Eric M. Scharf

Meeting the Challenge of Charging for ATM

ATM provides a flexibility that is considerably greater than that to which users of traditional telephony services are accustomed. This paper will endeavour to establish a framework within which ATM charging schemes may be devised and offered to the customer. We shall look at some of the many candidates for charging parameters and will suggest how they may be grouped in order to assist the development of charging schemes and the specification of suitable facilities in ATM switching fabrics. A charging scheme structure comprising parameters, algorithms and price plans will be suggested. This paper will also look at how one may assess the merits of a particular charging scheme. Some proposed schemes will be reviewed in the context of the suggested framework and the assessment criteria. The conclusions to the paper will also provide a brief forward look to the directions that ATM charging may be taking.

Introduction

We are now entering an era in which many trial ATM networks are being established, and some ATM networks are starting to come into revenue earning service. The operational flexibility offered by ATM networks provides a challenge of new dimensions to the network operators and service providers looking for suitable charging schemes to reflect costs and to match the wide variety of services that ATM can support. A further challenge is to create charging schemes that can accommodate this new-found flexibility while at the same time being customer and user-friendly and easy to implement.

Charging is a customer-oriented function of communications networks. The charging process for a communications service may be influenced by at least three factors: marketing, political and technical. Marketing decisions will often take account of the perceived value of a service to the user or the need to attract users to new services. Political factors are usually associated with regulatory constraints. The cost to the service or network provider of an offered service involves traffic measurement and is thus intimately bound up with the particular communications technology being used; this aspect is the focus of this paper.

Effective charging algorithms are required to support the accounting function within the overlying telecommunications management network (TMN) and the study of such algorithms has thus been supported within the user-oriented ACTS programme by projects such as CANCAN[1] and CASHMAN[2].

Eric M. Scharf: Department of Electronic Engineering, Queen Mary and Westfield College
Mile End Road, London E1 4NS, United Kingdom
e.m.scharf@elec.qmw.ac.uk

The Value Chain

The wealth of possibilities for charging for ATM-based services calls for a structure within which such schemes may be developed and appraised. After a look at the value chain, we shall then take a closer look at how an ATM charging scheme may be structured. This will be followed by a discussion of some typical charging parameters and of how they may be classified.

In the example in Figure 1, we can consider four groupings: an organisation, a service provider, a network operator and an infrastructure provider. Of particular interest here is the payment links between the organisation and its service provider and between the service provider and the network operator.

The organisation could be a factory producing some commodity. Staff (the users) within the organisation use telephones and fax machines. The finance department of the organisation is the customer of the service supplier whose operating section it pays. If the 'organisation' is a domestic consumer, both customer and user may be the same person, or, the customers may be the parents, and the users their children! In this example, the customer may deal with the network operator directly, without going through a service provider.

The service provider gives the organisation access to the network operator and the services the operator provides. The enabler (part of the service provider) will provide the access by the user to the service provider, and may get paid by the operator party (part of the service provider) for so doing. The network operator actually runs the communications networks and the services hosted thereon. The infrastructure grouping provides the hardware and infrastructure of the network and its access lines.

Figure 1—Example of a value chain

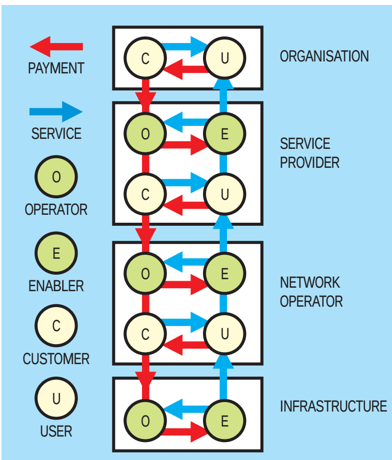

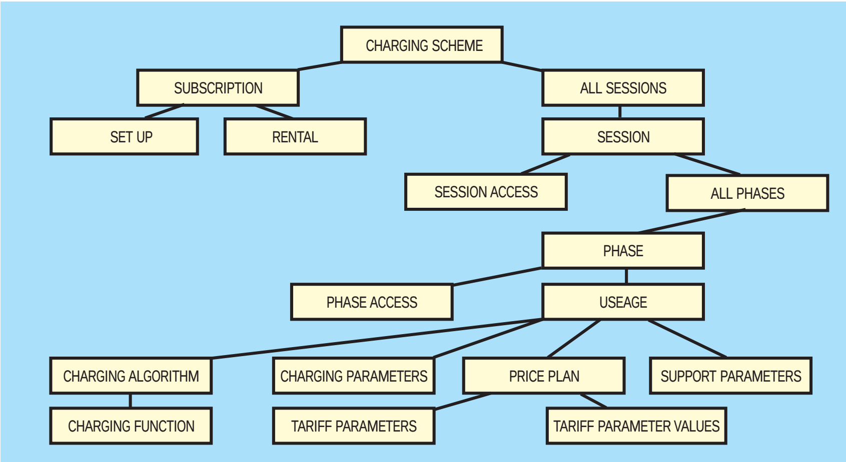

Figure 2—Charging scheme structure

Thus the enabler party within the infrastructure grouping will install the hardware and systems needed by the operator, including the hardware and systems needed to access the network.

Charging Schemes: Structure and Parameters

The subscription part of a charging scheme[3] consists of a set-up and a rental charge (Figure 2). The former is usually a flat charge and is sometimes termed the *joining fee* for a service. The recurring charge is often termed *rental* and is often a function of time. There may be several *sessions*[4] spread over the rental period. Each session may consist of several *phases*, each with its own set-up and usage charges. In this context, a phase[5] is a continuous period during which none of the following four parameters associated with an ATM connection change: quality of service (QoS), interruption of service (could be considered as a degraded QoS), time segment and traffic descriptor.

The usage part of a phase consists of a charging algorithm, one or more charging parameters, a price plan and, possibly, one or more support parameters. A *charging algorithm* consists of one or more charging functions in each of which the charge appears as a function of the charging, tariff and support parameters described below. The *charging parameters* quantify the 'amount' of service used. The price plan consists of *tariff parameters* that price the individual amounts per unit, and of tariff parameter values. *Support parameters* may be needed to complete the description of the charging scheme. An example, covering the phase access and usage part of a duration-based charging scheme is given in Table 1.

Charging Parameters are associated with network usage. Examples include volume, duration, effective bandwidth and contracted peak or sustainable cell rates. Charging parameters can have contracted values such as mean or peak allowable cell rates. These contracted values are known before a connection starts to operate and can thus be regarded as static, since the values do not change over at least the lifetime of a given phase of a connection. Other charging parameters can be measured and are regarded as dynamic (see below). Thus while the number of cells (volume of cells) *actually* passed by a connection may not be known before a connection starts to operate, this number may theoretically be measurable.

Support Parameters may be needed to complete the description of a charging scheme, but are not charged directly. Thus extra volume above some minimum agreed volume may be charged; however, the extra volume requires the measurement of the total volume which is not charged directly. Another example is instantaneous cell rate or bandwidth. This may not be charged, provided it does not exceed some agreed peak cell rate, itself a charging parameter; however, the instantaneous cell rate must be measured.

Tariff Parameters put monetary value onto service usage. Examples are Euro/cell, Francs/Mbit, and Dollars/second. The values of Tariff Parameters are influenced by market as well as technical considerations. As stated earlier, for a particular instantiation of a charging scheme, the collection of tariff parameters together with their values constitutes the *price-plan* for that scheme.

A further useful classification is in terms of static and dynamic parameters. *Static Parameters* are usually known at the start of phase and may also be known at the start of a session. Static parameters may also be fixed at predetermined times. Examples are contracted peak or sustainable cell rates. Static parameters are usually some function of:

- the objective ATM parameters as defined by the ATM Forum (for example, peak, sustainable and minimum cell rates, maximum burst size, cell delay variation, maximum cell transfer delay and cell loss ratio[6], and
- QoS classes.

Static parameters are usually based on assumptions about source statistics and are used by Connection Admission Control (CAC) to reach an accept/reject decision and to set up the Usage Parameter Control (UPC) (for example, leaky bucket) parameters.

Dynamic parameters are those that are *measured* while the connection is in progress or on the release of a connection. These parameters are used to establish the session-usage charge and examples include volume and duration. Dynamic parameters are usually measured at the point at which a user's traffic *enters* the network, but could also be measured at the receiving end.

Table 1: Typical sample of duration-based charging scheme

Entity	Symbol	Typical unit
charging function	$C = K \times t + S$	Dollars
charging parameter	t = duration	Second
tariff parameters	K = constant	Dollars/second
	S = access charge	Dollars

What Makes a Good Charging Scheme?

How do we assess the merits of a particular charging scheme? For the customer, ease of use and profitability are important considerations, while the operator may be concerned, not just with profitability but with the effect of a scheme on network performance and also with the ease with which a new charging scheme can be introduced[3,7].

Operator viewpoints

When introducing a new charging scheme, simplicity in terms of the required hardware and software will be important. How easily can the charging scheme be integrated into the existing environment comprising the host computer and the network itself? Are the measurement requirements complex or constrained by speed considerations? Does the implementation of the charging scheme depend on how it will interact with the CAC and UPC functions? Will the charging scheme entail high customer care and customer billing costs because of its complexity and the need to educate the customer about the basis for charging? Can the charging scheme be expanded easily to accommodate the changing aspirations and requirements of users, customers, network operators and service providers? To which service categories may the charging scheme be applied?

The use of a charging scheme should not adversely influence network performance. A charging scheme and its associated on-line displays require computation and network bandwidth resources. This may reduce the bandwidth and computational resources available to the other traffic carried by the network and may also influence network performance. Computation in real time may be needed to obtain usage parameters such as time and duration of connection, and mean and peak cell rate; computation may need to work with the UPC function. In addition code and data sizes and dynamic memory requirements related to the computation may need to be considered. Bandwidth may be required to request and obtain measurement information from, and to send results to, other locations in the network. How well does a charging scheme foster the sharing of network resources by a multiplicity of uses? Does the use of the charging scheme risk causing network traffic instability or congestion? The formulation of a charging scheme should help the operator to predict the effect a user will have on the network. Thus a band-limited scheme based on duration charging (for example, as used for POTS) gives the network operator information about the *maximum* requirement a user will make on bandwidth resources.

The ultimate purpose of a charging scheme is usually to generate revenue and earn profit for the supplier or operator. The scheme should therefore facilitate the identification and modification of the relevant parameters. Such mechanisms include peak-load pricing, discounts for high usage, ability to recover incremental costs, price discrimination and ability to response to elasticity in demand. A usage-sensitive charging scheme can help a network operator (NO) or service provider to relate the price he charges for a service to the use that the service makes of network resources. This can help the NO to set prices that will ensure he recovers his costs. The NO will usually appreciate a charging scheme that enables him to predict effectively the revenue generated by a customer's use of a particular charging scheme. The advantage to the NO is not only an accurate revenue forecast, but also a useful indication as to the values to apply to the appropriate tariff parameters. Uncertainty in establishing appropriate tariff parameter values could lead to values being set too high with a consequent customer rejection, or it could lead to the opposite extreme with consequent loss in profit. Indeed, this uncertainty could also lead to violent swings in the tariff parameter values, before a steady state is reached! Can the scheme be fooled by the user? How adaptable is it to changes in traffic conditions? Additionally, a scheme that depends on the content of a call can give the operator flexibility in adjusting prices.

Customer viewpoints

One customer survey (the CANCAN User Forum[8]), indicated that the criteria deemed as most important to the customer are the knowledge required, predictability, ability to audit the customer's bill, and whether or not the scheme is resource based.

How much does the user need to know about the workings of the charging scheme? Does the user need to pre-declare detailed traffic statistics, the complexity of which may not be commensurate with his mathematical knowledge? How much experience does the user need in order to maximise his benefit when employing the particular scheme?

The charging scheme should represent value to the customer in terms of quality of service and price. Knowledge of the nature of (see above), and experience in using, a particular charging scheme could help customers to maximise their gains when employing the charging scheme. This experience is fostered by factors such as: usage sensitivity, predictability, conformance to contract and ease of auditing the bill. Customers and users may like a scheme that gives them the possibility of tuning and re-negotiating prices for network resources according to demand. The prices to be paid for dynamic charging schemes are additional complexity, bandwidth requirement and computational overhead; these can increase the net costs to the customer. How secure is the scheme? For example, consider a charging scheme that samples bandwidth at regular intervals; an unscrupulous operator could ensure that traffic is only measured when traffic bursts occur. To what extent can the charge be estimated before start of the call? Is the tariff known before a call?

Some Candidate ATM Charging Schemes

We shall now turn our attention to some possible charging schemes for ATM. They all involve a need to measure charging parameters related to either time, cell volume, or both.

Three term charging

In this scheme we have a charge C composed of a flat rate term D, a duration-dependent term and a volume dependent term. In the duration dependent term E is the charge per unit time and t is the duration or lifetime of the connection. In the volume dependent term F is the charge per cell, and v is the total number of cells passed by the connection. The equation for this scheme is as follows:

$$C = D + E \times t + F \times v$$

While the tariff parameters D, E and F are likely to be contractual, the scheme, as presented, contains no contractual obligation regarding either the total number of cells to be passed by the connection or the maximum rate at which these may be passed. Hence the charging scheme, as it stands, can offer the customer surprises in the bill regarding the resultant total charge for the connection. In practice, the operator has to rely on mechanisms such as usage parameter control to limit a user's bandwidth. Such action provides a means of predicting the worse case effect a user might have on the network and can help to prevent network congestion and instability

Flat-rate charging

The flat-rate charging (parameter *D*) is very simple and has negligible computational and bandwidth overhead. It can thus easily be integrated into a computing environment and into the network. However, since it is not usage sensitive, there are no inducements in the scheme itself for the user to control his traffic volume. However, flat-rate charging is easy to explain to a customer who might find a flat rate charge—if is not excessive—quite attractive on account of its simplicity and potential to be profitable. It is a trivial task to audit the bill that the customer receives.

Duration charging

Duration charging, the second term, is familiar to customers since it is widely used in traditional telephone networks, albeit with bandwidth limiting. From the operator viewpoint, duration charging is very simple to set up and to explain to customers. It is best suited to CBR traffic. The bandwidth and computational overheads for running the scheme are low. Nothing is said in the contract about the traffic other than its duration and the bandwidth limiting imposed by the network. The customer gets charged for a maximum bandwidth, even if he or she does not use all of it. Hence the scheme can be profitable to the operator and it may also give the operator some flexibility in determining *E*, the price per duration unit. *E* could be based on some bandwidth other than the peak value.

Volume-based charging

In communications systems the commodity is information. One can try and quantify this by relating the amount of information conveyed by a connection to some function derived from counting the number of cells or bits actually sent in that connection. Volume-based charging may be attractive to customers since they are charged according to what they use. This is akin to the way customers are charged for mains electricity. It has been suggested[9] that volume-based charging could especially suit users of available bit rate (ABR) category services (for example, those transferring files or library information overnight) where access and transfer times are not important but space priority (no cell loss) is very important. Versions of volume-based charging are Walker's scheme[10] for wholesale charging and Viero's scheme[5] involving traffic sampling. Volume-based charging also features in schemes based on effective bandwidth and in dynamic charging schemes that are considered below.

An attractive feature of the three-term scheme is that it is simple enough to be implemented easily in silicon provided that the cell volume counters are able to accommodate the maximum expected cell rate; present technology may limit this to about 100 Mbit/s.

Effective bandwidth

The effective bandwidth of a variable bit rate (VBR) source aims to represent all the statistical characteristics of the source by only one parameter. The VBR source can then be charged as a CBR source. The concept of effective bandwidth is useful where a link has a common grade of service (GoS) parameter for each of the traffic connections and carries mainly VBR traffic, where each VBR connection expects multiplexing gains.

Possible candidates for this GoS parameter are cell-loss ratio (CLR) and probability of saturation. Irrespective of the basis for its calculation, the value of the effective bandwidth of a source usually lies *between* the source's peak and mean bandwidths. The long term average CLR is a natural choice, since cell loss is a phenomenon that users can readily appreciate and the CLR can be measured and thus can be audited after the call has finished. For bursty traffic, the probability of saturation (POS) is the probability that some arriving cells might be lost. While effective bandwidth based on the POS is easier to calculate, when compared with that derived using the CLR, it can be a very conservative estimate[5]. However, the EB based on the POS can be calculated[11,12] for a source, independently of the other sources that are going to mix with it, where all sources share a common resource such as an ATM link

The effective bandwidth is important for several charging schemes including Kelly's and Botvitch's schemes. Kelly's charging scheme[13] is based on the use of probability theory and requires knowledge of the statistics of the ATM cell traffic generated by the user's source. Kelly's algorithm assumes the user has a source for which he knows the peak bit rate as well as the probability distribution for the mean bit rate *M*, but not *M* itself. The network operator (NO) uses this information to construct a graph with axes for effective bandwidth *E*(*M*) and for the mean bandwidth *M* of the user. A price line—

$$P = a + bM$$

—may then be drawn on the same graph and moved across it in a vertical direction. *P* is a price proportional to the actual effective bandwidth of the user, and *a* and *b* are proportional to charge per second and charge per volume per second respectively. A remarkable result is that the declaration *m* is derived from the point at which the price line *P* just touches the curve *E(M),* where *E(M)* is the effective bandwidth derived the user's source statistics with *M* as the unknown. The NO assumes that the user, in order to minimise the charge levied on him, declares a mean volume rate m. The NO charges the user extra for any excess effective bandwidth the user uses above *E(m)*. The user may well end up paying the NO for bandwidth above the agreed *E(m)* value, and the NO himself can not predict the value of the mean bandwidth the user might actually use—there can be surprises for both user and operator! On the other hand, the user may use less than the declared *E(m)* and may not get a corresponding price reduction. The question of maintaining a record of the various values of *E(M)* over the lifetime of a connection for customer validation has not yet been addressed for Kelly's scheme. The NO must use an on-line effective bandwidth estimator[5] to determine the effective bandwidth actually used by the user.

The Botvich Effective Bandwidth Scheme[9] is based on the following formula:

$$K = \Sigma(T_i B_i) / \Sigma(T_i m_i)$$

where each connection i has a duration T_i, an effective bandwidth B_i and a mean bandwidth m_i. Σ indicates summation over all i terms. K is the aggregate factor of proportionality for i known connections. A new VBR connection with mean rate m will be charged in the same way as a CBR source whose rate is $B = K \times m$. The network has to measure on-line the mean rate of each new connection but the effective bandwidth is now derived from off-line measurements that are used to derive the factor K.

Envelope Charging

The Envelope Charging Scheme, due to Griffiths[5,9,14], offers no surprises to user or operator, while at the same time allowing for some burstiness in the user's ATM cell stream. The user is given a statistical envelope that he must not exceed. Poisson statistics are chosen because of their wide applicability and their linear additive property. The Poisson envelope can be enforced by a Poisson statistical filter, which can be implemented by a suitable combination of standard UPC components (two leaky buckets and a small buffer with negligible delay). Charging can be done on the basis of connect time, which facilitates validation; a record could also be kept of the number of occurrences of leaky bucket overflows (that is, contract violations by the user). Of all the charging schemes considered above, the envelope method is closest in concept to the original telephone duration-based charging scheme, while at the same time recognising the statistical multiplexing offered by ATM.

Dynamic charging

Dynamic charging may be defined in terms of dynamic pricing and on-line contract renegotiation[15]. On-line Contract renegotiation, in turn, can mean in-call renegotiation or spot-market re-negotiation.

Dynamic Pricing

The values of the tariff parameters of the charging scheme can vary with time. Prices varying slowly in response to inflation or competitive pressures or prices varying according to some peak/off-peak scheme are largely predictable by the customer. On the other hand, the NO may need to vary prices rapidly (for example, in response to the need for network congestion control), in order to maintain a customer's QoS, and this may lead to surprises on the customer's bill.

In-call contract renegotiation

In-call contract re-negotiation is usually a customer-controlled activity and may be viewed as a means of optimising charges and as a means of adapting the traffic contract to the bandwidth requirements of the application, possibly irrespective of the associated charges. Parameters whose values can be renegotiated may include the peak cell rate or a QoS attribute. The Distributed Pricing Scheme by Murphy at al. [16] for in-call contract renegotiation acts as a CAC mechanism and bandwidth allocation mechanism and is viewed as offering the possibility of eliminating both the CAC and UPC functions. Users (sources) compete for network resources according to their individual bandwidth requirements. For this reason, users may not be able to predict the final charge for any particular call. Network resources, such as signalling and processing influence the effectiveness of the renegotiation process[15,17]. Renegotiation with fixed phase lengths appears to be robust with respect to the exact phase length chosen, but advantages to both user and network operator vanish quickly with increasing renegotiation cost. Renegotiation with variable phase length is considerably more complicated than renegotiation with fixed phase length, but results in better transmission and renegotiation resource usage. These results are valid even if, before the call commences, there is no detailed information available on the bandwidth requirements that will arise during that call. The customer may be reluctant to renegotiate to reduce his bandwidth requirement if he perceives that any advantage thus gained is offset by the effort of the renegotiation process itself. Recently, Bigham et al.[18] have applied a system of distributed intelligent agents to the task of resource management and charging in ATM networks.

Spot Market

The telecommunications spot market is important for service providers and large corporate customers, since it will simplify buying and selling of telecommunication services, possible commodities being bandwidth/destination, time/destination and volume/destination[19]. Band-X, a London based Web exchange[20] and Rate Xchange in San Francisco[21] are leading the way in establishing spot markets for telecommunications bandwidth. A mathematical model of a telecommunication spot market has been developed[15]. Price movement in the spot market is modelled in a static fashion for discontinuous trading but this work provides a pointer to modeling the spot market in a dynamic fashion using stochastic models to represent continuous trading. A study of market-based call routing in ATM networks has been reported by Gibney et al.[22], who achieved adaptive pricing and real bidding by using distributed intelligent co-operating agents.

Conclusions

A charging scheme structure has been developed to include concepts such as subscription and usage as well as tariff, charging, static, dynamic and measured parameters. Operator viewpoints include the ease of introducing a charging scheme, its effect on network performance and its ability to generate revenue. A customer will view a charging scheme in terms of ease of understanding and profitability. Duration-based charging, especially envelope charging, looks attractive for DBR services, whereas VBR services could benefit from charging based on effective bandwidth. Purely volume-based charging may be more suited for ABR and UBR traffic. Dynamic charging is the focus of much research and could be attractive for in-call contract renegotiation as well as for spot-market renegotiation. The cost of renegotiation is an important consideration.

Acknowledgements

The author would like to acknowledge the work of colleagues in the CANCAN consortium (ACTS Project AC014). The project gave rise to a number of the concepts described in this paper.

References

1 CANCAN Consortium, Contract Negotiation and Charging in ATM

Networks, ACTS Project AC014, 1995-1998).
2 CASHMAN Consortium, Charging and Charging Schemes in Multi-Service ATM Networks, Contract Negotiation and Charging in ATM Networks, ACTS Project AC039, 1995-1998.
3 SCHARF, E. M.; and GRIFFITHS, J. M. Charging for ATM—Case Studies by the CANCAN Project, Third IFIP Workshop on the Management and Design of ATM Networks, London, 26-27 April 1999.
4 McGIBNEY et al. Contracts for ATM Services: A Structured Analysis. IEE Colloquium on Broadband Charging, IEE, London, 12 November 1996.
5 CANCAN Consortium, *ATM* Charging Schemes—Review of ATM Charging Schemes and their Performance Issues, CANCAN Deliverable D5, AC014/QMW/-/DS/P/005/b1, Document 312ds-b2.doc, October 1996.
6 LAMBARELLI, L. (Editor), ATM Service Categories: The Benefits to the User. White Paper, The European Market Awareness Committee, The ATM Forum, Brussels, May 1996.
7 SCHARF, *E. M.* First Steps towards the Development and Assessment of Effective ATM Charging Algorithms. IEE Colloquium on Broadband Charging, IEE, London, 12th November 1996.
8 CANCAN Consortium, Phase A Field Trial & User Forum Report, documents, 007ds-b2.doc and 07ads-b2.doc, October 1996.
9 CANCAN Consortium, Final Report on Static Charging Schemes and Their Performance, CANCAN Deliverable D9a, AC014/QMW/DS/P/350/a7, 1997.
10 CANCAN Consortium, Assessment of a Selection of Charging Schemes, CANCAN Deliverable D9b, AC014/LUN/DS/P/ 009/B2, May1997.
11 HUI, J. Y. Switching and Traffic Theory for Integrated Broadband Networks, 1st. Ed., Kluwer, Boston'90.
12 KELLY, F. P. Effective Bandwidths at Multi-Class Queues. *Queuing Systems Theory and Applications*, Vol. 9, No. 1-2, pp 5-15, October 1991.
13 KELLY, F. P. Tariffs, Policing and Admission Control for Multiservice Networks. 10th UK IEE Teletraffic Symposium, BT Laboratories, Martlesham Heath, April 1993.
14 GRIFFITHS, J. M.; and CUTHBERT, L. G. Very Selective ATM Filter, *Electronics Letters*, vol. 32, issue 7, 28 March 1996, pp. 637-8.
15 CANCAN Consortium, Dynamic Charging Schemes: Network Implications, CANCAN Deliverable D16, AC014/QMW/-/DS/I/370/A7, Document 370in-a7.doc, July 1998.
16 MURPHY, J.; and MURPHY, L. Bandwidth Allocation by Pricing in ATM Networks. IFIP Transactions C-24: Broadband Communications II, pp 333-351, North-Holland, 1994.
17 GRIFFITHS, J. M.; and CUTHBERT, L. G. The Impact of the Cost of Renegotiating ATM Rates, Proc. 5th IFIP ATM Workshop on Performance Modelling and Evaluation of ATM Networks, Ilkley, UK, 21–24 July 1997.
18 Bigham, J. *et al*., *Resource Management & Charging Using a Multi-Agent System*, 3rd IFIP Workshop on the Management and Design of ATM Networks, Queen Mary and Westfield College, London, April'99.
19 *Communications Week*, 3rd Nov. 1997.
20 Band-X sets up spot market in telephone minutes, Financial Times, Monday 29 June, 1998.
21 http://www.RateXchange.com.
22 GIBNEY, M. A. et al, Market-based Call Routing in ATM Networks using Adaptive Pricing and Real Bidding, 3rd IFIP Workshop on the Management and Design of ATM Networks, Queen Mary and Westfield College, London, April 1999.

Biography

Dr. Eric M. Scharf
Queen Mary and Westfield College

Eric Scharf is a lecturer in the Department of Electronic Engineering at Queen Mary & Westfield College, University of London. Since 1989 he has been involved with European Union funded projects on ATM technology. In particular he has worked in the areas of network control and measurement as well as charging. He has also been involved with the Distributed Summer School on Broad Band Technology, held annually from 1993 to 1996. His interests are communication networks and computing, and he has written and co-authored several papers in these areas. He lectures in the areas of programming, electronics and microprocessors.