PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A
MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Control and pricing for communication networks

Richard Gibbens

Phil. Trans. R. Soc. Lond. A 2000 **358**, 331-341
doi: 10.1098/rsta.2000.0534

THE ROYAL SOCIETY

Control and pricing for communication networks

By Richard Gibbens

Statistical Laboratory, University of Cambridge, 16 Mill Lane, Cambridge CB2 1SB, UK (r.j.gibbens@statslab.cam.ac.uk)

This paper surveys the field of congestion control in networks. It combines an historical perspective, looking in particular at the Internet, with an account of recent research work on network resource pricing. Finally, the paper concludes by speculating on how control and pricing in communication networks might lead to the efficient and fair allocation of scarce resources in tomorrow's Internet.

Keywords: Internet; rate control; shadow prices; charging; communication networks

1. Introduction

The rapid pace of developments in communication networks has produced outstanding advances. Two parallel developments have been taking place. Firstly, significant advances in the engineering of communication equipment, that is the switches, routers, optical fibres and all the software that coordinates and controls them, has led to a huge expansion in the bandwidths available in transmission networks. Secondly, developments in the computing technology used by the end systems have brought about a keen interest in multimedia applications as well as distributed computing; the growth of the World Wide Web in the Internet being amongst the most familiar example of these developments.

Furthermore, such developments have started to produce major changes in the structure of many industries and in the way in which commerce is taking place. These changes will continue as the cost of accessing all kinds of information continues to tumble.

All these exciting developments result from research and developments across a diverse range of fields, including mathematics, engineering, computer science and economics. This survey article, written for a general scientific audience, presents some examples of these recent advances and then develops a theme that has recently started to emerge (Gibbens & Kelly 1999*a*, *b*), which has the potential to transform the way in which we think about building communication networks in the future.

This latter approach is based on a view that the network need only convey feedback signals to users, perhaps in the form of marks attached to users' packets, so as to indicate the cost, measured in terms of resources consumed within the network. If a small charge is then incurred for each mark received, users would have the incentive and the necessary information to adapt their traffic demands to share the available resources efficiently and fairly. Speculating about the future, the article describes ways in which this approach, if adopted, would lead to simple and robust mechanisms for adding more elaborate notions of quality of service. Examples will be given in which it is possible to synthesize higher level service models associated with real-time services, such as telephony and video, from simple underlying packet networks.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

2. Current Internet

We begin our account of the current Internet by summarizing its basic components and properties. We shall then look at how it has been used and what mechanisms have been developed to support these uses. Any discussion of the development of the Internet will also need to consider its funding and the means by which revenue is extracted from the users. This section will then conclude with some of the fundamental difficulties that need to be overcome before further growth in the use of the Internet can succeed.

(a) *The basics of packet-switched networks*

The Internet is an example of a *packet-switched* network. The data to be transferred between two parties connected to such a network are not sent in one whole unit but are instead chopped up into smaller individual units called *packets*. Each separate packet is labelled with the source and destination addresses and is then dispatched between the end parties. Packets are transferred between the end parties along communication channels between intermediate network components (called *routers*) that perform routing (or switching) functions and which multiplex together slower individual communication channels onto faster channels. In this way, packets traverse a network between any two end parties. All the routing functionality is performed packet by packet using the addressing information that has been attached to the data within each packet. At the routers, packets may be delayed in buffers before they can be forwarded along the next stage. When there are too many packets for the routers to buffer, packets are dropped and do not succeed in reaching their destination.

This packet-switched approach adds several types of overhead. The first is the addressing information that is added to each of the packets and the second is the additional processing load that is required at the routing network components to forward each packet. Nevertheless, this packet-switched design emerged for the Internet since it is very simple and robust. The primary aims of the design were to provide *connectivity* and *resilience* to channel or router failures. Issues of *efficiency* and *qualities of service* (as measured by packet delays or packet drops) were of secondary importance to the design.

This idea of sending packets into the network at one point and receiving at another point those that were not lost at intermediate congested components along the way is the lowest level functionality of the Internet. This *best-effort* approach had the simplicity to enable many forms of communications to be built upon it. This basic functionality is provided by the protocol known as the *universal datagram protocol* (UDP).

One development was to provide a means for the end parties to monitor whether any packets had been lost and re-send them accordingly. This packet-loss detection was implemented by the sending party attaching sequence numbers along with the addressing information to the packets within a given connection. The receiving party was then required to send back packets of its own acknowledging the receipt of packets, where each packet could be identified by means of its sequence number. The lack of an acknowledgement by the sender within a certain timeout period was then deemed to signify that a particular packet had been lost and should

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

therefore be re-sent. The net effect was to provide an error-free end-to-end communication channel. This is the essence of how the *transmission control protocol* (TCP) operates.

Several points are worth noting. Firstly, the TCP has to establish a connection between the two end parties *before* any actual data packets are transferred (a signalling protocol exchanging information is used to *open* a connection, and later, when the connection is no longer required, to *close* a connection). Secondly, the TCP is entirely implemented within the end parties' networking software. In particular, it requires no support from the network routers either in the form of additional packet processing or in terms of storing state information about the fate of packets according to the many connections that may pass through any given router.

(*b*) *Flow control*

Early applications of the Internet protocols included the exchange of email and files using protocols, known as the *simple mail transfer protocol* (SMTP) and *file transfer protocol* (FTP), respectively, built on top of TCP connections. Congestion was soon a noticeable problem and at the heart of this issue lay the problem of the sending parties needing to be informed about the rate they should be sending in order to share the resources without needing to re-send too many packets. A flow-control strategy was devised for the TCP (Jacobson 1998), which used the discovery of dropped packets to trigger the sender to reduce its sending rate. The sending party then slowly increased its sending rate again until further packet losses caused it to reduce and the cycle of slow increase and rapid decrease repeated. In this way, the sending parties, which share common congested network resources, had a mechanism to adapt their sending rates to balance the conflict between wishing to achieve high throughput of packets as well as low packet-loss rates.

Additionally, a mechanism is provided at the start of a connection to allow the sender to *rapidly* increase its sending rate prior to beginning the cyclic phase of slowly increasing and rapidly decreasing; strangely, this became known as the *slow-start* mechanism.

Again, it is worth noting that this flow-control mechanism could be implemented entirely within the end parties and did not require any additional functionality to be added to the network routers.

(*c*) *Applications*

Many other applications have since been developed that operate using the Internet to exchange information. The most familiar application today is the World Wide Web, which provides for distributed access to information that can be displayed by means of browser software. The browser software running on a client end system makes a TCP connection to a server process running on another end system and thereby exchanges the information to render a document on the browser's display. Often, a document has embedded references to other documents that are retrieved by the client in the same fashion. Thus, a single request to display a page may produce multiple TCP connections to be opened with an arbitrary collection of servers. This simple architecture has proved highly effective at enabling producers and consumers of rich multi-media information to flourish.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Another application that has received much attention recently is that of Internet telephony (also referred to as *voice over IP*). This application sends and receives packets between two end parties that encode speech information. For this to provide real-time speech between two humans, the packet-sending rates should not be too low otherwise the conversation breaks up. Thus, this application is usually implemented not by using TCP connections with their associated flow-control mechanisms, but simply by means of a UDP connection. The senders do *not* reduce their rates to share the other users' demands on the common, scarce, network resources.

Much of the attention on Internet telephony has focused around cost and quality issues. Since Internet telephony calls avoid parts of the public switched telephony network, they also avoid their charges. However, the quality of the connections rarely matches that of the public switched telephony networks due to the impact of the packet-loss rates induced by the congestion levels, which are, to some extent, self-caused by their own lack of rate adaption. It can be argued that to some extent the users of Internet telephony adapt over longer time-scales than those within TCP, in the sense that they will defer making calls until the congestion levels are within acceptable bounds.

(*d*) *Pricing*

In discussions about the Internet, the question of who pays for it often occurs. Attempts to answer this question typically lead to unconvincing explanations at best, and, more usually, to quite a lot of confusion!

Several features can be described (Walker *et al.* 1997). Broadly, there are two types of users of the Internet. The first type are the residential users with access to the Internet through *Internet service providers* (ISPs) over the regular telephone (or cable TV) networks. Such users connect through a modem, which necessarily limits their access bandwidth. Access speeds of 56 kbit s^{-1} are now commonly available. The second type of user is more usually found in businesses or academic environments, where a local-area network of computers interconnects with the Internet. Normally, though not necessarily, this interconnection is at a higher bandwidth than that for the residential customers. Although the interconnection is again provided by an ISP, the contract is with the central management for a group of users and not normally with the individual users themselves.

The standard form of pricing seen by residential customers is a monthly subscription charge (a typical charge being £10 or less) to the ISP together with any telephony charges associated with the dial-up connection. In the UK, the dial-up connection is normally charged at the local call rate. In the US, local calls are normally free. Thus, in the UK there is a usage-based component to the user's charge (dependent on the *time* spent connected though not the actual *volume* of data transferred), whereas in the US, there would not normally be any usage-based component to the charge. The phenomenon whereby US customers dial-up and stay connected through a local telephone circuit for periods of days at a time is not unknown and has caused much concern to US telephone operating companies.

Recently in the UK, several ISPs have dropped the subscription part of the charge and just rely on receiving a proportion of the local call rate deriving from interconnection payments between the telecom operators.

(*e*) *The JANET example*

The UK academic network (known as JANET) provides a very interesting and topical case study. Since the adoption of IP protocols at the end of the 1980s, the traffic volumes have grown at a high rate (estimates suggest a trebling each year). Between the various UK academic institutions there has been relatively little difficulty in keeping pace with this growth. The bottlenecks have been the much more expensive transatlantic connections. Here there have been serious difficulties in supplying sufficient bandwidth to meet the demand for resources. Approaches to these problems not only provide much-appreciated relief to UK academic users (of the World Wide Web especially), but have also focused attention on research into the growing multidisciplinary field of Internet economics, which lies at the heart of this article.

One approach to the problems of handling huge numbers of accesses to Web pages has been the important development of Web caching systems and services. Another approach has been the introduction of a usage-based charge for traffic in-bound to the UK academic network. This charging began in August 1998 and consists of a *volume* charge of 2 p per Mb of data transferred (except during a low-charge period for several hours at night when it is currently free). The aggregated charge is billed to separate institutions and for an initial transition period it is centrally subsidized, reducing the charge from 2 p to 1 p per Mb. These levels were set in order to recover the shortfall between fixed funding levels and the total cost of providing the necessary bandwidth. (Note that the unit cost of providing the bandwidth is falling, but not at a sufficient rate to compensate within a constant budget for the growth in traffic levels.) It is currently too early to report other than anecdotal effects arising from these charging mechanisms. Certainly, institutions have begun to monitor more closely their use of network facilities, but, as yet, the majority of standard users seem at best only dimly aware that such charging is taking place. This could easily change very quickly and it will be fascinating to see the variety of responses taken; both by institutions and by users!

(*f*) *Difficulties*

There are two major obstacles that are causing fundamental concerns about how well the current Internet will scale up in the future. One concern is the increasing degree of diversity, or heterogeneity, in the applications being used over the Internet. It is no longer the case that the majority of the traffic is confined to just a few well-understood types, such as telnet, FTP or SMTP traffic. Now, we need to add Web traffic, Internet telephony traffic, and many more types that few have any doubt will be invented and just as rapidly adopted. The second concern is that along with this diversity at the traffic level, there is also increasing diversity in value, or utility, terms. Not every packet transferred is equally valued. Some packets in an Internet telephony connection may well be lost without the users noticing, whereas lost data in a file transfer would normally trigger data to be re-sent. The value attached to the data does not just depend on the application that generated them, but it will also increasingly depend on the context and the human user's perceptions. Consider how a frivolous browse of a Web page by one person may correspond for another person to a much desired purchasing decision in some e-commerce transaction.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

These two concerns, of increasing diversity in traffic types and user valuations for quality, lie at the heart of much of the research into the way the Internet will evolve. For further discussion on the future design of the Internet, see Shenker (1995) and Shenker *et al.* (1996).

3. IETF proposals

The *Internet Engineering Task Force* (IETF) is the official body that coordinates research and development for the Internet. Two of its activities are directed at tackling some of the problems raised in the previous section.

The proposals under the umbrella of *integrated services* concern the introduction of different service classes describing a user's connection. The different service classes would then, potentially, be handled in different ways appropriate to their characteristics. One class would correspond to the existing best-effort behaviour, whereas another would provide statistical guarantees on the quality of service seen by the connections within that class.

A second approach under investigation by the IETF is known as *differentiated services* (Clark 1996). In this approach, the packets sent across the Internet are labelled or tagged with some priority information. This might take the form of just one or two extra bits of information added to each packet header in order to describe whether the network should regard this packet as either high or low priority (or somewhere in between). High-priority packets might be processed through queues in the routers that take precedence over the queues for low-priority packets. In this way, real-time voice communications would not see the congestion that low-priority email or Web traffic would otherwise cause.

Both of these proposals strive to solve the difficulties of allocating scarce resources between competing heterogeneous users with diverse notions of quality.

4. Pricing proposals

We now turn to some of the pricing proposals that have recently been made within the emerging field of Internet economics.

(*a*) *Smart market*

Seminal work by MacKie-Mason & Varian (1995) describe a *smart-market* approach to allocating scarce network resources to the users who value them the most.

The approach can briefly be described as follows. Each user adds a value attribute to their packets that specifies the user's willingness to pay for the transportation of this individual packet through the network. At individual resources within the network the various competing packets are sorted according to their willingness-to-pay attribute. The resource then drops all those low-valued packets that lie beyond its capacity threshold. In order to deter users from simply attaching higher and higher willingness-to-pay values to their packets, each user is charged when their packets are carried by the resource. The amount they are charged is not their actual willingness to pay but the willingness to pay of the highest-valued packet dropped. The economic language to describe this procedure is to say that the resource conducts an *auction* (Vickrey 1961) for the available bandwidth. The choice of the charge that

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

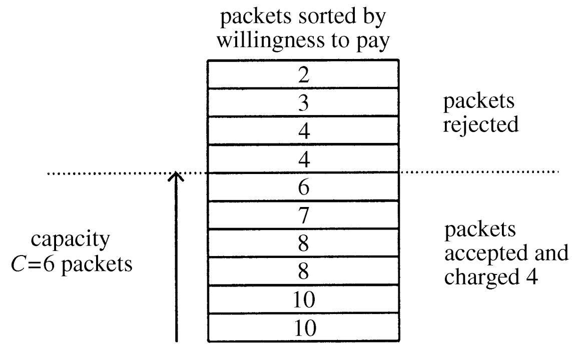

Figure 1. Smart market. This figure shows ten packets simultaneously offered to a resource with a capacity for six packets. The willingness to pay ranges from 2 to 10 units. Packets valued 6 and above are accepted while those valued 4 or below are rejected. The accepted packets are all charged an amount of 4 units.

each user is required to pay when their packet is carried has been shown to be such as to provide an incentive for users to choose their true willingness to pay when they send their packets. Figure 1 shows an example with ten packets offered to a network resource whose capacity is only six packets.

This *smart-market* approach, which uses the tools of economic theory for its definition, has been a source of much inspiration to the engineering community wrestling with the interrelated objectives of achieving network efficiency together with ease of implementation. Conducting an auction on a packet-by-packet basis within a router (observe that this requires a computationally expensive sorting operation) is not considered viable to implement as such. Its importance lies in describing the goal that more easily implemented approximate schemes must seek to achieve.

(*b*) *Paris Metro pricing*

One very pragmatic proposal is made by Odlyzko (1997), and it resembles the pricing used at one time on the Paris Metro. The proposal is to segregate networks into two subnetworks and to charge two different prices. Capacities being equal, the users who are willing to pay more for better quality would choose the high-price network, where bandwidth should be less scarce (and, hence, connections will be less congested), since the remaining users with lower willingness to pay are deterred from entry. Notice that, just like the first-class carriages on the Paris Metro, there is no technical difference between the subnetworks. Improved quality is achieved by the natural behaviour of the self-interested users and not by any clever or complex actions on the part of network resources to respect priority marks on the packets or service classes.

This proposal has many merits and is likely to be the source of many further investigations to determine its eventual role. One such investigation by Gibbens *et al.* (1998) considers the way in which schemes of this type would be affected by competing network suppliers. Would both the competing networks wish to operate a Paris Metro style of pricing or would they collapse to head-to-head competition on price alone? The initial conclusions from these studies suggest that Paris Metro pricing would not be expected to emerge in a competitive market situation. How-

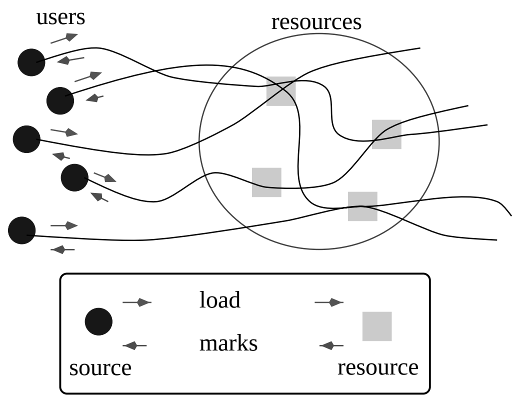

Figure 2. Packet marking. This figure shows the overall system architecture of users and resources. The users produce packets that act as load for resources along a route to the destination node. Each resource receives the aggregated load from all users whose routes traverse the given resource and selectively mark packets in order to indicate the level of congestion (shadow price). The marks are returned to the users who adjust their packet sending rate in accordance with their own interests.

ever, much work remains to be done in this area to determine to what extent these conclusions remain under less-severe modelling assumptions.

(*c*) *Packet marking*

The final pricing proposal considered in this article results from the work of Kelly and co-workers (Kelly 1997; Kelly *et al*. 1998; Gibbens & Kelly 1999*a*, *b*). In this proposal, users adapt their sending rates much as they do in the existing TCP algorithms of Jacobson. However, there are two differences to the existing framework. The first difference is that the network provides a more refined notion of congestion than simply that of packet drops. The resource feeds back a measure of the true congestion: a quantity that mathematicians and economists call the *shadow price* for the resource. The second difference is that the users no longer need to behave identically when faced with congestion information. Instead, they are not just allowed but indeed *encouraged* to follow their own interests in the way that they adapt the rate of sending packets into the network. Using the principle that the users are charged the shadow price for congestion, the users will have both the knowledge and the incentive to jointly vary their behaviour to use the network most efficiently. The overall system architecture of users and resources is shown in figure 2.

Thus the network has an additional job to do and the users have greater freedom in their choice of actions. An important contribution of the research on this topic has been to determine answers to how difficult or complex these jobs are for the designers of network routers and for the software developers of the future seeking to provide integrated multi-media applications at lowest (network) cost.

For the network resources, their job can be accomplished by setting marking information on the packet that is conveyed back using the same acknowledgement procedure that TCP uses to indicate receipt of packets. The average rate of packet marking would then signify the shadow price. Research into the appropriate algorithms for choosing which packets to mark and which packets to leave unmarked is still continuing but early work shows that very simple strategies on the part of the routers, requiring minor extra overheads (certainly as compared with, say, packet sorting), can produce adequate descriptions of the shadow prices.

User strategies have also been the subject of much investigation. Research has shown that when combined with shadow-price information it is possible to make minor modifications to the TCP algorithms that lead to dramatically different allocations of resources between user connections. In particular, users can be responsive to the congestion levels in a way that leads to the resources being allocated according to the expressed willingness to pay, whatever those values actually happen to be. Thus, a truly differentiated services network (Crowcroft & Oechslin 1997) can be formed without the network resources having to understand or interpret any classifications made by the users (or network designers on their behalf).

5. Future developments

Using the packet-marking framework of the previous section, it is possible to build more sophisticated notions of quality than simply that of the packet-drop rate. An example is provided by Gibbens & Kelly (1999*a*), who show how it would be possible to develop a gateway mechanism that effectively translates between the world of packets and marks (which incur a fixed, possibly notional, charge per mark) and the world of non-adaptive sources (such as the current Internet telephony users, who prefer to be told whether their entire connections are accepted or blocked). The gateway could be an extra component that monitors the level of congestion within the network (by means of observing the process of marks flowing back from the resources) and takes decisions concerning the acceptance or rejection of requests from Internet telephony users. It functions as a form of risk taker on behalf of the users.

The gateway is really a virtual device and need not be constructed as a separate component in hardware but can, instead, be implemented as a distributed software process running at a variety of convenient locations.

The approach outlined in this paper describes a public network together with a system for micro-payments. It is worth remarking that, within a private network of cooperative users, the feedback of shadow-price information to the end systems using packet marking may alone be sufficient to allow the efficient operation of the network resources.

In a public network of non-cooperative users, the addition of micro-payments would appear to be essential to provide the necessary incentives. There has been some speculation (Kirkby 1997) that once such a payment system were in place, then ISPs would have a natural means to bill for a range of further services, including content, and a whole system for e-commerce would follow.

Concern about the potential transaction costs associated with micro-payments has led to interesting suggestions for probabilistic payment methods (Wheeler 1996) for use in the case in which there is a large number of very small transactions.

It is hard to speculate with too much certainty about future Internet developments, but it would seem reasonable to suggest that there will be greater opportunities for advancement if the economic and teletraffic problems are resolved via a system of congestion pricing that is built from simple mechanisms. This would ensure the efficiency, robustness and strength with which to support our communication needs for the start of the new millennium and beyond.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

References

Clark, D. D. 1996 Adding service discrimination to the Internet. *Telecomm. Policy* **20**. http://ana-www.lcs.mit.edu/anaweb/abstracts/TPRC20.html.

Crowcroft, J. & Oechslin, P. 1997 Differentiated end to end Internet services using a weighted proportionally fair sharing TCP. ftp://cs.ucl.ac.uk/darpa/multcp.ps.

Gibbens, R. J. & Kelly, F. P. 1999*a* Distributed connection acceptance control for a connectionless network. In *Proc. 16th Int. Teletraffic Congress, Edinburgh, June 1999*, pp. 941–952. http://www.statslab.cam.ac.uk/~frank/dcac.html.

Gibbens, R. J. & Kelly, F. P. 1999*b* Resource pricing and the evolution of congestion control. *Automatica* **35**, 1965–1985. http://www.statslab.cam.ac.uk/~frank/evol.html.

Gibbens, R., Mason, R. & Steinberg, R. 1998 Multiproduct competition between congestible networks. University of Southampton Discussion Paper on Economics and Econometrics, no. 9816, September 1998.

Jacobson, V. 1998 Congestion avoidance and control. In *Proc. ACM SIGCOMM '88*, pp. 314–329. (An updated version is available via ftp://ftp.ee.lbl.gov/papers/congavoid.ps.Z.)

Kelly, F. P. 1997 Charging and rate control for elastic traffic. *Euro. Trans. Telecomm.* **8**, 33–37. http://www.statslab.cam.ac.uk/~frank/elastic.html.

Kelly, F. P., Maulloo, A. K. & Tan, D. K. H. 1998 Rate control in communiction networks: shadow prices, proportional fairness and stability. *J. Oper. Res. Soc.* **49**, 237–252. http://www.statslab.cam.ac.uk/~frank/rate.html.

Kirkby, P. A. 1997 Business models and system architectures for future QoS guaranteed Internet services. In *IEE Coll. 'Charging for ATM—the reality arrives', November 1997*. London: Savoy Place.

MacKie-Mason, J. K. & Varian, H. 1995 Pricing the Internet. In *Public access to the Internet*. Prentice Hall.

Odlyzko, A. 1997 A modest proposal for preventing Internet congestion. AT&T research, http://www.research.att.com/~amo/doc/recent.html.

Shenker, S. 1995 Fundamental design issues for the future Internet. *IEEE J. Selected Areas Comm.* **13**, 1176–1188.

Shenker, S., Clark, D., Estrin, D. & Herzog, S. 1996 Pricing in computer networks: reshaping the research agenda. *Telecomm. Policy* **20**, 183–201.

Vickrey, W. 1961 Counterspeculation, auctions and competitive sealed tenders. *J. Finance* **16**, 8–37.

Walker, D., Kelly, F. P. & Solomon, J. 1997 Tariffing in the new IP/ATM environment. *Telecomm. Policy* **21**, 283–295.

Wheeler, D. 1996 Tranactions using bets. ftp.cl.cam.ac.uk/users/djw3/tub.ps.

AUTHOR PROFILE

R. Gibbens

Richard Gibbens was born in 1962 in London. He studied mathematics at the University of Cambridge, graduating with an MA in mathematics, a diploma in mathematical statistics, and a PhD in 1988 in the Statistical Laboratory. His doctoral thesis work, a collaboration between the University of Cambridge and British Telecom, was on the design and analysis of dynamic routing strategies in telecommunications networks, resulting in the development of the dynamic alternative routing strategy, which is now in operation in the British Telecom trunk network.

He has worked in the area of modelling of communications networks mainly at the Statistical Laboratory, University of Cambridge, but was also a visiting consultant at AT&T Bell laboratories, Murray Hill, NJ, during 1989. He was appointed to a Royal Society University Research Fellowship in 1993.