Pricing For ATM Network Efficiency

Liam Murphy[*] John Murphy[†]

Abstract

ATM networks are expected to accommodate a wide range of users including some who can tolerate a certain amount of cell loss and/or delay. There are also likely to be some users who can modify their traffic inputs in response to feedback signals from the network. We propose a feedback scheme to increase network efficiency by taking advantage of this flexibility. Our scheme is based on economic principles of pricing for resource allocation. The pricing framework also provides the basis for a fast reservation scheme for adaptive users who require cell loss guarantees.

1 Introduction : Integrated Services Networks

Asynchronous Transfer Mode (ATM) is expected to be the basis for networks capable of carrying a wide range of user applications in an integrated way. Such a network is usually referred to as an *Integrated Services Network*. To some, this means a single network which supports some or all of the services currently offered by separate networks, such as voice and data. To others, it means a network which can carry all current and future user applications. Before we discuss the problem of resource allocation in ATM networks, therefore, we first outline the kind of network – and user – that we envisage.

1.1 The user perspective

We discuss the user perspective on Integrated Services Networks first to emphasize that user preferences *should* be the primary consideration. Once the network is in place only the users directly benefit from using it, by running applications which achieve higher–layer communication goals. Network owners and operators benefit indirectly, by providing services that users want or are willing to pay for.

---

[*] Department of Computer Science and Engineering, Auburn University, AL 36849, USA; **lmurphy@eng.auburn.edu**

[†] School of Electronic Engineering, Dublin City University, Glasnevin, Dublin 9, Ireland; **murphyj@eeng.dcu.ie**

From the user's point of view, there should be as few restrictions as possible on the communication services they obtain from the network. In particular :

- a user should be able to specify **any values** for the various traffic parameters the network may ask for prior to accepting a connection; or simply specify nothing about the traffic their connection will generate.
- a user should be able to demand **any values** of the various Quality Of Service (QOS) parameters that the network has defined; or simply tell the network to provide the best possible service for a given price, with no guarantees required.
- a user should be able to **adjust connection parameters dynamically during the connection lifetime, if desired.** For example, many data transfer applications are flexible regarding the delay incurred in completing the transfer.
- perhaps most importantly, the user should have a **simple interface** to the network to conduct these negotiations.

However some processing power is assumed to be available on the user side of the user–network interface, so that the details are handled by application hardware or software. Therefore the restriction to a simple user interface to the network can be relaxed to a simple *human* interface. Any complicated processing necessary to translate user commands into actions affecting their network connection is assigned to the local processing element[1].

This flexibility in user service characteristics is motivated by the observation that *it is becoming more and more difficult to accurately define a "typical" user.* There is already a spectrum of such user traffic characteristics as mean bit–rate or peak–to–mean bit–rate ratio. In addition, technological advances may continue to change the requirements for present–day services, for instance by reducing the bandwidth needed for voice or VCR–quality–video calls. There is also a wide range of user QOS requirements even within many of the service classes proposed in the literature. For example, some "videophone" users may require a high–quality reliable connection while others may be satisfied with poorer–quality or interruptible connections.

1.2 The network perspective

An Integrated Services Network could range from a Local Area Network (LAN) to a world–wide Wide Area Network (WAN), and could be private (all the applications controlled by one organization) or public.

---

[1] We will return later to the question of how much local processing power might be needed.

The operation of a public network may be the responsibility of several organizations, within each of which the operational functions may be automated and / or distributed. Conceptually, however, the control and management functions of an Integrated Services Network can be associated with a *network operator* as if one entity was responsible for controlling and operating the network.

The network operator tries to satisfy three competing objectives (Figure 1) :

- **performance guarantees** should be provided and fulfilled for those users who require them. These guarantees may be deterministic or statistical.
- **all services** demanded by the users should be supported, including (ideally) future services with as–yet unknown characteristics.
- network operation should be an **efficient utilization** of network resources, such as link capacities and node buffers. At the very least, network operation should be more efficient than a *circuit–switched* or *peak–allocation* strategy in which the maximum resources required by a connection are reserved for the duration of that connection, regardless of whether or not the resources are actually used continuously.

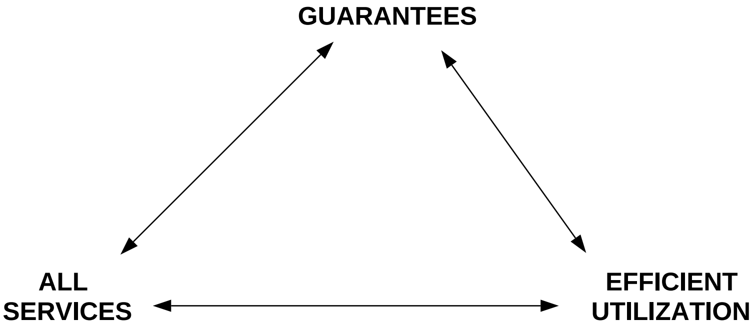

Figure 1: Network objectives

The difficulty with these objectives is that any two of them can be achieved relatively easily – conceptually at any rate – but simultaneously achieving all three is still an open problem. More specifically :

- *a wide range of services can be supported with efficient use of network resources, provided no guarantees have to be made to the users.* Typical values of network performance measures may be a good indication of the expected QOS, aggregated over time and all users, but some user applications require more specific guarantees.
- *guarantees can be made to the users and network operation can be efficient, provided only one (or a narrow range) of service types have to be supported.* Focusing on one type of service allows the network

to be optimized to efficiently deliver that service in the ways required by the users. This is essentially the traditional telephone network model.

- *a wide range of services can be supported, and guarantees demanded by the users can be offered and fulfilled, provided efficient network operation is not important.* This is usually achieved by an overprovisioning of resources, such as reserving the peak bandwidth required by a connection.

The network operator should also be able to offer more customized service to individual users than is usually available today. This means that traditional network performance metrics such as average delay or packet loss may not be fine–grained enough, since typically a user cares only about the QOS their connections receive. Advances in user hardware and software, and the development of a competitive network provider industry, will require network operators to **focus on satisfying the communications needs of an individual user**, regardless of the size of the user's connection.

2 ATM–based networks

ATM is a form of fast packet switching which uses short fixed–length packets called **cells** [1]. The fixed size and reduced header functionality of ATM cells is intended to facilitate the high–speed switching necessary in networks with link speeds of 155 – 622 Mbps. The short cell length (48–byte information field + 5–byte header) is intended to keep the packetization delays for real–time services within acceptable limits. ATM is connection–oriented and supports hierarchical addressing through the use of Virtual Channel (VC) and Virtual Path (VP) connections. A VC identifier (VCI) and VP identifier (VPI) are contained in the header of each cell of an end–to–end connection. Hop–by–hop routing is implemented by changing the VPI and/or the VCI at intermediate ATM switches.

2.1 The promise of ATM

ATM has been adopted by ITU (formerly CCITT) as the transfer mode for the Broadband Integrated Services Digital Network (BISDN), a service–independent network capable of supporting all the communication services that users now require or may require in the future. ATM is also emerging as a local area networking technology, since it provides flexible bandwidth–on–demand and internetworking capabilities for conventional data communications. The data communications and computer industries are now participating in ATM specification and development through the ATM Forum.

If ATM can in fact transport any service – irrespective of its bit rate, bursty nature or QOS requirements – it offers the possibility of a **unifying communications paradigm**, enabling not only integrated applications but also integrated operating environments and administrative domains [2]. However ATM development and implementation is a dynamic area, with a wide range of proposed solutions but little experience with actual users. Whether ATM can live up to its promise as a universal transfer mode is still an open question.

2.2 Defining services in an ATM network

An ATM–based network is expected to allow users to negotiate their own service definitions. In this way a user can get a service appropriate to their individual needs. This assumes that some users will take a more active role in defining their network services than in present–day networks : these users are assumed to enter into negotiations with the network at the **connection level**, whenever a new connection is requested.

Some applications require guarantees from the network, on loss or delay or both. The user–network relationship for these applications in many of the proposed Connection Admission Control (CAC) schemes is a **contract** : users describe their traffic and make QOS demands, and the network provides a stated level of service while enforcing the user commitments. When a new connection is requested, the network must decide whether or not to accept the connection; and if so, how to route it through the network and what resources to reserve for it. The connection request is refused by the network if accepting it would lead to QOS degradation for one or more currently–active connections which are complying with their contracts. Some CAC schemes provide feedback on the reason for the refusal [3], which guides the user in submitting a revised request.

Some proposed CAC schemes reserve the peak bandwidth needed by an accepted connection – *deterministic multiplexing* – as required for constant–bit–rate (CBR) sources. The gain in network utilization possible by taking advantage of the statistical nature of variable–bit–rate (VBR) sources has led to many proposed schemes for *statistical multiplexing*. Such schemes assign less than the peak bandwidth required, and therefore may introduce cell loss and/or delay.

The aim of a preventive CAC scheme is to balance the QOS provided to admitted connections against network utilization by limiting the number of connections using the network. Most proposed CAC schemes decide whether or not to accept a connection request based on knowledge of the *traffic description*, the user's *QOS requirements*, and the *current state* of the network :

- ideally a user requesting a connection would give a complete statistical description of their traffic, but in practice only a limited indication of expected behavior is feasible. Connection behavior is described by a set of parameters called traffic descriptors, such as mean bit–rate, peak bit–rate, maximum burst length, probability of cell arrival in a fixed interval, and so on.

- user QOS requirements are usually expressed by networks researchers in terms of acceptable cell loss, delay and jitter. How to relate these quantities to parameters the user is concerned with (such as picture quality in a video call) is an active area of study.
- the current state of the network can be determined by monitoring the utilization of network resources and/or by characterizing the behavior of connections already admitted.

Based on this knowledge, CAC schemes have been proposed in which an effective bandwidth [4] is associated with each source in order to meet its QOS requirements while still permitting a statistical multiplexing gain.

Accepted connections are provided with a traffic contract and as long as they comply with that contract, their cells are treated as high priority by the network. Cells submitted in excess of their contract may be discarded at the network access point, or may be marked as low priority and discarded if the network is congested. Therefore the network provides no guarantees for low–priority traffic.

Not all applications require guarantees from the network. For example, some types of non–real–time data transfers can tolerate a certain amount of cell loss; or the application may recover from cell loss or delay at a higher layer of the protocol stack. These **best–effort** applications request an Unspecified QOS service [5], in which the network always accepts the connection request but may assign a different Peak Cell Rate (PCR) than what was requested (for example, a connection may be assigned a PCR of zero). Cells which are submitted in excess of the PCR are marked as low priority.

3 Adaptive users

In the previous Section, network services were divided into "guarantees–required" and "best–effort", and an enforceable traffic description was needed from users who wanted guarantees from the network. This breakdown of services does not capture all possible user desires now, let alone in the future, because it does not account for **adaptive** users. We say that a user is adaptive if they are able *and* willing to respond to feedback signals from the network during their connection lifetime by changing their offered traffic. The addition of willingness–to–respond reflects the fact that some users may run applications which are capable of adaption, but do not want to be concerned with network feedback during the connection. A user who is willing to wait (based on network feedback) to set up their connection but who wishes to be insulated from such feedback once their transmission starts is adaptive only at the connection level, and we do not regard such users as adaptive in this paper.

Adaptive users can help to increase network efficiency if they are given appropriate feedback signals. When the network load is high, the feedback should discourage adaptive users from injecting cells; when the load is low, the feedback should encourage these users to send any cells they have ready to transmit. Up to now, adaptive users have been regarded as lower priority than non–adaptive users who want guarantees from the network. Typically the latter are served first if possible, and any remaining bandwidth is shared out among the adaptive users on a best–effort basis. There is an implicit assumption in such schemes that the more demanding users pay premium prices in order to get the service they require, and the adaptive users pay lower prices as a reward for their flexibility.

Whether or not a network operator gives higher priority to more demanding users is a policy or business decision, not a technical one, and we do not intend to argue for one side or the other. However, technical problems do arise when some adaptive users also require guarantees on cell loss. For example, data transfers using compression may be very flexible with respect to delivery time but require zero cell loss to ensure data integrity, if there is little or no redundancy in the transmitted data. The classical approach to such users is to retransmit when errors occur, but in congested conditions this approach tends to increase the network load and worsen the congestion.

Another approach is to provide guarantees to adaptive users who require them. This is a difficult problem, since an adaptive user cannot give an accurate traffic description at connection setup–the traffic they send depends in some way on network conditions at the time their application generates it. Before we describe our proposed solution to this problem, we first discuss a general classification of user types which includes adaptive users.

3.1 Possible types of user

We consider four criteria for classifying user types (other criteria may also be useful and we do not claim to fully describe all possible user behavior) :

- a traffic description–in terms of network traffic descriptors–is available at connection setup;
- guarantees are required on cell delay. Since the user knows (or can determine) their own processing delays, the application's delay tolerance can be translated into a maximum acceptable delay within the network;
- guarantees are required on cell loss. In practice this means that the network guarantees to treat the user's cells as high priority;
- the user is adaptive, in the above sense

This classification is summarized in Figure 2.

Type	Adaptive	Traffic Descriptor	Loss Guarantee	Delay Guarantee
1	Yes	Yes	Yes	Yes
2				No
3			No	Yes
4				No
5		No	Yes	Yes
6				No
7			No	Yes
8				No
9	No	Yes	Yes	Yes
10				No
11			No	Yes
12				No
13		No	Yes	Yes
14				No
15			No	Yes
16				No

Figure 2: Classification of user types

We can rule out some of the entries in Figure 2 as follows :

- for non–adaptive users, a traffic description is required in order to get guarantees on loss and/or delay. Therefore rows 13, 14 and 15 can be eliminated;
- for non–adaptive users who provide suitable traffic descriptors to the network, we combine the cases where they require guarantees. Therefore rows 9, 10 and 11 can be combined into a single user type;
- for users who do not require guarantees, no traffic descriptor is needed (we can think of these users as marking all their cells to be low priority, as if they had been given Unspecified QOS service with a zero PCR). Therefore rows 4 and 12 can be eliminated;
- by definition, adaptive users do not have a traffic descriptor at the start of a connection. Therefore rows 1–4 can be eliminated;

- adaptive users who require guarantees on loss and delay have to decide a minimum acceptable bandwidth. This essentially converts their traffic into a non–adaptive component which has a traffic descriptor and requires loss and delay guarantees (row 9), and an adaptive component which requires no loss guarantees (row 7). Therefore row 5 can be eliminated.

Type	Old Type	Adaptive	Traffic Descriptor	Loss Guarantee	Delay Guarantee
1	9, 10, 11	No	Yes	Yes	
2	16		No	No	No
3	8	Yes	No	No	No
4	7				Yes
5	6			Yes	No

Figure 3: Expected user types

When these entries are eliminated, we are left with the user types shown in Figure 3 :

1. *non–adaptive, requires guarantees, has a traffic description at connection setup.* This is one of the user types addressed by various proposed CAC schemes, so we will not discuss it further in this paper.
2. *non–adaptive, no guarantees required.* This is non–adaptive best–effort traffic, such as the traffic a user injects in excess of their contract with current CAC schemes.
3. *adaptive, not delay–sensitive, no loss guarantees required.* This type of user waits until feedback from the network indicates that the load is low, then sends cells on a best–effort basis. This could be implemented by negotiating a traffic contract specifying zero mean and zero peak rates. Then there is no reservation of resources and all cells will be marked as low priority. We call these users **adaptive best–effort** users.
4. *adaptive, delay–sensitive, no loss guarantees required.* This type of user responds to network feedback by modifying how many cells they transmit, but sends them immediately on a best–effort basis. Otherwise they are similar to adaptive best–effort (type 3) users, and we call them **inelastic** users.
5. *adaptive, not delay–sensitive, requires loss guarantees.* This type of user waits until feedback indicates the network is lightly loaded, then transmits and requires that their cells are not lost in the network. We call these **elastic** users.

Traditionally the view has been that without an enforceable traffic description, no guarantees can be expected of the network. One possibility is that bursty sources can be well–described during a burst. The need for burst–level CAC has been widely discussed, e.g. [6], [7]; however burst start time and duration can be unpredictable.

We take the view that adaptive users can be accommodated with a **fast reservation** scheme. We propose a form of in–call negotiation on a timescale shorter than the dynamics of connection setup and teardown but longer than cell–scale effects. Elastic users who comply with these dynamic contracts transmit high–priority cells in the interval between successive feedback signals from the network.

3.2 Feedback–based fast reservation

Ideally, feedback would be provided continuously to adaptive users. In practice the signalling and computation required means that time must be divided into *feedback intervals*. We envisage feedback intervals on the order of a millisecond, so the adaptive user responses would have to be automated.

At the start of each feedback interval, the network sends a signal to each adaptive user. The responses of inelastic and adaptive best–effort users were outlined in the previous Section. Elastic users presumably have an application–dependent deadline by which transmission must be completed or the application will be aborted; we assume this deadline is long–term compared to the feedback dynamics and therefore is not a factor in their response to the feedback signals.

Based on the feedback signal and their application requirements, each elastic user signals the network with the number of cells they wish to transmit in this feedback interval. The network uses buffers to take care of possible cell–scale congestion, and decides if it can accommodate each elastic user's request with zero cell loss. If not, the request is denied and the user waits until the next feedback interval to try again (they can also transmit the cells immediately, at low priority, but that essentially converts them into adaptive best–effort users). If an elastic user's request is accepted, then they transmit their cells with high priority during the feedback interval and are guaranteed that the network will not lose those cells. A mechanism is probably needed to enforce the user requests, since otherwise users who are accepted but decide not to transmit unnecessarily block other users; we return to this point in Section 4.1.

The main problem with this kind of dynamic contract negotiation is that the network has to do the corresponding computations every feedback interval. If the users contending for shared network resources are geographically separated, the time to do the required signalling may be excessive even if the necessary computational power was available, and the feedback scheme would be potentially unstable.

One possibility is to localize the contention, for example by deterministically multiplexing the Virtual Paths (VPs) within the network. This reduces the network to a set of “virtual links” between source–destination pairs : the capacity of each such link is fixed for the feedback interval, and interactions within the network between traffic streams corresponding to different source–destination pairs are eliminated. Therefore the only contention between users is at the network access points. While we assume this model as a first cut, two points should be noted :

- the loss of efficiency by not statistically multiplexing the VPs within the network may be outweighed by the increased efficiency obtained by using feedback to multiplex the users at the network edges;
- the VP capacities could also be updated, on a longer timescale than the feedback intervals, to match overall dynamics of network traffic.

4 Price as a feedback signal

There are two distinct scenarios that need to be considered when deciding on an appropriate feedback signal : private and public networks. In a private network (such as a LAN/MAN or a company–wide network) the "users" are applications owned and controlled by one organization. Therefore the users are cooperative, since their responses to feedback can be programmed to obtain a desirable traffic mix. In a public network the users must be considered as separate entities, with their own private rules for deciding their traffic inputs. The network cannot assume the users will be cooperative without the right incentives.

A number of proposals have been made concerning the use of feedback in ATM networks, e.g. [8], [9], [10]. A crucial issue ignored in most feedback proposals is the basis on which decisions should be made, both by the network and by users. The scheme we propose addresses this issue by developing an economic framework in which incentives are provided to enable rational decisions and resource allocations. These incentives could correspond to actual money or could simply be control signals. For example, in a private network the prices are control signals which summarize the state of network resources such as bandwidth or buffer space.

4.1 Economic framework

In this Section we outline one possible method for computing adaptive prices. The users discussed below are thus adaptive users (either adaptive best–effort, inelastic, or elastic). A more detailed description can be found in [11], [12]. Other pricing schemes for communication networks have been suggested, e.g. [13], [14].

The network and its users are considered to form an *economy* or economic system. The system has various resources such as link bandwidths and buffer spaces that can be used to meet user demands for service. Network constraints such as buffer sizes or link capacities are translated into cost functions on the demands for resources. This reflects the fact that one user's consumption of bandwidth or buffer space gives rise to a "cost" (in terms of longer delays, less available bandwidth, longer buffer occupancies, etc) which is borne by all users.

Each adaptive user is viewed as placing a benefit, or willingness–to–pay, on the resources they are allocated. Given a price per unit of bandwidth or buffer space, a user's benefit function completely determines that user's traffic input. A benefit function could follow the usual economic assumption of diminishing incremental benefit as more of the resource is consumed (Figure 4(a)). Or it could be a simple threshold rule, or series of threshold rules, for deciding how much of the resource to request based on the current price (Figure 4(b),(c)). Users are allowed to change their benefit functions every feedback interval so the examples in Figure 4 are for a particular interval.

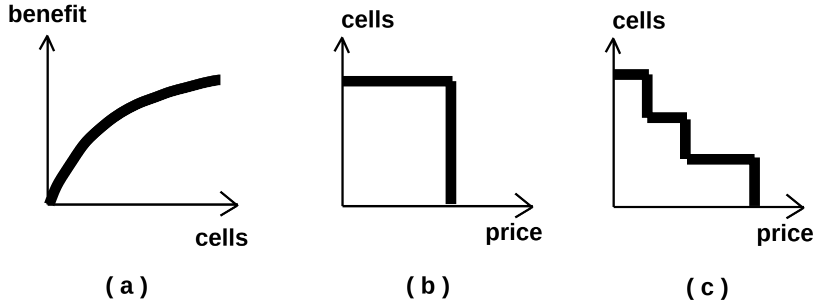

Figure 4: Possible user benefit functions

The network operator sets the prices so that the marginal benefit the users place on their resource allocation is equal to the marginal cost of handling the resulting traffic in the network[2]. The basic requirement is that price should go to infinity as usage of the resource approaches capacity. The network operator dynamically adjusts the prices based on current network conditions. It turns out that it is not necessary for the network operator to know the user benefit functions; therefore this pricing scheme is suitable for both public and private networks.

A distributed iterative pricing algorithm has been developed [12]. The distributed nature of the pricing algorithm suggests that it may be possible to meet the real–time feedback requirement. In addition, the computation required per iteration at each user and ATM access switch is simple, which suggests that inexpensive processing elements may be sufficient in executing the algorithm.

In our simulations, resource costs were chosen as barrier functions [15] to give a smooth increase to infinity as resource usage approached its limit. With a mix of voice, video and data sources sharing a small access buffer, the pricing scheme resulted in all traffic being accepted into the network without loss. A leaky–bucket scheme with the same sources resulted in 7.5% cell loss spread across all sources.

In a feedback–based fast reservation scheme, it may be necessary to enforce the user requests to avoid

---

[2] These prices only address the variable costs corresponding to network constraints.

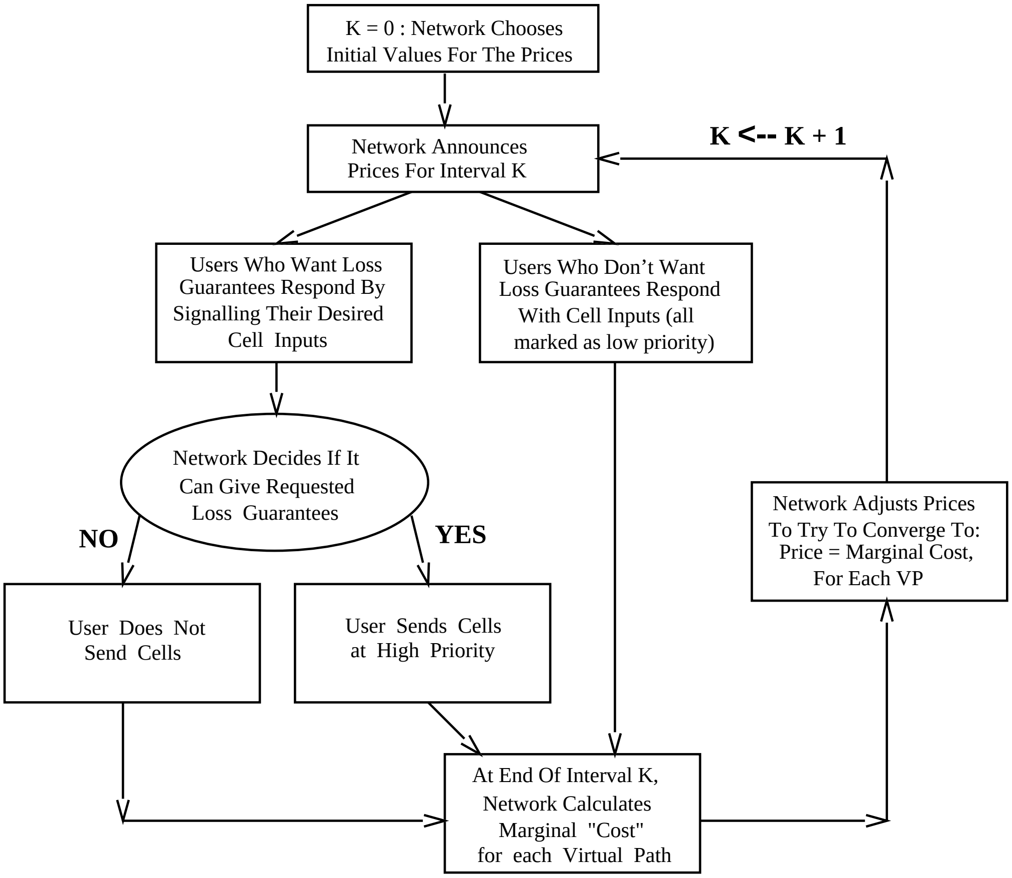

Figure 5: Distributed iterative pricing algorithm

unnecessarily blocking other users. One suggestion to provide financial incentives to users to encourage truthful traffic descriptors (for both adaptive and non–adaptive users) was presented in [16]. In the scheme shown in Figure 5, the network operator could charge each elastic user according to their request if their actual cell input is less than that.

In [11] the pricing algorithm of Figure 5 was extended to include adjustment of VP capacities using a similar iterative scheme. The VP adjustments are not purely local, since contention between users is not confined to the network access points as before. On the other hand the VP adjustments have a longer time period in which to carry out their computations.

5 Conclusions

We have proposed a feedback scheme for adaptive users which, by taking advantage of their flexibility, increases the efficiency of network operation. Our proposal should be viewed as a framework rather than a suggested implementation. There are many unresolved questions concerning this framework, including whether the minimum feasible pricing interval can capture the network traffic dynamics and whether it is applicable to more realistic, complicated network environments. Nevertheless it may provide another tool to network operators to assist them in coping with the increasing range of user demands in the Integrated Services Networks of the future.

Acknowledgements

We would like to thank Jeff MacKie-Mason for helpful discussions about this work. The second author thanks Prof. Charles McCorkell of Dublin City University for his continuing support and encouragement.

References

[1] M. de Prycker, *Asynchronous Transfer Mode : Solution for Broadband ISDN*, 2nd Ed., Ellis Horwood, 1993.

[2] I. M. Leslie, D. R. McAuley and D. L. Tennenhouse, 'ATM Everywhere ?', *IEEE Network*, Vol. 7, No. 2, pp 40-46, March 1993.

[3] D. Ferrari, J. Ramaekers and G. Ventre, 'Client-Network Interactions in Quality of Service Communication Environments', *IFIP Transactions C (Communication Systems)*, Vol. C-14, pp 221-234, 1993.

[4] F. P. Kelly, 'Effective Bandwidths at Multi-Class Queues', *Queueing Systems Theory and Applications*, Vol. 9, No. 1-2, pp 5-15, Oct. 1991.

[5] *ATM Forum*, 'ATM User–Network Interface Specification', Prentice Hall, 1993.

[6] J. S. Turner, 'Managing Bandwidth in ATM Networks with Bursty Traffic', *IEEE Network*, Vol. 6, No. 5, pp 50-58, Sept. 1992.

[7] J. W. Roberts, 'Variable-Bit-Rate Traffic Control in B-ISDN', *IEEE Commun. Mag.*, Vol. 29, No. 9, pp 50-56, Sept. 1991.

[8] H. Kanakia, P. P. Mishra and A. Reibman, 'An Adaptive Congestion Control Scheme for Real-Time Packet Video Transport', *Computer Communication Review*, Vol. 23, No. 4, pp 20-31, Oct. 1993.

[9] , Y. Gong and I. F. Akyildiz, 'Dynamic Traffic Control Using Feedback and Traffic Prediction in ATM Networks', *Proc. of IEEE INFOCOM '94*.

[10] N. Yin and M. G. Hluchyj, ‘On Closed–Loop Rate Control for ATM Cell Relay Networks’, *Proc. of IEEE INFOCOM '94*.

[11] J. Murphy, L. Murphy and E. C. Posner, ‘Distributed Pricing For Embedded ATM Networks’, *Proc. of International Teletraffic Congress ITC–14*, Antibes, France, June 1994.

[12] J. Murphy and L. Murphy, ‘Bandwidth Allocation By Pricing In ATM Networks’, *Proc. of IFIP Broadband Communications '94*, Paris, France, March 1994.

[13] J. K. MacKie-Mason and H. R. Varian, ‘Pricing the Internet’, presented at *Public Access to the Internet*, JFK School of Government, Harvard University, May 26-27, 1993.

[14] S. Low and P. Varaiya, ‘A Simple Theory of Traffic and Resource Allocation in ATM’, *Proc. of Globecom '91*, pp 1633-1637, Dec. 1991.

[15] D. G. Luenberger, *Linear and Nonlinear Programming*, 2nd Ed., Addison-Wesley, 1984.

[16] F. P. Kelly, ‘Tariffs, Policing and Admission Control for Multiservice Networks’, *10th UK Teletraffic Symp.*, BT Laboratories, Martlesham Heath, April 1993.