solution of several design issues, both within the subnetworks run by individual telephone companies, and at the interfaces between these subnetworks. Some of these issues are of an engineering nature, concerning stability and control over fast timescales. Others are of an economic nature, involving coordination between distinct commercial entities. In the development of standards, technical and commercial aspects are often inextricably linked.

The emergence of broadband networks has made more acute the need for a better understanding of these issues. Developments in telecommunication technology are leading towards networks which will allow several widely disparate traffic streams to share the same broadband channel. A call, which might be a mixture of voice, video and data, would appear to the network as a stream of cells, and the hope is that calls with a broad range of burstiness characteristics and quality of service requirements can be efficiently integrated to share common resources. Coordination issues now extend beyond the network to include the form of contracts and control signals between the network and users.

My aim in this lecture is to show how mathematical models can help make precise some of the questions in these areas, and how analysis of models may provide important insights into what are often otherwise rather ill-defined and contentious issues. The topics I have chosen to discuss are some that I have found especially interesting, and illustrate different aspects of the central challenge: to understand how large-scale networks may be designed to function coherently. No attempt will be made to survey all of the important models that arise in the study of communication networks: a fuller perspective and some indication of the huge scope for mathematical modelling may be found in the wide-ranging recent collections edited by Mitra (1995) and Steenstrup (1995).

The topics chosen are all outcomes of collaborations with colleagues at the Statistical Laboratory in Cambridge and BT Laboratories at Martlesham Heath, and I especially acknowledge my debt to Richard Gibbens, Philip Hunt, Peter Key, Neil Laws, David Songhurst, Stephen Turner and Martin Whitehead, who have been authors and co-authors of the original papers on which much of this lecture is based. I also take this opportunity to thank Sir Robert Clayton, then Technical Director of GEC, for his thoughtful and prescient introduction, eleven years ago, to the group at BT Laboratories.

(a) *Outline*

The organization of the paper is as follows. In §2 there is described an abstract mathematical model of a system where arriving demands require simultaneous use of a number of scarce resources. The classical example of this model is a telephone network, but much of the recent mathematical development has been motivated by applications to local area networks, mobile radio and broadband networks. The breadth of application is not, of course, entirely fortuitous: the model emerged as an abstraction of the common features of resource allocation problems across a range of technologies.

Dynamic routing plays a major role in the reliable operation of large-scale networks. Section 3 outlines some of the dangers associated with the development of dynamic routing schemes, in particular the danger of instability, and describes some simple decentralized schemes designed to implicitly seek out near optimal routing patterns.

The development of dynamic routing schemes for international networks raises several further issues. Joint action by a group of operating companies may lead to a surplus of benefits over costs, but how should these benefits be divided, and which mechanisms for division encourage stable cooperative routing and capacity management? Section 4 describes how the theory of games can shed light on some of these issues.

Any quantitative model of a broadband network requires some measure of the resource usage of a connection. Section 5 outlines how the notion of an effective bandwidth may be developed to provide such a measure. The effective bandwidth of a connection depends sensitively upon the statistical properties of the connection as well as its quality of service requirements, and an interesting issue concerns how much of the effort of statistical characterization should fall upon the network and how much upon the user responsible for the connection. A charging mechanism is described which encourages the cooperative sharing of information and characterization effort between users and the network.

2. Loss networks

This section describes how several disparate models can be cast within a common mathematical framework describing the simultaneous possession of a subset of resources by a call. In some examples the resources are physical, for example transmission capacity over a link, but in some examples flexibility of routing allows physical resources to be pooled, and the essential constraints are on logical resources representing pooled subsets of physical resources.

(a) Erlang's formula

In 1917 the Danish mathematician A. K. Erlang published his famous formula,

$$E(\nu, C) = \frac{\nu^C}{C!} \left[\sum_{n=0}^C \frac{\nu^n}{n!} \right]^{-1}, \quad (2.1)$$

for the loss probability of a telephone system (Brockmeyer *et al.* 1948). The problem considered by Erlang can be phrased as follows. Calls arrive at a link as a Poisson process of rate ν . The link comprises C circuits, and a call is blocked and lost if all C circuits are occupied. Otherwise the call is accepted and occupies a single circuit for the holding period of the call. Call holding periods are independent of each other and of arrival times and are identically distributed with unit mean. Then *Erlang's formula* (2.1) gives the proportion of calls that are lost, and for over eighty years this simple formula has helped telecommunication engineers determine the amount of capacity, C , that is required for a given level of offered traffic, ν , if the loss probability is to be acceptable.

But what happens if the system consists of many links, and if calls of different types (perhaps voice, video or conference calls) require different resources? The next subsection describes a generalization of Erlang's model which treats a network of links, and which allows the number of circuits required to depend upon the call. The classical example of this model is a telephone network, and it is natural to couch its definition in terms of calls, links and circuits. The model also arises naturally in the study of local area networks, multiprocessor interconnection architectures, optical networks, mobile radio and broadband networks, and much of the modern interest in the model is a consequence of developments in these technologies. The term 'circuit-switched' is common in some application areas, where it is used to describe systems in which, before a request (which may be a call, a task or a customer) is accepted, it is first checked that sufficient resources are available to deal with each stage of the request. The essential features of the model is that a call makes simultaneous use of several resources and that blocked calls are lost. In §2c we indicate something of the breadth of application of the basic model, through a brief discussion of some simple examples of repacking, cellular radio and queueing networks.

(b) A network model

Consider a network with J links, labelled $1, 2, \dots, J$, and suppose that link j comprises C_j circuits (figure 1). A call on route r uses A_{jr} circuits from link j , where $A_{jr} \in \mathbb{Z}_+$. Let \mathcal{R} be the set of possible routes. Calls requesting route r arrive as a Poisson stream of rate ν_r , and as r varies it indexes independent Poisson streams. A call requesting route r is blocked and lost if on any link j , $j = 1, 2, \dots, J$, there are less than A_{jr} circuits free. Otherwise the call is connected and simultaneously holds A_{jr} circuits from link j , $j = 1, 2, \dots, J$, for the holding period of the call. The call

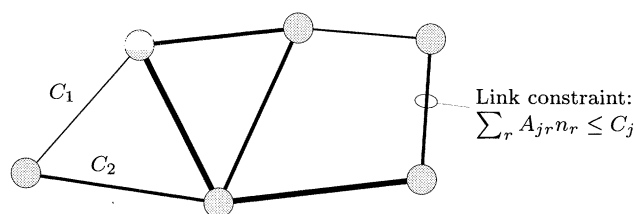


Figure 1. A telephone network. The linear constraint $\sum_r A_{jr} n_r \leq C_j$ expresses the condition that the number of circuits required from link j by calls in progress cannot exceed the capacity of link j .

holding period is independent of earlier arrival times and holding periods; holding periods of calls on route r are identically distributed with mean μ_r .

Let $n_r(t)$ be the number of calls in progress at time t on route r , let $\alpha_r = \nu_r / \mu_r$, and define the vectors $n(t) = (n_r(t), r \in \mathcal{R})$ and $C = (C_1, C_2, \dots, C_J)$. Then the stochastic process $(n(t), t \geq 0)$ has a unique stationary distribution and under this distribution $\pi(n) = P\{n(t) = n\}$ is given by

$$\pi(n) = G(C)^{-1} \prod_{r \in \mathcal{R}} \frac{\alpha_r^{n_r}}{n_r!}, \quad n \in \mathcal{S}(C), \quad (2.2)$$

where

$$\mathcal{S}(C) = \{n \in \mathbb{Z}_+^{\mathcal{R}} : An \leq C\} \quad (2.3)$$

and $G(C)$ is the normalizing constant (or partition function)

$$G(C) = \left(\sum_{n \in \mathcal{S}(C)} \prod_{r \in \mathcal{R}} \frac{\alpha_r^{n_r}}{n_r!} \right). \quad (2.4)$$

This result is easy to check in the case where holding times are exponentially distributed: then $(n(t), t \geq 0)$ is a Markov process and the distribution (2.2) satisfies the detailed balance conditions,

$$\pi(n) \cdot \nu_r = \pi(n + e_r) \cdot (n_r + 1), \quad n, n + e_r \in \mathcal{S}(C), \quad (2.5)$$

where $e_r = (\mathbb{I}[r' = r], r' \in \mathcal{R})$ is the unit vector describing just one call in progress on route r .

Most quantities of interest can be written in terms of the distribution (2.2) or the partition function (2.4). For example let L_r be the stationary probability that a call requesting route r is lost: since the arrival stream of calls requesting route r is Poisson,

$$1 - L_r = \sum_{n \in \mathcal{S}(C - Ae_r)} \pi(n) = G(C)^{-1} G(C - Ae_r). \quad (2.6)$$

Observe that the distribution (2.2) is simply that of independent Poisson random variables truncated to a linearly constrained region (2.3): thus from expression (2.6) we obtain Erlang's formula (2.1) in the case of a single truncated Poisson random variable. The mathematical analysis of the above model has been the subject of intensive research over the last decade. For more complex networks the explicit form (2.6) may be hard to compute (Louth *et al.* 1994), but there now exist many methods of approximation and analysis: these extend to generalizations of the model which allow more complex routing choices, and permit consideration of issues such as dynamic routing and network planning. For reviews see Kelly (1991b), Ross (1995).

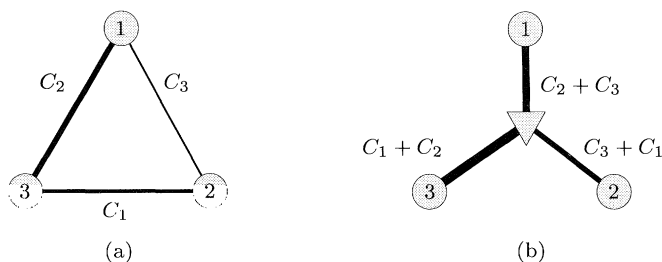


Figure 2. Repacking. If calls in network (a) may be repacked, then the essential capacity constraints attach not to the links of the network, but rather to the cut capacities out of each node, and network (a) becomes equivalent to network (b).

(c) Examples

We next discuss some simple examples, which help to indicate the breadth of application of the above network model.

(i) Repacking

Consider first the network of figure 2a, where a call between two nodes requires a single circuit from the direct link joining these nodes, or a single circuit from each of the links on the two-link path joining these nodes. Suppose that calls may be *repacked*: that is, a call may be shifted from one route to another while the call is in progress, and that an arriving call is accepted provided the calls already in progress can be repacked to provide space for the arriving call. If $n_{\alpha\beta}$ is the number of calls in progress between nodes α and β , then $n = (n_{12}, n_{23}, n_{31})$ behaves exactly as the stochastic process n describing calls in progress in the loss network illustrated in figure 2b. Thus under repacking, the essential capacity constraints are not attached to the links representing the *physical* resources of network 2a, but rather to *logical* resources representing the cut capacities out of the nodes of network 2a.

For more general network topologies the identification of logical resources may be less intuitively obvious, but remains possible: see Kelly (1991b, §3.3) for a fuller discussion.

(ii) Cellular radio

As our next example, consider the cellular radio network of figure 3a. We suppose there are C frequencies in total, and that a frequency may not be simultaneously used in two adjacent cells. Suppose that frequencies may be reallocated while calls are in progress, if this allows an additional call to be accepted. If n_j is the number of calls in progress in cell j , then the essential constraint on the vector $n = (n_j, j = 1, 2, \dots, 7)$ is that

$$n_\alpha + n_\beta + n_\gamma \leq C$$

for all triples (α, β, γ) such that the cells labelled α, β, γ are mutually adjacent (that is, meet at a point). In terms of our original loss network model, an equivalent system can be constructed as follows: place a logical 'link' with a capacity of C circuits at each of the vertices of the graph appearing in figure 3a, and suppose that a call arriving at a cell requires one circuit from each of the six 'links' located at the vertices of that cell. Again the essential capacity constraints are linear, and expressed in terms of logical resources.

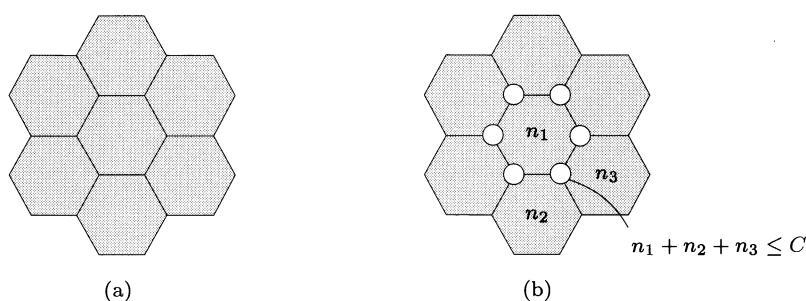


Figure 3. A cellular radio network. There are C frequencies in total in network (a), and re-use is possible in non-adjacent cells. The network is equivalent to network (b), in which for example a call in the centre cell requires a 'circuit' from each of six logical 'links' each of capacity C .

(iii) A queueing network

Our final example of this section is a queueing network, represented diagrammatically in figure 4. Suppose that a call (or connection) consists of a stationary stream of packets, to be sent along a defined route through the network. (In an asynchronous transfer mode network, packets are of fixed length and are called, confusingly in view of the previous example, *cells*.) Suppose that calls arrive randomly, but are accepted only if the resulting traffic intensity at each queue along the call's route remains less than one. Then this network operates essentially as a loss network, with a single linear constraint attached to each queue. Of course the condition that traffic intensities be held less than one ensures only stability of the queues: if buffers are finite, or delays have to be bounded, then tighter conditions will be necessary, as we discuss in §5.

Within a queueing network, repacking corresponds to an ability to send packets within a connection along different routes. For a discussion of the fascinating resource pooling issues that can arise in queueing networks, see Laws (1992), Kelly & Laws (1993).

3. Dynamic routing and stability

Dynamic routing schemes have an important role to play in ensuring the reliable operation of large-scale networks. In particular, routing schemes should allow the network to respond robustly to failures and overloads, lessen the impact of forecasting errors, and permit flexible use of network resources. The problem is to decide which resources have spare capacity and which are overloaded: this may vary over time in response to the random nature of traffic, and over the network as a consequence of asymmetries in traffic and capacity patterns.

The two major concerns associated with the development of dynamic routing schemes are instability and complexity. Suppose that 'intelligent' nodes within a network react to blocked routes by rerouting calls along more resource-intensive paths. This in turn may cause later calls to be rerouted, and the cascade effect may lead to a catastrophic change in the network's behaviour. Section 3*a* discusses this effect for a symmetric fully connected network. Of course real networks are not symmetric, and a major aim of dynamic routing schemes is to cope with mismatches between traffic and capacity patterns. In §3*b* we describe a scheme which attempts to use simple local rules to solve this complex matching problem. In §3*c* we describe some extensions of the basic ideas underlying this scheme to other network architectures.

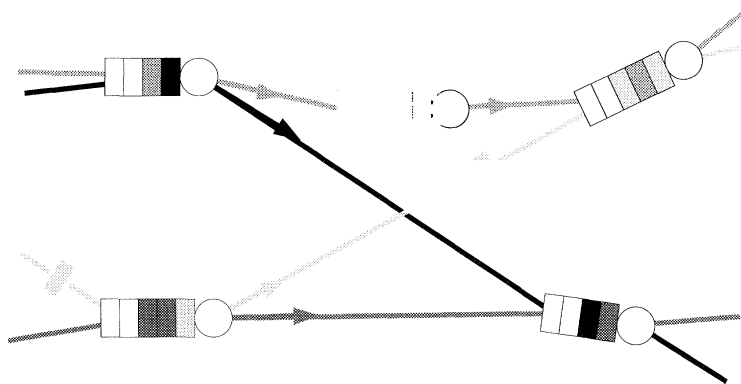


Figure 4. A queueing network. A connection produces a stream of cells, to be sent along a route through the network. For the traffic intensity at each queue to be less than one, the numbers of calls in progress must satisfy a linear constraint for each queue. If buffers are finite or delays have to be bounded then tighter conditions will be necessary, as discussed in §5. Propagation times may be much longer than queueing times: for transatlantic voice calls, there may be space for hundreds of cells within a buffer while tens of thousands of cells may be in flight between queues.

(a) *Instability*

Suppose that K nodes are linked to form a complete graph. Between any pair of nodes calls arise at rate ν , and there is a link of capacity C . If there is a spare circuit on the link joining the end points of a call then the call is accepted and carried by that circuit. Otherwise the call chooses at random one of the $K - 2$ two-link paths joining its end points: the call is accepted on that path if both links have a spare circuit, and is lost otherwise. Accepted calls have holding periods with unit mean.

Let B be the probability that a link is full. Then as the number of nodes K grows, the probability B approaches a solution of the equation,

$$B = E(\nu[1 + 2B(1 - B)], C), \quad (3.1)$$

where E is Erlang's formula (2.1). This simple and important conclusion is a consequence of limit theorems only recently established by Crametz & Hunt (1991) and Graham & Méléard (1993).

Equation (3.1) has multiple solutions over significant ranges of the parameters ν and C . This observation, first made by Nakagome & Mori (1973), and its relationship with the underlying model, have been investigated extensively by means of simulation and approximation (see, for example, Ackerley 1987; Gibbens *et al.* 1990). For certain parameter values, there are two distinct locally stable modes of behaviour. There is also a hysteresis effect: if ν is varied slowly the mode which obtains may depend not just on the current value of ν but also upon whether ν approached this value from above or below. An intuitive explanation is easy to provide. One solution corresponds to a mode in which blocking is low, calls are mainly routed directly and relatively few calls are carried on two-link paths. Another solution corresponds to a mode in which blocking is high and many calls are carried over two-link paths. Such calls use two circuits each, and this additional demand on network resources may cause a substantial number of subsequent calls also to attempt two-link paths. Thus a form of positive feedback may keep the system in the high blocking mode.

When a circuit is used by a call, the expected net effect upon the network, measured

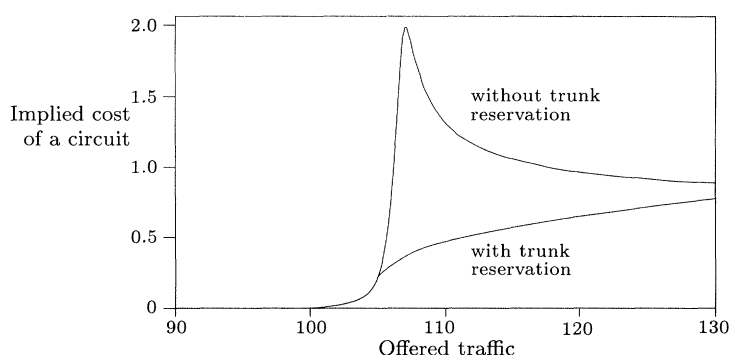


Figure 5. Implied costs. The implied cost of a circuit is the expected net effect upon the network, measured in terms of calls lost that would otherwise have been carried, of removing a single circuit for one unit of time. Note that without trunk reservation the implied cost can rise well above 1. Figures 5 and 6 are calculated for a network with a large number of nodes which is symmetric, fully connected, and where each link has capacity 120 circuits; the figures are adapted from Kelly (1990) and Gibbens & Kelly (1990) respectively.

in terms of calls lost that would otherwise have been carried, is called the *implied cost* of a circuit. The upper curve in figure 5 illustrates the implied cost of a circuit for a network with links of capacity $C = 120$. Observe the steep rise of the implied cost, up to a value of almost 2 at an offered traffic of about $\nu = 107$. At this point carrying an additional directly routed call causes the subsequent loss of nearly *two* calls. The additional call can cause this much damage, on average, because it can force a later call to be rerouted: this uses more network resource and may in turn cause further calls to be rerouted. This cascade effect can be modelled as a branching process: instability corresponds to supercriticality of the branching process, and to an *infinite* implied cost. The behaviour of implied costs is thus a much more sensitive indicator of inefficiency than the onset of instability – indeed in a network with links of capacity 120 where a call is allowed just one alternative route there is no instability.

A natural control is to allow a link to reject alternatively routed calls if the number of idle circuits on the link is less than or equal to a certain value, s , say. This method of giving priority at a link to certain traffic streams is known as *trunk reservation*, and the parameter s is called the trunk reservation parameter for the link. With trunk reservation in place alternative routing is capable of improving performance, and trunk reservation is widely used in telephone networks (Songhurst 1980). Houck *et al.* (1994) provide a recent review of several other forms of overload control, as well as a fascinating discussion of the AT&T network failure of 15 January 1990.

Figure 6 shows the minimal loss probability using the best choice of trunk reservation parameter for various values of the number of alternative routes allowed. The curve labelled 0 illustrates Erlang's formula, while the curve labelled ∞ corresponds to an unbounded number of alternative routes. Trunk reservation is not only simple to define and implement, it has also been shown recently to be asymptotically optimal as the number of nodes increases (Hunt & Laws 1992, 1993), in the sense that the curve labelled ∞ is a lower bound on the loss probability under *any* dynamic routing scheme, and is approached by a trunk reservation policy as the number of nodes and alternative routes increases. Also shown in figure 6 is a lower bound on the loss probability under any dynamic routing scheme, when repacking is allowed. This bound, also, is approached by a simple threshold routing scheme, as the number of nodes and alternative routes increase (Kelly 1994a, §3).

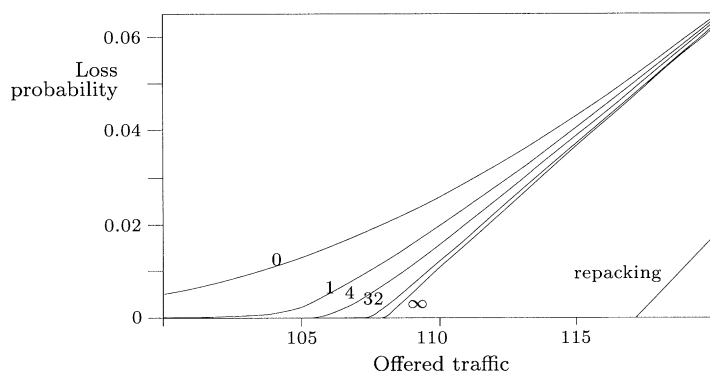


Figure 6. Minimal loss probabilities. The curves are labelled according to the number of alternative routes (0, 1, 4, 32, ∞) to which a call has access. In figures 5 and 6 the offered traffic measures the expected number of calls that request connection between any two given nodes per mean call holding time.

(b) Dynamic alternative routing

Trunk reservation can thus control instability and, in symmetric networks, approach optimality. But when traffic patterns and capacity are asymmetric, which routes should be used as alternatives? We next describe dynamic alternative routing (DAR), a simple and effective mechanism for making these choices, developed by researchers at Cambridge and BT Laboratories at Martlesham.

Suppose there are K nodes in the network, with the link $\{i, j\}$ joining nodes i and j having capacity C_{ij} . Each link is assigned a trunk reservation parameter s_{ij} , and each source-destination pair (i, j) stores the identity of its current tandem $k(i, j)$ for use in two-link alternative routes. A call between nodes i and j is first offered to the direct link and a call is always routed along that link if there is a free circuit. Otherwise, the call attempts the two-link alternative route via tandem node k with trunk reservation applied to both links. If the call fails to be routed via k , this call is lost and, further, the identity of the tandem node is reselected (at random perhaps) from the set $\{1, 2, \dots, K\} \setminus \{i, j\}$. Note especially that the tandem node is not reselected if the call is successfully routed on either the direct link or the two-link alternative route. Mees (1986) has coined the term *sticky random routing* to emphasize this property of the scheme. In practice it has been found simpler to reselect a tandem node by cycling around a fixed random permutation; the point is that reselection is not based on any collected data, only the important information that a call has just failed.

Let $p_k(i, j)$ denote the long-run proportion of calls between i and j which are offered to tandem node k , and let $q_k(i, j)$ be the long-run proportion of those calls between i and j and offered to tandem node k which are blocked. Then, under uniform reselection,

$$p_a(i, j)q_a(i, j) = p_b(i, j)q_b(i, j) \quad a, b \neq i, j.$$

Observe that this simple ergodic result is exact for either random reselection or reselection using a fixed permutation. More generally, suppose the DAR mechanism for reselection of the tandem node between i and j chooses node k with long-run frequency f_k where $\sum_{k \neq i, j} f_k = 1$. Then each selection of node k is paired with a

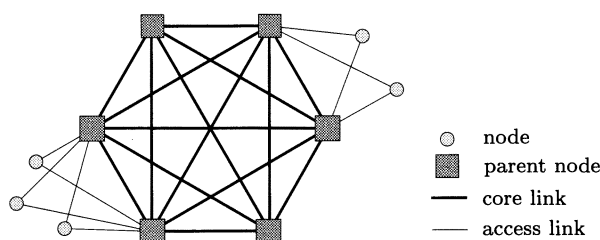


Figure 7. A dual-parented network. Two nodes with no parent nodes in common have four three-link paths between them.

failed call via node k , and so

$$p_a(i, j)q_a(i, j) : p_b(i, j)q_b(i, j) = f_a : f_b \quad a, b \neq i, j. \quad (3.2)$$

Observe that if the blocking $q_k(i, j)$ is high on the path through the tandem node k , then the proportion of overflow routed via node k will be low. This gives some insight into the means by which the routing scheme implicitly adapts to overloads and failures.

For further insight consider next the effects of mismatches between traffics and capacities. Suppose that traffic is greater than capacity on some links, and less than capacity on others. Is it possible for the excess traffic from overloaded links to be assigned to alternative routes which do not themselves clash with one another; and, if so, is it possible for this to be achieved by a simple algorithm?

The answer is that a simple greedy algorithm will perform well for this problem (Hajek 1987), and that DAR acts as a dynamic version of such an algorithm. It also has several similarities with probabilistic hill-climbing techniques such as simulated annealing, and Gibbens (1988) and Mitra & Seery (1991) have discussed its use as an algorithm to investigate random graph and network routing problems. Consider the multi-commodity flow problem of routing $n(n-1)$ distinct flows between each ordered pair of n nodes. As a linear program this problem has $O(n^3)$ variables, since each of $n(n-1)$ ordered node pairs has one direct route and $(n-2)$ two-link alternatives available, and $O(n^2)$ constraints, associated with the $\frac{1}{2}n(n-1)$ links and distinct flows (cf. Gibbens & Kelly 1990). The problem has a natural mapping onto a computing engine with $O(n^2)$ parallel processors. DAR tackles the nonlinear stochastic version of this problem using a probabilistic hill-climbing technique performed by $O(n^2)$ parallel processors, namely the occupancy levels of the $\frac{1}{2}n(n-1)$ links and the $n(n-1)$ tandem pointers $k(i, j)$. In this way the network itself operates as a distributed computer, executing a highly parallel randomized algorithm to solve a complex optimization problem.

The analytical methods referred to in §2, together with simulation, have been very important in the investigation of the performance of DAR under a wide range of failure and overload conditions. For an account of this work the reader is referred to Stacey & Songhurst (1987), Gibbens (1988), Gibbens *et al.* (1988, 1995), Key (1988), Key & Whitehead (1988).

(c) Extensions

The sticky principle underlying DAR can be readily extended to other network architectures. In this section we briefly sketch two examples taken from Gibbens *et al.* (1988, 1993, 1995) and Wroe *et al.* (1990).

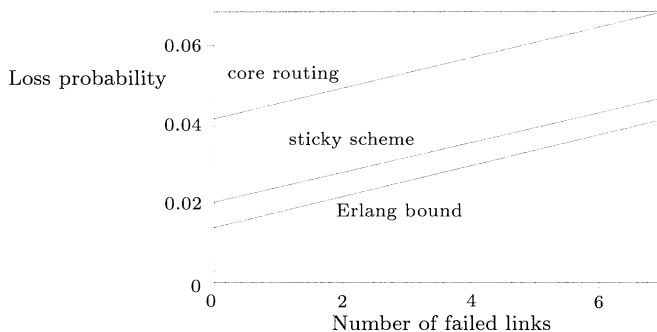


Figure 8. Performance under link failures. Note that the sticky scheme, operating with little explicit information on network state, outperforms the optimal core routing scheme, and indeed is not far from the Erlang bound upon the performance of any dynamic routing scheme. This example is taken from Gibbens *et al.* (1993): there are 24 nodes in the core network and a further 216 nodes in the access network.

(i) *Multiparented networks*

Consider the network structure illustrated in figure 7. There are two levels: the *core* network, and the *access* network. The nodes of the core network are fully connected, while nodes in the access network are connected to two or more nodes (termed *parents*) in the core network.

We shall consider three families of dynamic routing schemes for this network architecture. The first family assumes that a call must enter the core network through a single defined parent, and must leave the core network through a single defined parent. The core network may, however, route the call between these two nodes as it pleases: either on the direct single link joining the parents, or on any of the two-link or longer paths through the core network between the two parents. The second family of schemes assumes that the network is able to route a call by any route through the core network connecting any parent of the originating node with any parent of the destination node.

Dynamic routing schemes from the first family essentially ignore the multiparented structure, and are often used in networks where the access network and the core network belong to different operating companies. Schemes from the second family are able to exploit the additional routes made available by multiple parents, and can thus achieve a very high degree of resource pooling: indeed their performance approaches the bound given by Erlang's formula, with arguments the sum of the offered traffics to the core network, and the sum of the link capacity of the core network. However, to achieve this very good performance requires the originating node for a call to route the call to a parent whose identity depends on the instantaneous state of various links in the core network.

The third family of schemes is based on two simple principles: the *lastchance* principle, whereby a link gives priority, through trunk reservation, to calls which will be lost if they are blocked; and the *sticky* principle, whereby a route which successfully carries a call is left as a preferred route, while a route which is unable to carry a call is replaced.

In figure 8 we illustrate the performance of schemes from these three families when randomly chosen links are failed, one after another. The top curve describes the performance of an optimal scheme from the first family, while the bottom curve describes the Erlang bound, approached by schemes from the second family. It is remarkable

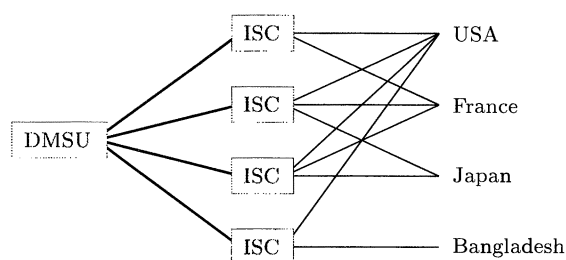


Figure 9. International access network. Nodes of the core network, termed digital main switching units (DMSUs), are connected to international switching centres (ISCs). Every DMSU is connected to every ISC, but a foreign carrier may be connected to only some of the ISCs.

that the sticky scheme, operating with no detailed knowledge of network state other than that conveyed by call blocking events, significantly outperforms optimal core routing: indeed its performance is closer to the Erlang bound than to optimal core routing. Our conclusion is that simple schemes, easily implemented and analysed, are able to achieve most of the additional advantages allowed to dynamic routing schemes by multiparenting. For further discussion of sticky schemes for multiparented networks, including a description of their behaviour under *node* failure events, see Gibbens *et al.* (1993).

(ii) International access networks

BT has several international switching centres (ISCs). These connect the nodes of the core network (termed digital main switching units, or DMSUs) to foreign carriers. For reliability, countries with large traffic volumes (for example, the US, France, Germany) are connected to three or four ISCs – see figure 9. Every DMSU is connected to every ISC, and the network between DMSUs and ISCs is called the *international access network*. The routing problem within the access network is to spread traffic over ISCs in such a way that the (relatively) expensive international circuits are used efficiently. This endeavour is made particularly difficult by the many traffic profiles the network must handle (those to the US and Japan are illustrated in figure 10), and general uncertainty about which ISCs foreign carriers will use for incoming traffic.

It is interesting to contrast two simple routing schemes. Under the first, called *crankback*, a call is offered in turn to each ISC to which the destination country's ISCs are connected. Besides the switching overhead, this scheme can cause localized areas of congestion, with calls being sent relentlessly to ISCs whose outgoing links are blocked. Traffic which must use these ISCs will suffer a poor quality of service. Under the second scheme a call is offered to just one current-choice ISC. If the call is successful, then the current choice is retained; otherwise the call is rejected and the current choice is reset. Note that, although each call has just one choice, the scheme will tend to balance load across ISCs.

The relative performance of these two schemes, and the perhaps surprisingly good performance of the second scheme, are discussed in Wroe *et al.* (1990) and Gibbens *et al.* (1995). One of the interesting effects noted is the counterintuitive behaviour possible under crankback, where adding capacity to the network may cause the performance to become worse. This phenomenon, known as Braess's paradox, has been observed in several network flow models, and is discussed more fully in Kelly (1991c).

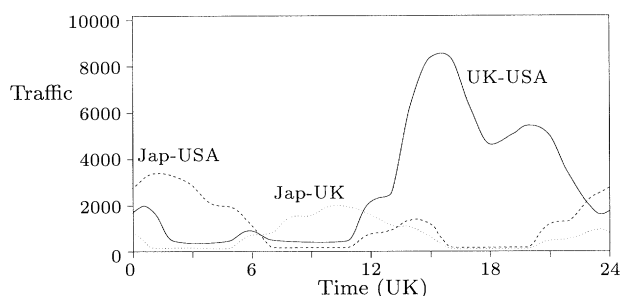


Figure 10. Daily traffic profiles between Japan, UK and USA, where traffic is measured in occupied circuits. The marked non-coincidence of busy hours is apparent.

4. Coalitions and global routing

Present usage of the international telecommunications network is inefficient. Routing generally takes the form of direct routes between countries, using circuit capacity jointly owned by the origin and destination carriers. Limited rerouting of traffic via transit countries takes place only under network management control, responding to link failures or other exceptional conditions. Yet timezone effects produce a marked non-coincidence of busy hours (see figure 10), and there exists the potential for large savings in network costs and improvement in the quality of network service.

The joint control of a large network when parts of the network belong to different operating companies raises several new issues concerning the partitioning of benefits. Joint action will lead to a surplus of benefits over costs, relative to the present arrangements. How should these benefits be divided in such a way that the contributions of individual carriers are respected, and which mechanisms for division encourage stable, cooperative routing and capacity management? In this section we describe how a game-theoretic approach, first applied in this area by Gibbens *et al.* (1991) and more recently by van Golstein Brouwers (1994), can help provide insight into the issues involved.

(a) Dominant coalitions

Let N be the set of players, and for subsets $S \subseteq N$ define the *value* of the coalition S to be

$$V(S) = (\text{cost separate}) - (\text{cost with coalition, } S).$$

Here 'cost separate' is the cost to the members of the coalition of providing links between themselves if they route all traffic between themselves directly. In contrast 'cost with coalition' is the cost to members of the coalition of providing links between themselves if they route traffic between themselves optimally. The optimal routing pattern varies according to the time of day, but can be calculated straightforwardly from a multi-commodity network flow model. (The stochastic effects of random call arrivals remain important for the design of dynamic routing schemes, but, provided a stable and efficient scheme is used, flows are sufficiently large that stochastic effects have little consequence for cost calculations. For a recent evaluation of dynamic routing schemes, including DAR, for a global network, see Ash & Huang (1994).) Thus $V(S)$ is the potential saving to be had by the formation of the coalition S . Table 1 shows $V(S)$ for the most advantageous coalitions of three, four and five players from the set $N = \{A, C, F, J, UK, USA\}$, representing Australia, Canada, France, Japan, UK and USA.

Table 1. *Relative values of the most advantageous coalitions of three, four and five players* (Figure 10 and tables 1 and 2 are abstracted from Gibbens *et al.* (1991), where the data sources and methods used are described more fully. (The first comprehensive collection of international traffic profiles and circuit costs has recently been published by Schmitt & Watanabe (1994).))

subset S	separate	coalition	$V(S)$	saving (%)
J UK USA	13895	11134	2761	20
A UK USA	12610	10406	2204	17
F J USA	6904	5609	1295	19
A F USA	5600	4801	799	14
C J UK	3995	3199	796	20
A C UK	3869	3127	742	19
A J UK USA	18558	13573	4985	27
F J UK USA	20248	16733	3515	17
A F J USA	18847	15982	2865	15
C J UK USA	15860	13044	2816	18
A F J UK USA	24990	19188	5802	23
A C J UK USA	20667	15570	5097	25

Some important features are apparent immediately from table 1. For example, note the large potential savings to be had by a three-member coalition between Japan, UK and USA, or between Australia, UK and USA. Yet the four-member coalition of Australia, Japan, UK and USA achieves a saving even greater than the combined saving of these two three-member coalitions. Thus it is reasonable to view Australia and Japan as complementary rather than substitutes.

(b) *The core*

Next we consider how the potential saving, $V(N)$, might be divided between the members of the set N . Define an *imputation* $(x_i, i \in N)$ to be a vector of real numbers such that $\sum_{i \in S} x_i = V(S)$, representing a division between the members of the set N of the value $V(N)$. Define the *core* to be the set of imputations $(x_i, i \in N)$ such that

$$\sum_{i \in S} x_i \geq V(S) \quad \forall S \subseteq N. \quad (4.1)$$

It is thus the set of imputations which leave no coalition in a position to improve the payoff to each of its members.

The core is typically quite large, and we next consider a subset defined by a tightening of condition (4.1). The *strong ε -core* is the set of imputations $(x_i, i \in N)$ such that

$$\sum_{i \in S} x_i \geq V(S) + \varepsilon \quad \forall S \subseteq N, S \neq \emptyset, N.$$

We may regard ε as an additional inducement to the formation of a coalition smaller

Table 2. *Extreme points of the least core, with $N = \{A, F, J, UK, USA\}$*

(Note the unfavoured position of France within the least core. The imputation determined by the Shapley value (Jones 1980) is also shown.)

A	F	J	UK	USA	winner
409	409	2528	409	2048	J,USA
409	409	2528	2048	409	J,UK
1877	409	886	2221	409	A,UK
1877	409	2528	409	579	A,J
1877	409	2528	579	409	
429	409	886	3670	409	UK
409	409	906	3670	409	
826	415	1076	1479	2003	Shapley value

than N . As ε increases the strong ε -core shrinks: define the *least core* to be the smallest non-empty strong ε -core (Shubik 1982).

For the set of players $N = \{A, F, J, UK, USA\}$, the least core is the convex hull of the set of extreme points shown in table 2. Any imputation within this convex hull is stable against a subset deciding to form a smaller coalition, even if the subset is offered an additional inducement of 409 for doing so: 409 is thus a measure of the additional effort involved in extending the coalition $\{A, J, UK, USA\}$ to include France that would make it not quite worth the effort. (Of course this example is just one illustration, with one particular choice of the set N ; note that the countries A, J, UK, USA have overlapping communities of interest through language, migration and financial sector activities.)

(c) *Transit payments*

For a coalition between Japan, UK and USA the potential capacity savings to Japan and the USA are greater than to the UK. Correspondingly the UK carries a greater peak transit traffic, and also a higher volume of transit traffic. It seems reasonable to expect the asymmetric allocation of the benefits of capacity savings to be offset to some greater or lesser extent by payments for transit traffic. But how should transit payments be calculated? We next consider whether allowing individual carriers to set their own transit charges, or to devise their own charging structures, may impair the efficiency of a coalition.

Consider a three-node network, where each carrier acts in turn as a monopoly provider of spare capacity between the other two nodes at different times of day, and assume there is ample capacity for transit traffic, as might be the case in the early period of a coalition. We consider three price structures and determine the actions of the transit carrier that maximize its revenue, assuming a cost minimizing response on the part of other carriers.

In the first price structure the transit carrier announces in advance a fixed transit charge per time unit per circuit. ('In advance' means that the potential payers of this charge have sufficient time to expand their own capacity, if they so choose, in response to the announced charge.) In the second price structure, a charge is levied proportional to the maximum number of calls simultaneously carried over some period – for example, circuits may be leased. Finally a hybrid pricing structure is described, which has the twin properties that it maximizes the transit revenue of each carrier, and encourages the collectively optimal use of the network.

(i) *Transit charges*

Suppose the transit carrier announces a transit charge per occupied circuit per unit of time ('per paid hour'). What will the transit carrier do to maximize its transit revenue, assuming a cost minimizing response on the part of the other carriers to whatever transit charge it announces?

Let the *load curve* $G(x)$ be the number of hours over a day that transit traffic is at a higher level than x . We assume that if the transit charge is set at t then sufficient additional direct capacity will be provided so that on average transit overflow will occur for z or less hours per day, where $zt = c$. (It is cheaper to build an extra circuit, at cost c per day, than to pay a transit charge of t per hour for more than z hours per day.) The revenue collected from transit traffic will thus be

$$t \int_{G^{-1}(c/t)}^C G(x) dx,$$

where, by our assumption of ample capacity for transit traffic, $G(C)$ is zero. Let \hat{t} maximize this expression, and let $\hat{z} = c/\hat{t}$. Then the transit carrier will maximize its revenue by choosing a transit charge of \hat{t} ; in calculating \hat{t} it reasonably assumes that the carriers paying this charge will expand their direct capacity so that they offer transit traffic for a maximum of \hat{z} hours per day.

(ii) *Circuit levies*

Suppose next that a charge is levied proportional to the maximum number of calls simultaneously carried over some period. If the charge levied is d per circuit per day, then the rational response for the payers of this levy is as follows:

- (i) if $d > c$ then offer no overflow – provide direct capacity instead;
- (ii) if $d < c$ then offer overflow freely up to a maximum of $G^{-1}(d/w)$, where w is the benefit per paid hour of carrying a call. Above this level calls should be lost – it costs more in transit costs to carry them than the benefit they generate.

The transit carrier should then choose $\hat{d} \leq c$ so as to maximize its transit revenue

$$dG^{-1}(d/w).$$

(iii) *A hybrid pricing scheme*

Neither transit charges nor circuit levies will, in general, lead to collectively optimal behaviour: under either of these price structures less overflow traffic will be carried than under a joint network optimum. It is interesting, however, that a hybrid price structure can induce optimal behaviour. Suppose that the total transit revenue is determined as

$$cG^{-1}\left(\frac{c}{w}\right) + w \int_{G^{-1}(c/w)}^C G(x) dx.$$

This charge has the following interpretation: up to a number of circuits $G^{-1}(c/w)$ a levy of c per circuit per day is extracted; for traffic above this level, a charge of w per paid hour is made. This price structure simultaneously possesses the following two properties: the payers of the charge have no incentive to depart from the network optimal policy, either by building extra capacity, or by blocking calls; and no other policy can achieve a greater transit revenue for the transit carrier. There may thus be collective advantage in allowing transit carriers to independently devise their own, possibly complex, pricing strategies. A transit carrier will find it in its interest to move to a scheme similar in effect to the above scheme (since this maximizes its revenue), and such a scheme encourages maximum collective network efficiency.

(d) Network evolution

The earlier discussion assumes ample spare capacity for transit traffic. As, over time, traffic loads increase, the question of where to add capacity arises. In a minimal cost network the three cut sets identified in figure 2 would each just saturate at some time in the day. Let us make the simplifying assumption that different cut sets saturate at different times, and that circuit costs and call benefits are uniform over the network. Let $G_i(x)$ be the number of hours over the day that offered traffic to and from node i exceeds the level x . Then the collectively optimal network has

$$C_1 = (G_2^{-1} + G_3^{-1} - G_1^{-1})(c/2w).$$

Some consequences are immediate from this calculation. With direct routing only, the best choice of capacity for a link would have the link saturated for c/w hours per day. With the use of alternative routing, the best of choice of capacity for links would have the cut sets from nodes saturated for $c/2w$ hours per day.

The key point identified in §4c(iii) carries over to the network just described, and is familiar from economic theory (Varian 1992): if a monopolist has sufficient freedom over pricing then it can maximize net benefit and appropriate this benefit for itself. What makes our model unusual, and potentially attractive as a means of allocating benefits within the coalition, is that each of the three players gets to act as the monopolist at different times of day!

Interesting further issues arise with more than three players, and with uncertainty and differential information, and there is some hope that more explicit solutions may be obtained than exist for the general dynamic game. Just as the value function and core treated in §§4a, b are not arbitrary, but are derived from multicommodity flow problems, so the dynamic game inherits substantial structure from the network itself.

5. Broadband traffic and statistical sharing

Within a broadband network, the usage of a network resource may not be well assessed by a simple count of the number of bits carried. For example, to provide an acceptable performance to bursty sources with tight delay and loss requirements it may be necessary to keep the average utilization of a link below 10%, while for constant rate sources or sources able to accommodate substantial delays it may be possible to push the average utilization well above 90%.

What is needed is a measure of resource usage which adequately represents the trade-off between sources of different types, taking proper account of their varying characteristics and requirements. In recent years the notion of an *effective bandwidth* (Hui 1988; Gibbens & Hunt 1991; Guérin *et al.* 1991; Kelly 1991a), reviewed in

§§ 5 *a, b*, has been developed to provide such a measure. The effective bandwidth of a source depends sensitively upon the statistical properties of the source as well as its quality of service requirements, and thus the issue becomes how much of the effort of statistical characterization should fall upon the network and how much upon the user responsible for a source.

Within the telecommunications and computer industries it is possible to discern two extreme approaches to this issue. One approach insists that a user provide the network with a full statistical characterization of its traffic source, which is then policed by the network. Another approach stresses the difficulty for a user of providing any information on traffic characteristics, and expects the network to cope nevertheless. These descriptions are, of course, caricatures. Note, though, that both approaches recognize the benefits of statistical sharing: they differ in how much characterization effort is required, and how this effort should be distributed over the combined system comprising users *and* network. The correct balance will necessarily involve trade-offs between the user's uncertainty about traffic characteristics and the network's ability to statistically multiplex connections in an efficient manner. Section 5 *c* describes a charging mechanism that makes these trade-offs explicit, and which encourages the cooperative sharing of information and characterization effort between users and the network.

(a) *Effective bandwidths*

Let $X(t)$ be the amount of work that is produced by a source in a random interval of length t , and let

$$\alpha(z, t) = \frac{1}{zt} \log \mathbb{E}[e^{zX(t)}]. \quad (5.1)$$

As the parameter z increases over the range $(0, \infty)$, the function $\alpha(z, t)$ increases from the mean of $X(t)/t$ to the maximum of $X(t)/t$. We shall call $\alpha(z, t)$ the *effective bandwidth* of the source, although it might more properly be called the *effective bandwidth function*: the free parameters z and t are left undetermined, and their choice at any given resource will depend upon characteristics of the resource such as its capacity, buffer space and scheduling policy.

Some examples may help to illustrate the concept (for generalizations of these examples, and more detailed analysis, see Kelly (1996)).

(i) *Periodic sources*

For a source which produces b units of workload at times $\{Ud + nd, n = 0, 1, \dots\}$, where U is uniformly distributed on the interval $[0, 1]$,

$$\alpha(z, t) = \frac{b}{t} \left\lfloor \frac{t}{d} \right\rfloor + \frac{1}{zt} \log \left[1 + \left(\frac{t}{d} - \left\lfloor \frac{t}{d} \right\rfloor \right) (e^{bz} - 1) \right]. \quad (5.2)$$

This function is plotted in figure 11, with parameters $b = d = 1$.

(ii) *On-off fluid sources*

Consider a stationary fluid source described by a two-state Markov chain. The transition rate from state 1 to state 2 is λ and the transition rate from state 2 to state 1 is μ . While the Markov chain is in state 1 no workload is produced; while it

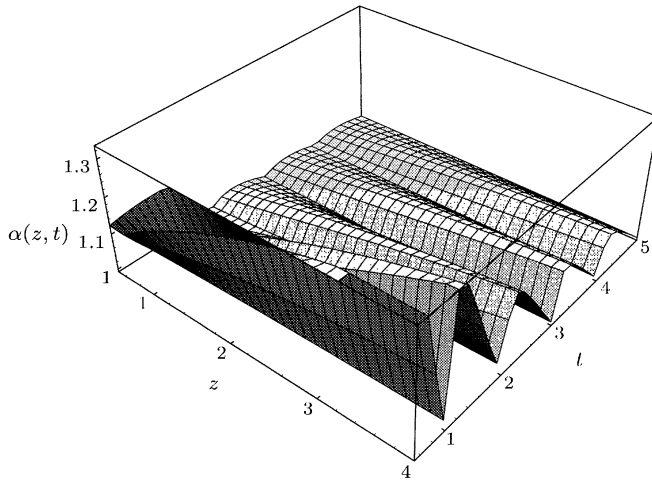


Figure 11. Effective bandwidth of a periodic source. The source produces a single unit of workload at the end of every unit interval, but the phase of the source is random. Note the rapid growth of the effective bandwidth over intervals shorter than the period of the source.

is in state 2 workload is produced at a constant rate h . Then

$$\alpha(z, t) = \frac{1}{zt} \log \left\{ \left(\frac{\mu}{\lambda + \mu}, \frac{\lambda}{\lambda + \mu} \right) \exp \left[\begin{pmatrix} -\lambda & \lambda \\ \mu & -\mu + hz \end{pmatrix} t \right] \begin{pmatrix} 1 \\ 1 \end{pmatrix} \right\},$$

and

$$\lim_{t \rightarrow \infty} \alpha(z, t) = \frac{1}{2z} (hz - \mu - \lambda + ((hz - \mu + \lambda)^2 + 4\lambda\mu)^{1/2}),$$

an expression discussed by Gibbens & Hunt (1991), Guérin *et al.* (1991), and Chang (1994).

(iii) Gaussian sources

Let

$$X[0, t] = Z(t) + \lambda t,$$

where $Z(t)$ is normally distributed with zero mean and arbitrary variance. Then

$$\alpha(z, t) = \lambda + \frac{z}{2t} \text{Var } Z(t).$$

The case $\text{Var } Z(t) = \sigma^2 t$ commonly arises from heavy traffic models (Harrison 1985). The more general case $\text{Var } Z(t) = \sigma^2 t^{2H}$ arises when the process Z is fractional Brownian motion, with Hurst parameter $H \in (0, 1)$, when

$$\alpha(z, t) = \lambda + \frac{1}{2} \sigma^2 z t^{2H-1}. \quad (5.3)$$

For $H > \frac{1}{2}$ the process exhibits long range dependence, and has been proposed and discussed as a model for Ethernet traffic data (Norros 1994; Willinger *et al.* 1995; Willinger 1995). Figure 12 illustrates the form (5.3) with parameters $H = 0.75$, $\lambda = 1$, $\sigma^2 = 0.5$.

(iv) General on/off sources

Suppose next that a source alternates between long periods in an 'on' state, where it behaves as a source with effective bandwidth $\alpha_1(z, t)$, and long periods in an 'off'

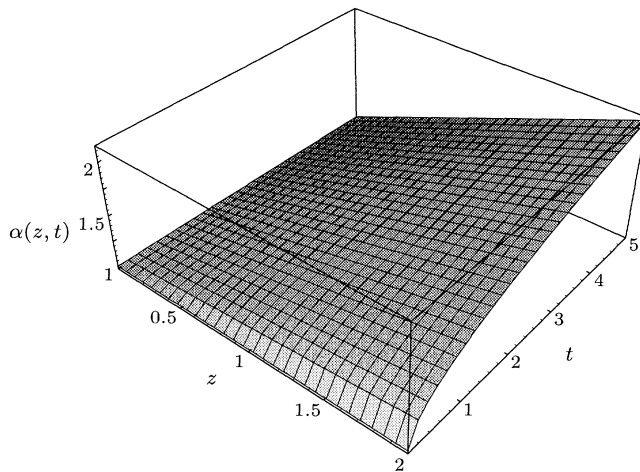


Figure 12. Effective bandwidth of a Gaussian source. This example has Hurst parameter $H = 0.75$: long range order is indicated by the growth of the effective bandwidth with t .

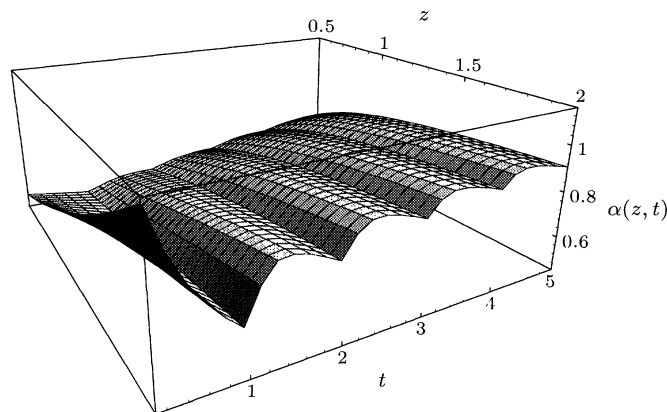


Figure 13. Effective bandwidth of an on/off periodic source. Note the increase of the effective bandwidth as t either decreases below the period of the source, or increases towards the interval over which the source remains 'on' or 'off'.

state, where it produces no workload. Let p be the proportion of time spent in the 'on' state. Then for values of t small compared with periods spent in an 'on' or 'off' state,

$$\alpha(z, t) = \frac{1}{zt} \log[1 + p(\exp(zt\alpha_1(z, t)) - 1)]. \quad (5.4)$$

Figure 13 illustrates this function when $p = 0.4$ and $\alpha_1(z, t)$ is given by expression (5.2). This example shares similarities with examples (i) or (iii) over short or longer time scales respectively.

(b) Multiplexing models

Suppose that J sources share a single resource, and let $X_j[\tau, \tau + t]$ be the load produced by source j over the time period $[\tau, \tau + t]$, where the processes $X_j, j = 1, 2, \dots, J$ are independent and stationary. Can the resource cope with the superposition of the J sources?

For a variety of different resource models, the answer to this question is, by now,

well understood (Botvich & Duffield 1994; Courcoubetis & Weber 1994; De Veciana & Walrand 1994; Hsu & Walrand 1994; Kesidis *et al.* 1993; Key 1995; Simonian & Guibert 1994; Whitt 1993). For several simple models there are constants z^* , t^* and C^* such that if

$$\sum_{j=1}^J \alpha_j(z^*, t^*) \leq C^*, \quad (5.5)$$

then a defined quality of service guarantee, expressed in terms of cell loss or delay, can be assured. For more complex models there may be several constraints of the form (5.5) corresponding to different physical or logical resources within the network. For example, suppose a single resource gives strict priority to sources $j \in J_1$, which have a strict delay requirement, but also serves sources $j \in J_2$ which have a much less stringent delay requirement. Then two constraints of the form

$$\sum_{j \in J_1} \alpha_j(z_1, t_1) \leq C_1, \quad \sum_{j \in J_1 \cup J_2} \alpha_j(z_2, t_2) \leq C_2 \quad (5.6)$$

will generally be needed to ensure that both sets of requirements are met (cf. Bean 1994; Elwalid & Mitra 1994). If the less stringent delay requirement becomes very weak, corresponding to a *very* large buffer and almost *no* sensitivity to delay, then z_2 will approach zero, and $\alpha_j(z_2, t_2)$ will approach $\mathbb{E}(X_j[\tau, \tau + 1])$, the mean load produced by source j . The second constraint of (5.6) then becomes the simple constraint that the mean loads of all sources should not exceed the capacity of the resource, the minimal constraint necessary for the infinite buffer queue to be stable (cf. figure 4, and the discussion of § 2 c (iii)).

In § 2 we have seen how fundamental constraints may be logical rather than physical. Our earliest examples, illustrated in figures 2 and 3, arose from routing flexibility, and the above example arises from scheduling flexibility. Another interesting example can be constructed as follows. Suppose that several independent periodic sources are superimposed. Then the resulting process approximates a Brownian bridge (Roberts 1993, p. 119; Hajek 1994, p. 150). Suppose that all of the sources have the same period d , and that sources of type j share the same block size b_j . Then, under the Brownian bridge approximation, the probability the queue size exceeds a level B is held below $e^{-\gamma}$ if and only if

$$\sum_j n_j b_j \leq Cd \quad (5.7)$$

and

$$\sum_j n_j \left(b_j + b_j^2 \frac{\gamma}{2B} \right) \leq B + Cd. \quad (5.8)$$

The constraint (5.7) is straightforward: no more work should arrive per period than can be dispatched per period. The second constraint (5.8) is again linear, but is now a logical constraint arising from the probabilistic guarantee on buffer overflow.

Under specific modelling assumptions on sources it is often possible to refine constraints such as (5.5)–(5.8) by detailed numerical computations. The effective bandwidths defined by the more refined constraints may no longer have the simple analytical forms of the above examples, but share a qualitatively similar dependence on the statistical properties and performance requirements of the sources.

To develop some understanding of this dependence, let us consider the very simple

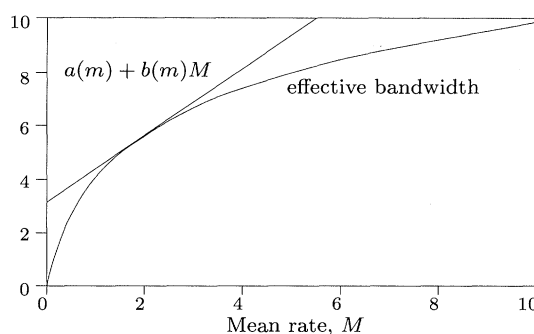


Figure 14. Implicit pricing of an effective bandwidth. The effective bandwidth is shown as a function of the mean rate, M . The user is free to choose the declaration m , denominated in, for example, megabits per second, and is then charged an amount $a(m)$ per second, and an amount $b(m)$ per megabit. The values of $a(m)$ and $b(m)$ are determined from the tangent to the effective bandwidth function at the point $M = m$.

case of an on/off source which produces workload at a constant rate h while in an 'on' state, and produces no workload while in an 'off' state. Suppose the periods spent in 'on' and 'off' states are large, so that the effective bandwidth is given by expression (5.4) with $\alpha_1(z, t) = h$ and $p = m/h$. Here m and h are respectively the mean and peak rate of the source. Let $s = z^*t^*$, and rewrite the effective bandwidth as

$$B(h, m) = \frac{1}{s} \log \left[1 + \frac{m}{h} (e^{sh} - 1) \right], \quad (5.9)$$

where the notation now emphasizes the dependence of the effective bandwidth on the mean and peak rate of the source. Observe that for fixed h this function is increasing and concave in m , while for fixed m it is increasing and convex in h . As $s \rightarrow 0$ (which corresponds, in this example, to a very large resource capacity in relation to the peak h), the effective bandwidth approaches m , the mean rate of the source. However, as s increases (corresponding to a larger peak h in relation to the resource capacity) the effective bandwidth increases to the peak rate h of the source.

(c) Charging mechanisms

The effective bandwidth of a source depends sensitively upon its statistical characteristics. The source, however, may have difficulty providing a complete characterization: for example it may know its peak rate but not its mean rate. Can tariffs be designed that perform the dual role of conveying information to the network that allows more efficient statistical sharing, and feedback to the source about the congestion that it is causing?

One possible charging mechanism might measure the effective bandwidth of a connection, perhaps by estimating expression (5.1) using an empirical averaging to replace the expectation operator. There is, however, a simpler indirect mechanism, which has important additional advantages in the coordination of information and characterization effort between users and the network.

To illustrate the mechanism, consider the case where a call comprises an on/off source with a known (and possibly policed) peak rate h , but with a mean rate that may not be known with certainty, even to the user responsible for the source.

Suppose that, before the call's admission, the network requires the user to announce a value m , and then charges for the call an amount $f(m; M)$ per unit time, where M is the measured mean rate for the call. We suppose that the user is risk-

neutral and attempts to select m so as to minimize the expected cost per unit time: call a minimizing choice of m , \hat{m} say, an *optimal* declaration for the user. What properties would the network like the optimal declaration \hat{m} to have? Well, first of all the network would like to be able to deduce from \hat{m} the user's expectation of M . A second desirable property would be that the expected cost per unit time under the optimal declaration \hat{m} be proportional to the effective bandwidth of the call (or, equivalently, *equal* to the effective bandwidth under a choice of units). In Kelly (1994*b*) it is shown that these two requirements essentially characterize the tariff $f(m; M)$ as

$$f(m; M) = a(m) + b(m)M, \quad (5.10)$$

defined as the tangent to the curve $B(h, M)$ at the point $M = m$ (see figure 14). Note the very simple interpretation possible for the tariff (5.10): the user is free to declare a value m , and then incurs a charge $a(m)$ per unit time, and a charge $b(m)$ per unit of volume carried. (Detailed examples and further discussion may be found in Kelly (1995).)

The properties characterizing the tariff f have many interesting and desirable consequences. For example, suppose that a user can, with some effort, improve its prediction of the statistical properties of a call. A crude method of deciding upon the declaration m might be to take the average of the measured means for the last n calls, but more sophisticated methods are possible. If the user is an organization containing many individuals, the user might observe the identity of the individual making the call, the applications active on that individual's desktop computer, as well as the called party, and utilize elaborate regression aids to make the prediction m . Is it worth the effort? In Kelly (1994*b*) it is shown that improved prediction reduces the expected cost per unit time of the connection by exactly the expected reduction in the effective bandwidth required from the network. This is an important property: users should not be expected to do more work determining the statistical properties of their calls than is justified by the benefit to the network of better characterization.

The overall aim of this work, which is as yet in only its preliminary stages, is to design simple, robust tariff structures, admission controls and routing schemes which will permit an entire system, comprising many networks and huge numbers of users, to function coherently.

6. Conclusion

The topics discussed have, I hope, illustrated the varied and powerful role of mathematical models in the design of communication networks. I hope also to have conveyed something of the theoretical challenge posed to mathematicians, engineers and economists by the practical need to cope with the forthcoming proliferation of interconnecting networks and the increasing variety of telecommunication traffic.

A theme running through the lecture has been the importance of an interdisciplinary view. Traditionally stability and robustness have been considered engineering issues, requiring an analysis of randomness and feedback operating on fast timescales, while the coordination of players has been considered an economic issue, involving discussion of longer-term strategy and of externalities. These distinctions are disappearing, and the insights of at least these two disciplines will be needed to understand how large-scale networks may be designed to function coherently. For example, the approach to charging mechanisms in a broadband network described in § 5 gives an

illustration of how software agents might be provided with feedback about the resource implications of their actions: this approach benefits from both the economist's insights on incentive-compatible tariffs in a stochastic environment, and the control engineer's insights on the hierarchical control through Lagrange multipliers of the combined system comprising users and network. Throughout there are questions that provide impetus for the development of new mathematical and statistical methods.

While the topics discussed have not been that close to electrical science, I suspect that Clifford Paterson, who did so much to establish international standards in his field in the early part of this century, would have been interested in how the scope of standards has extended from physical measures such as luminance or voltage, to information structures, such as the packet length of ATM cells, or formats for the exchange of dynamic routing and charging information.

The Clifford Paterson Lecture was delivered at the Royal Society in London in June 1995, at Heriot-Watt University in Edinburgh in August 1995 as part of a Royal Statistical Society Research Workshop in Stochastic Networks, and at GEC's Hirst Research Centre in Borehamwood in September 1995. I am grateful to all the participants in these gatherings for their very helpful and stimulating comments. During the preparation of the lecture and the paper the author was supported by a Royal Society Leverhulme Trust Senior Research Fellowship. The figures for this paper were prepared by Damon Wischik, while supported by EPSRC grant GR/J31896 and Commission of the European Communities ACTS project AC039. Further information on the lecture and pointers to related information resources will be available for a period on <http://www.statslab.cam.ac.uk/~frank/>.

References

- Ackerley, R. G. 1987 Hysteresis-type behaviour in networks with extensive overflow. *Br. Telecom Technol. J.* **5**, 42–50.
- Ash, G. R. & Huang, B. D. 1994 Comparative evaluation of dynamic routing strategies for a worldwide intelligent network. In Labetoulle & Roberts (1994), pp. 677–686.
- Bean, N. 1994 Effective bandwidths with different quality of service requirements. In *IFIP Transactions, Integrated Broadband Communication Networks and Services* (ed. V. B. Iversen), pp. 241–252. Amsterdam: Elsevier.
- Botvich, D. D. & Duffield, N. 1994 Large deviations, the shape of the loss curve, and economies of scale in large multiplexers. Dublin Institute for Advanced Studies.
- Brockmeyer, E., Halstrom, H. L. & Jensen, A. 1948 *The life and works of A. K. Erlang*. Copenhagen: Academy of Technical Sciences.
- Chang, C.-S. 1994 Stability, queue length, and delay of deterministic and stochastic queueing networks. *IEEE Trans. Automatic Control* **39**, 913–931.
- Courcoubetis, C. & Weber, R. R. 1994 Buffer overflow asymptotics for a switch handling many traffic sources. University of Cambridge.
- Crametz, J.-P. & Hunt, P. J. 1991 A limit result respecting graph structure for a fully connected loss network with alternative routing. *Ann. Appl. Prob.* **1**, 436–444.
- De Veciana, G. & Walrand, J. 1994 Effective bandwidths: call admission, traffic policing and filtering for ATM networks. *Queueing Syst.*
- Elwalid, A. & Mitra, D. 1994 Analysis, approximations and admission control of a multi-service multiplexing system with priorities. AT&T Bell Laboratories.
- Gibbens, R. J. 1988 Dynamic routing in circuit-switched networks: the dynamic alternative routing strategy. Ph.D. thesis, University of Cambridge.
- Gibbens, R. J. & Hunt, P. J. 1991 Effective bandwidths for the multi-type UAS channel. *Queueing Syst.* **9**, 17–28.
- Gibbens, R. J., Hunt, P. J. & Kelly, F. P. 1990 Bistability in communication networks. In *Disorder in physical systems* (ed. G. R. Grimmett & D. J. A. Welsh), pp. 113–128. Oxford University Press.

- Gibbens, R. J. & Kelly, F. P. 1990 Dynamic routing in fully connected networks. *IMA JI Math. Cont. Information* **7**, 77–111.
- Gibbens, R. J., Kelly, F. P., Cope, G. A. & Whitehead, M. J. 1991 Coalitions in the international network. In *Teletraffic and Data Traffic in a Period of Change: Proc. 13th Int. Teletraffic Congress* (ed. Jensen & Iversen), pp. 93–98. Amsterdam: Elsevier.
- Gibbens, R. J., Kelly, F. P. & Key, P. B. 1988 Dynamic alternative routing – modelling and behaviour. In *Proc. 12th Int. Teletraffic Congress, Turin* (ed. M. Bonatti). Amsterdam: Elsevier.
- Gibbens, R. J., Kelly, F. P. & Key, P. 1995a Dynamic alternative routing. In Steenstrup (1995), pp. 13–47.
- Gibbens, R. J., Kelly, F. P. & Turner, S. R. E. 1993 Dynamic routing in multi-parented networks. *IEEE/ACM Trans. Networking* **1**, 261–270.
- Graham, C. & Méléard, S. 1993 Propagation of chaos for a fully connected loss network with alternate routing. *Stoch. Proc. Appl.* **44**, 159–180.
- Guérin, R., Ahmadi, H. & Naghshineh, M. 1991 Equivalent capacity and its application to bandwidth allocation in high-speed networks. *IEEE JI Selected Areas Commun.* **9**, 968–981.
- Hajek, B. 1987 Average case analysis of greedy algorithms for Kelly's triangle problem and the independent set problem. In *26th IEEE Conference on Decision and Control*. New York: IEEE.
- Hajek, B. 1994 A queue with periodic arrivals and constant service rate. In *Probability, statistics and optimization* (ed. F. P. Kelly), pp. 147–157. Chichester: Wiley.
- Harrison, J. M. 1985 *Brownian motion and stochastic flow systems*. Kreiger.
- Houck, D. J., Meier-Hellstern, K. S., Saheban, F. & Skoog, R. A. 1994 Failure and congestion propagation through signalling controls. In Labetoulle & Roberts (1994), pp. 367–385.
- Hsu, I. & Walrand, J. 1995 Admission control for ATM networks. In Kelly & Williams (1995), pp. 411–427.
- Hui, J. Y. 1988 Resource allocation for broadband networks. *IEEE JI Selected Areas Commun.* **6**, 1598–1608.
- Hunt, P. J. & Laws, C. N. 1992 Least busy alternative routing in queueing and loss networks. *Prob. Engng Inf. Sci.* **6**, 439–456.
- Hunt, P. J. & Laws, C. N. 1993 Asymptotically optimal loss network control. *Math. Oper. Res.* **18**, 880–900.
- Jones, A. J. 1980 *Game theory: mathematical models of conflict*. Chichester: Wiley.
- Kelly, F. P. 1985 Stochastic models of computer communication systems. *J. R. Statist. Soc. B* **47**, 379–395.
- Kelly, F. P. 1990 Routing and capacity allocation in networks with trunk reservation in networks with trunk reservation. *Math. Oper. Res.* **15**, 771–793.
- Kelly, F. P. 1991a Effective bandwidths at multi-class queues. *Queueing Syst.* **9**, 5–16.
- Kelly, F. P. 1991b Loss networks. *Ann. Appl. Prob.* **1**, 319–378.
- Kelly, F. P. 1991c Network routing. *Phil. Trans. R. Soc. Lond. A* **337**, 343–367.
- Kelly, F. P. 1994a Bounds on the performance of dynamic routing schemes for highly connected networks. *Math. Oper. Res.* **19**, 1–20.
- Kelly, F. P. 1994b On tariffs, policing and admission control of multiservice networks. *Oper. Res. Lett.* **15**, 1–9.
- Kelly, F. P. 1995 Charging and accounting for bursty connections. In *Internet economics* (ed. J. Bailey & L. McKnight). MIT Press.
- Kelly, F. P. 1996 Notes on effective bandwidths. In *Stochastic networks: theory and applications* (ed. S. Zachary, I. B. Ziedins & F. P. Kelly). Oxford University Press.
- Kelly, F. P. & Laws, C. N. 1993 Dynamic routing in open queueing networks: Brownian models, cut constraints and resource pooling. *Queueing Syst.* **13**, 47–86.
- Kelly, F. P. & Williams, R. J. (eds) 1995 *Stochastic networks*. New York: Springer.
- Kesidis, G., Walrand, J. & Chang, C. S. 1993 Effective bandwidths for multiclass Markov fluids and other ATM sources. *IEEE/ACM Trans. Networking* **1**, 424–428.

- Key, P. B. 1988 Implied cost methodology and software tools for a fully connected network with DAR and trunk reservation. *British Telecom Technol. JI* **6**, 52–66.
- Key, P. 1990 Optimal control and trunk reservation in loss networks. *Prob. Engng Inf. Sci.* **4**, 203–242.
- Key, P. 1995 Connection admission control for ATM networks. *British Telecom Technol. JI* **13**, 52–66.
- Key, P. B. & Whitehead, M. J. 1988 Cost-effective use of networks employing Dynamic Alternative Routing. In *Proc. 12th Int. Teletraffic Congress, Turin* (ed. M. Bonatti). Amsterdam: Elsevier.
- Labetoulle, J. & Roberts, J. W. (eds) 1994 *The Fundamental Role of Teletraffic in the Evolution of Telecommunication Networks: Proc. 14th Int. Teletraffic Congress*. Amsterdam: Elsevier.
- Laws, C. N. 1992 Resource pooling in queueing networks with dynamic routing. *Adv. Appl. Prob.* **24**, 699–726.
- Louth, G. M., Mitzenmacher, M. & Kelly, F. P. 1994 Computational complexity of loss networks. *Theor. Computer Sci.* **125**, 45–59.
- Mees, A. I. 1986 Simple is the best for dynamic routing of telecommunications. *Nature, Lond.* **323**, 108.
- Mitra, D. (ed.) 1995 Advances in the fundamentals of networking. *IEEE JI Selected Areas Commun.* **13**, 933–1362.
- Mitra, D. & Seery, J. B. 1991 Comparative evaluations of randomized and dynamic routing strategies for circuit-switched networks. *IEEE Trans. Commun.* **39**, 102–116.
- Nakagome, Y. & Mori, H. 1973 Flexible routing in the global communication network. In *Proc. 7th Int. Teletraffic Congress*. Amsterdam: North-Holland.
- Norros, I. 1994 A storage model with self-similar input. *Queueing Syst.* **16**, 387–396.
- Roberts, J. (ed.) 1992 *Performance evaluation and design of multiservice networks*. Luxembourg: Office for Official Publications of the European Communities.
- Ross, K. W. 1995 *Multiservice loss models for broadband telecommunications networks*. London: Springer.
- Ryde, J. W. 1949 Clifford Copland Paterson. *Obituary Notices of Fellows of the Royal Society* **6**, 479–501.
- Shubik, M. 1982 *Game theory in the social sciences*. MIT Press.
- Simonian, A. & Guibert, J. 1994 Large deviations approximation for fluid sources fed by a large number of on/off sources. In Labetoulle & Roberts (1994), pp. 1013–1022.
- Schmitt, J. A. & Watanabe, Y. 1991 Models and results for planning the evolution to worldwide intelligent network dynamic routing. In Labetoulle & Roberts (1994), pp. 687–698.
- Songhurst, D. J. 1980 Protection against traffic overload in hierarchical networks employing alternative routing. In *Telecommunications Network Planning Symp., Paris*.
- Stacey, R. R. & Songhurst, D. J. 1987 Dynamic Alternative Routing in the British Telecom trunk network. *Int. Switching Symp., Phoenix, Arizona*.
- Steenstrup, M. (ed.) 1995 *Routing in communication networks*. New Jersey: Prentice-Hall.
- van Golstein Brouwers, W. G. 1994 Cost reductions by coalitions in the international network. In Labetoulle & Roberts (1994), pp. 699–708.
- Varian, H. R. 1992 *Microeconomic analysis*, 3rd edn. New York: Norton.
- Whitt, W. 1993 Tail probabilities with statistical multiplexing and effective bandwidths in multi-class queues. *Telecommun. Syst.* **2**, 71–167.
- Willinger, W. 1995 Traffic modelling for high-speed networks: theory versus practice. In Kelly & Williams (1995), pp. 395–409.
- Willinger, W., Taqqu, M. S., Leland, W. E. & Wilson, D. V. 1995 Self-similarity in high-speed packet traffic: analysis and modelling of ethernet traffic measurements. *Statist. Sci.* **10**, 67–85.
- Wroe, G. A., Cope, G. A. & Whitehead, M. J. 1990 Flexible routing in the BT international access network. In *Proc. 7th UK Teletraffic Symp.* London: IEE.

Lecture given 14 June 1995