

Decentralized Control Scheme For Large-Scale ATM Networks*

Lorne G. Mason, Anne Pelletier, Eric Létourneau

INRS-Télécommunications

16, Place du Commerce

Verdun, Québec, Canada, H3E 1H6

Tél: (514) 765-7836, Fax: (514) 761-8501

lorne@inrs-telecom.quebec.ca

Abstract

Two approaches are proposed in this paper in order to facilitate network access control in large-scale ATM networks. We consider the centralized isarithmic control scheme that controls the total number of cells in the network using a permit mechanism. An implementation problem of this window-based scheme in large-scale networks is the significant control overhead associated with the flow of permits. The first approach is a partially decentralized control architecture that extends the centralized isarithmic scheme by means of an hierarchical zoning structure. An adaptive isarithmic controller is located in each zone with a distinct permit class and the zones interact through user traffic and management information flows. The decentralized architecture allows a substantial reduction of the permit flows and numerical results show that the network optimal operating point can be reached. Moreover, the level of management flows can be limited without deteriorating the network performance by using the approximate global information structure, based on mean delays. The second approach consists of decreasing the window size by letting one permit correspond to C credits. The control overhead is then reduced by a factor of C , the credit/permit ratio. Simulation results show that the network performance can even be improved compared to the centralized isarithmic scheme.

Keywords

IP/ATM, adaptive flow control

1 INTRODUCTION

A challenge in designing ATM networks is the synthesis of effective control architectures and algorithms for managing heterogeneous traffic and network resources such that quality of service constraints are met while achieving efficient use of the network resources. In this paper we address the issue of

*This work was supported by a grant from the Canadian Institute for Telecommunications Research under the NCE program of the Government of Canada.

network configuration management in large-scale ATM networks. Network configuration management refers to the functions of virtual network topology and bandwidth allocation which enable the partition of the aggregate network capacity into its component virtual networks. The control architecture considered here and outlined in [1] supports both switched connection-oriented services and connectionless services: switched VBR connections are carried on a virtual network while delay-insensitive traffic (either connectionless or connection-oriented) is handled via the residual network component. The logical bandwidth allocated to the delay-insensitive traffic must be enforced through either rate-based or credit-based congestion control schemes (see [2] and [3] and references therein).

In this paper, we consider the adaptive isarithmic flow control scheme [4] which regulates the delay-insensitive traffic flow entering the residual network based on its optimal operating point using a permit mechanism. The resulting window size and permit distribution determine the configuration and allocated bandwidth of the residual network. To ensure that the performance of time-critical VBR traffic is not degraded, we employ the selective window control with priority scheduling [4].

Isarithmic flow control was proposed by Davies [5] for network access control in packet-switching networks. This scheme limits the total number of packets that can be in transit within the network. This is accomplished by allowing packets to enter the network only after they secure a permit. When the packet/permit combination arrives at the destination node, the permit is released and becomes available for packets entering at that node. It was found that this approach can result in degraded performance under asymmetric traffic conditions. Mason and Gu [6] proposed to employ a controller for permit distribution. In the centralized version, a controller distributes the released permits to the source queues associated with each node. Another approach is decentralized where a controller is associated with each destination node [7]. For both cases, the authors proposed several learning automata control schemes to track the optimal operating point of the network. Cotton and Mason [8] designed an adaptive central controller for fast packet-switching networks. When compared to a rate-based scheme, the isarithmic scheme has been shown to have similar performance in terms of throughput and fairness [9].

Two approaches are proposed in this paper to enable the use of the isarithmic scheme in large-scale ATM networks. The first approach is a partially decentralized control architecture which extends the centralized isarithmic approach by means of a hierarchical zoning structure. The zones are sub-networks of the given network and are connected via gateway nodes to other zones. Each zone operates a centralized adaptive isarithmic control within its boundary for a distinct permit class and it interacts with the other zones through user traffic and management information flows. The effects of zone

partitioning and partial information structure on the resulting network performance are studied here.

The second approach consists of reducing the number of permits. As discussed in [10], the window size can be reduced by using one permit to control a group of cells instead of individual cell. We propose a different mechanism where we let one permit correspond to C credits. When a permit arrives at a node, the value of the counter increases by C . As in the original scheme, each admitted cell decreases the value of the counter by one. At the exit nodes each cell adds one to a counter of credits and when C are accumulated one permit is sent to the controller for dispersal. The permit flow in the network as well as the action rate of the controller are reduced by a factor of C .

We would like to mention that the control architecture described in [1] can be extended to the case of managing multiple virtual private networks which are expected to proliferate with the emergence of ATM networks – LAN interconnection being an important early example. In this concept different corporate users can be offered features similar to those of a dedicated private network while retaining the efficiency advantages of resource sharing and connectivity typical on an integrated public network. The multi-chain approach proposed in [6] allows the residual network to be subdivided into multiple such virtual networks carrying delay-insensitive traffic.

The paper is organized as follows. In Section 2, we present the performance criterion employed in this paper. Section 3 reviews the centralized adaptive isarithmic scheme and the selective windowing mechanism with priority. The decentralized control architecture is proposed in Section 4 and simulation results show that the network optimal operation point can be achieved. In Section 5, we study the effect of reduced control overhead by the introduction of the credit/permit ratio. Section 6 concludes.

2 PERFORMANCE MEASURE

A performance criterion for congestion control is the power function proposed by Kleinrock [11]. Defined as the ratio of average throughput over average delay, the power function unifies the antithetic objectives of high throughput and low delay. However, as Bharath-Kumar and Jaffe pointed out in [12], the power function is an unfair measure for network access control since it can happen that some users are denied access to the network. To prevent unfairness we employ the product of powers defined as [12]:

$$P = \prod_{o,d|\Lambda_{o,d} \neq 0} \frac{\lambda_{o,d}}{T_{o,d}} \quad (1)$$

where

$\Lambda_{o,d}$: exogenous traffic from node o to node d ;

$\lambda_{o,d}$: throughput of o - d traffic;

$T_{o,d}$: mean end-to-end delay of o - d traffic.

A proof that the product of powers is indeed a reasonable design criterion was provided by Mazumdar and *al.* [13]. The maximal product of powers point is unique and corresponds to the Nash arbitration strategy for the users: at optimum, if one user improves his power, it will adversely affect at least one other user power.

3 REVIEW OF CONTROL FRAMEWORK

In this section, we will describe in general the centralized adaptive isarithmic flow control scheme proposed by [6] [8] for high-speed data networks and the selective window control introduced by [4] to support different classes of traffic.

3.1 Centralized adaptive isarithmic flow control

In the centralized version of the adaptive isarithmic scheme, a controller is implemented to distribute the released permits as illustrated in Fig. 1. A packet is accepted to enter the network only if there is a permit available at the source queue. The packet/permit combination is routed through the network and at the destination node, the permit is sent to the central controller. The central controller is composed of two parts: the window size and the permit distribution controllers. The latter distributes adaptively the released permits while the window size controller adjusts the total number of permits, W , according to network conditions.

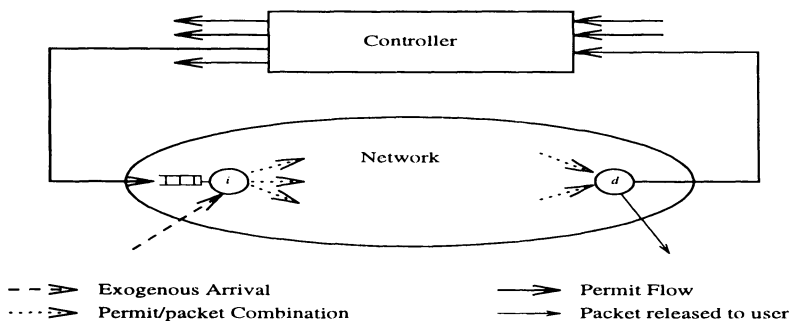


Figure 1 Centralized adaptive isarithmic flow control

The permit distribution controller is a *learning automaton* based on the structure proposed by Mason [14]: it is an L_{R-I} type automaton in an S -

model environment. The distribution probabilities are updated using a *learning algorithm* [8].

The second component of the central controller is the window size controller. The population of permits must also be adjusted since no value is optimum for all loads. Several methods were proposed for adaptive window size adjustment (see [15] and the references therein). A method to optimize the window size for a virtual circuit was proposed by Mitra and Seery [16]. It is based on an explicit expression relating the optimal window size to the mean round trip response time. For the isarithmic scheme, a self-optimizing method employing a cross-correlation technique for process identification was proposed by [17]. Another method to evaluate the sensitivity of the product of powers with respect to the window size is based on the phantom RPA method for estimation of derivatives [18].

3.2 Selective window control

The window size adaptation in ATM networks is more complicated than in packet-switching networks because of the frequent changes in the isochronous traffic level. To adapt the window size a fast algorithm is required which could be difficult to realize. The selective window control illustrated in Fig.2. was proposed by Liao and Mason [4] to avoid this adaptation problem.

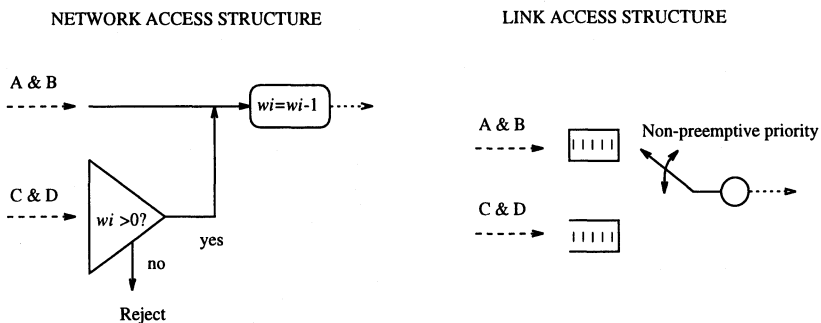


Figure 2 Selective window control with a priority mechanism

At each node, the value of counter w_i is initially set such that $\sum_i w_i = W$. It is assumed that an appropriate preventive control is implemented for isochronous traffic cells (Classes A and B in ITU-T standards). The selective window control is transparent for the isochronous traffic, but each admitted cell at node i decreases by 1 the value of the counter w_i . A delay-insensitive cell (Classes C and D) is admitted only if the counter is positive, hence this scheme adjusts the level of data traffic as a function of the level of isochronous traffic. The theoretical studies and simulation results presented in [4] have

shown that the optimal window size is independent of the level of isochronous traffic. The link access structure consists of two buffers, one for each type of traffic and the service discipline gives non-preemptive priority to isochronous traffic. Simulation results have also shown that the window size has a negligible effect on the performance of the isochronous traffic because of this priority scheme. The network efficiency is further improved by using finite buffers for the isochronous cells. This prevents using the network resources to carry isochronous cells with excessive delay.

Since the selective window controls only the delay-insensitive traffic and since the priority scheme assures that the window control has a negligible effect on the isochronous traffic, we will only consider in this paper the performance of the delay-insensitive traffic as measured by the product of powers.

4 DECENTRALIZED CONTROL ARCHITECTURE

Some practical problems arise when one wants to implement a centralized control scheme in a large-scale network. On the one hand, it implies linking all the nodes to a single controller which can become fairly impossible when the number of nodes explodes. When considering a window-based control scheme, such as the centralized isarithmic control scheme described in Section 3.1 the large window size means significant control overhead. In this section we propose a partially decentralized control architecture that extends the centralized isarithmic approach and we study some design issues that are related to it.

The network management function in most existing and planned networks is either hierarchical or completely centralized. The method we proposed here assumes an hierarchical zoning structure, where each zone is a subnetwork of the given network. In each zone, one node serves as a gateway to communicate with the other zones. In the hierarchical structure considered here, a low-level zone is connected to one higher-level zone only. The number of levels as well as the number of zones connected to a higher-level zone are general. An adaptive isarithmic controller is implemented in each zone with a distinct permit class within its boundary.

The permits for a given zone must circulate in that zone only. This implies that for cells crossing zone boundaries at the gateway nodes, the attached permits are routed to the associated controller when the cells arrive at the gateway of the origin subnetwork. There is no access control for cells arriving at a gateway from another subnetwork. They are treated as endogenous arrivals at the gateway. In other words, when a cell gains access to a subnetwork it also has access to the entire network.

One attractive advantage of the centralized adaptive isarithmic control is that it automatically ensures that the global information flow, necessary for fair-efficient operation, is available without the need for further protocols and interfaces. Under the decentralized control architecture described above, management information flows must circulate between subnetworks to achieve the

network optimal operating point. In each zone, the adaptive isarithmic controller will adjust the window size and the permit distribution probabilities in order to maximize the product of powers of the entire network. The permit flows in one zone provide the associated controller with the measurements (delays and throughputs) for the traffic originating and terminating in that zone. The management information flows arriving from other zones will provide the other measurements required.

We will now concentrate on the window size adjustment which requires a single number, i.e. the product of powers of all $o-d$ pairs in the network. The permit distribution adjustment problem under the decentralized control architecture will be addressed in a further study.

4.1 Information structures

Different information structures correspond to what delay and throughput information is made available to the adaptive controllers. There are three basic cases, namely *local* (o and d in same zone), *exact global* and *approximate global* (o and d in all zones). Under the local information structure, only the $o-d$ pairs originating and terminating in a given zone are considered by the adaptive isarithmic controller to adjust the zone's window size. The performance criterion is not the product of powers of the complete network thus it implies suboptimality (see Section 4.2 below). In the exact global information structure, all necessary information is passed across zone boundaries in order to achieve the optimal performance for the complete network. Approximations of the delay measurements are employed in the approximate global information structure to reduce the amount of management information flows circulating between the zones.

The product of powers of the complete network requires the knowledge of the throughputs and delays of all $o-d$ pairs in the network. Each controller needs this information to adjust the number of permits in its zone. At first glance, one would think that all these parcels of information must circulate from one zone to the others but the nature of the performance criterion avoids this situation. In fact very low information overhead is required.

A partial product can be computed in a zone for all $o-d$ pairs originating in the zone using throughput and delay information available locally while management information flows provide the delays measurements for the pairs terminating remotely. The product of powers of all $o-d$ pairs originating in a zone is a single number per zone. The zone's controller multiplies this number with the partial products received from the other zones to obtain the network product of powers. A given zone z receives one partial product from the zone situated above. This partial product is the product of powers of all $o-d$ pairs originating in this part of the hierarchy. Zone z also receives one partial product from each zone below. In return the given zone transmits to

the higher-level zone one partial product for the $o-d$ pairs originating in zone z and in the zone beneath. It also sends a partial product to each lower-level zone. In short, one number is sent by a zone to the zone above it and one number is sent to each zones below it.

Apart from this information, the management flows carry the delay measurements of the pairs that terminate remotely. Under the exact global information structure the number of delay values sent by zone z of level l to the zone of level $l+1$ is:

$$N_z(N - N_z) \quad (2)$$

where N_z is the total number of nodes in zone z and in the zones under it, and N is the total number of nodes in the network. The higher-level zone sends the same amount of information to zone z . As we move up in the hierarchy, the number of values to exchange increases which can become an implementation problem. The approximate global information structure is proposed to avoid it. Rather than carrying all the delay measurements, the management flows could carry mean delays only. Each zone computes the arithmetic mean of the delays between the gateway and each node in the zone. This value is made available to the other zones via the management flows which also carry the gateway-to-gateway delays. The amount of information grows linearly with the network's size instead of exponentially.

4.2 Simulation results

In this section, we present simulation results for an ATM network operating under the decentralized control architecture. We will compare the network performance for each information structure described above (local, exact global and approximate global). We also compared it to the performance under the centralized isarithmic control, which corresponds to one zone spanning the whole network.

The topology of the simulated networks $N1$ and $N2$ is depicted at Fig. 3. It has one level consisting of two zones (subnetworks A and B) with a gateway in each and a connection between gateways. The capacity of each unidirectional link is $1.7Gbits/s$. The two networks have the following characteristics:

$N1$: All the links cover a distance of 13 Km.

$N2$: Links $A1-A2$ and $B1-B2$: 16.7Km, links $A1-A4$ and $B1-B4$: 9.3Km (all other links remaining at 13Km).

The two traffic types simulated are the isochronous traffic (Type 1) and the delay-insensitive data traffic (Type 2). The isochronous traffic (Classes A and B in ITU-T standards) is modeled by two-state Markov Modulated Poisson Process (MMPP) which approximates the superposition process of bursty sources. The MMPP parameters are the arrival rates during states 1 and 2 ($\lambda_{1,1}$, $\lambda_{1,2}$) and the mean sojourn times in states 1 and 2 ($r_{1,1}^{-1}$, $r_{1,2}^{-1}$).




Figure 3 Simulated network topology

In the rest of this paper, the time unit is the cell transmission time. The arrival rates are:

$$\lambda_{1,1}(i, j) = \begin{cases} 0 & i = j \\ 0.16 & i, j | i, j \in \text{subnetwork A} \\ 0.16 & i, j | i, j \in \text{subnetwork B} \\ 0.08 & \text{otherwise} \end{cases}$$

$$\lambda_{1,2}(i, j) = \begin{cases} 0 & i = j \\ 0.04 & i, j | i, j \in \text{subnetwork A} \\ 0.04 & i, j | i, j \in \text{subnetwork B} \\ 0.02 & \text{otherwise} \end{cases}$$

The total offered isochronous traffic to each node is $\Lambda_1(i) = 0.5$ cell per cell transmission time and overall, $\Lambda_1 = 4$. All sources have $r_{1,1} = r_{1,2} = 0.005$. Due to the selective window control with the priority mechanism, we are assured that no isochronous cell will be refused. The residual bandwidth will be allocated to the delay-insensitive data traffic. In the simulation, the finite buffers for isochronous traffic have 100 waiting places.

The data traffic (Classes C and D in ITU-T standards) is assumed to arrive according to Poisson process. The arrival rates for Type 2 traffic are:

$$\lambda_2(i, j) = \begin{cases} 0 & i = j \\ 0.225 & i, j | i, j \in \text{subnetwork A} \\ 0.225 & i, j | i, j \in \text{subnetwork B} \\ 0.1125 & \text{otherwise} \end{cases}$$

which implies $\Lambda_2(i) = 1.125$ and $\Lambda_2 = 9$.

Table 1 presents the simulation results obtained with the decentralized implementation for each information structure (exact global, approximate global, and local). During each simulation, the window sizes in zones A and B, denoted W_A and W_B , are identical. We report the network performance, as measured by the product of powers, for a range of window sizes. Because of the magnitude of the product of powers its n^{th} root is used, where n is the number of o - d pairs in the network (here $n = N(N - 1) = 56$). The per-

Case	W_A, W_B	Exact global	Approx. global	Local
		$\sqrt[56]{P}$	$\sqrt[56]{\tilde{P}}$	$\sqrt[12]{P_{A,loc.}}$
N1	1200	2.42×10^{-5}	2.41×10^{-5}	5.43×10^{-5}
N1	1230	2.64×10^{-5}	2.63×10^{-5}	6.15×10^{-5}
N1	1260	2.87×10^{-5}	2.87×10^{-5}	6.83×10^{-5}
N1	1290	2.99×10^{-5}	2.99×10^{-5}	7.58×10^{-5}
N1	1320	2.79×10^{-5}	2.78×10^{-5}	8.41×10^{-5}
N1	1350	1.47×10^{-5}	1.47×10^{-5}	9.25×10^{-5}
N2	1200	2.45×10^{-5}	2.43×10^{-5}	5.51×10^{-5}
N2	1230	2.67×10^{-5}	2.66×10^{-5}	6.24×10^{-5}
N2	1260	2.90×10^{-5}	2.89×10^{-5}	6.92×10^{-5}
N2	1290	2.98×10^{-5}	2.97×10^{-5}	7.70×10^{-5}
N2	1320	2.86×10^{-5}	2.86×10^{-5}	8.58×10^{-5}
N2	1350	1.47×10^{-6}	1.47×10^{-6}	9.41×10^{-5}

Table 1 Simulation results with decentralized control architecture

formance measure is related to the information structure considered. For the exact global information structure, we have the exact product of powers P , whereas \tilde{P} is computed under the approximate global information structure. In the last column, we report $P_{A,loc.}$ the product of powers of $o-d$ pairs that originate and terminate in zone A. (Due to the symmetry, the local value for zone B is similar.)

For both $N1$ and $N2$, the optimal operating point is $W_A^* = W_B^* = 1290$ since the product of powers P is then maximized. The simulation results show that the approximate product of powers \tilde{P} is a very good estimate of the exact product of powers. Under the approximate global information structure the optimal operating point is also $W_A^* = W_B^* = 1290$. We can see that the product of powers is slightly underestimated when the window sizes are under their optimal values. However the results obtained for the local information structure are not satisfactory. The local product of powers for zone A indicates that the optimal operating point would be achieved for very large window sizes. But if we consider $W_A = W_B = 1350$, we can observe that the exact global product of powers is not optimal. The local information structure considers only 12 $o-d$ pairs thus it is an inadequate performance estimate.

While it was found by simulation that the optimal operating point under centralized control is $W^* = 3300$ for $N1$ and $N2$, the optimal window size for the decentralized control is $W_A^* + W_B^* = 2580$. This reduction in the total population of permits is possible because permits are returned to their controller when the cells cross zone boundaries. To conclude, these simulation results show that the decentralized approach is very attractive. The total window

size is smaller and the approximate global information structure allows very low information overhead without impairing the performance.

5 CREDIT/PERMIT SCHEME

In this section we present our second approach to implement the isarithmic scheme in large-scale networks. This approach is named the credit/permit scheme. In order to reduce the flow of permits in the network, as well as to reduce the real time processing load at the adaptive controllers one could let a permit correspond to C credits. When a permit arrives to a source queue it is replaced by C credits where an admitted cell decreases by 1 the counter of credits, as previous. At the exit nodes each cell adds 1 to a counter of released credits and when C credits are accumulated a permit is sent to the adaptive controller for dispersal. The action rate of the controllers is then reduced by a factor of C as well as the signaling traffic (i.e. permits).

Because of the selective window control with priority, the credit/permit scheme does not affect the isochronous traffic and simulations were performed to assess its impact on the data traffic. The simulated network is identical to one of the subnetworks illustrated in Fig. 3 and all the links cover $13Km$. As in the previous case, the isochronous traffic is modeled by an MMPP and the data traffic by a Poisson process. The total offered traffics are: $\Lambda_1 = 2$ and $\Lambda_2 = 4.5$. Under the centralized isarithmic control, the optimal operating point of the network is $W^* = 2350$ and $\sqrt[12]{P^*} = 4.76 \times 10^{-4}$.

Two implementations of the credit/permit scheme have been simulated. In version A, when a permit arrives at a source queue, it is immediately replaced by C credits. In version B, the credits are generated in a smooth manner that will be described below. Results for both implementation are plotted at Fig. 4.




Figure 4 Version B: Product of powers versus credit/permit ratio

In this version, the network performance is even better than the original isarithmic scheme for small values of the credit/permit ratio. Since the credits

are generated smoothly, the bursts of data cells are avoided. Moreover, the data traffic is shaped such that its mean delay is smaller. Again when C is large, the throughput of data traffic decreases and this degradation is more important than with version A. But in terms of the product of powers, version B outperforms version A due to the low delays involved.

6 CONCLUSIONS

This paper proposed two approaches to facilitate network access control in large-scale ATM networks. We consider the selective window control scheme with priority mechanism. It can be difficult to implement this window-based scheme in large-scale networks because of the control overhead associated with the flow of permits.

First, we proposed a partially decentralized control architecture which extends the centralized isarithmic scheme by means of an hierarchical zoning structure. An adaptive isarithmic controller is located in each zone with a distinct permit class and the zones interact through user traffic and management information flows. It was demonstrated that the decentralized architecture allows a substantial reduction of the flow of permits. Numerical results showed that the network optimal operating point can be reached under the decentralized implementation. Moreover, the level of management flows can be limited without deterioration of the global performance by employing the approximate global information structure, based on mean delays.

The second approach proposed is the credit/permit scheme that reduces the control overhead by a factor of C , the credit/permit ratio. This second approach can be used in conjunction with the first approach to further reduce the management flow overhead.

The simulation results presented here were for a hypothetical network example with a topology that lends itself naturally to a hierarchical zoning structure. The results have shown the validity of the approach in this simple case. The methodology, proposed for decentralization of access flow control is however not restricted to the simple case simulated, and results for more general topologies are currently under study and will be reported upon in due course. The merit of the proposal we believe is chiefly that it provides a potential alternative access flow control scheme to the various VC based ABR proposals for efficient transparent support of connectionless data services in a multi-service ATM network, by means of an overlay virtual network as proposed in [1] and elaborated in [19].

REFERENCES

- [1] L.G. Mason, Z. Dziong, K.-Q. Liao, and N. Tétreault (1990) Control Architectures and Procedures for B-ISDN, in *Proc. of 7th ITC Spe-*

- cialists' Seminar, Morristown, New Jersey.*
- [2] H.T. Kung and R. Morris (1995) Credit-Based Flow Control for ATM Networks, *IEEE Network*, **March/April**, pp. 40–49.
 - [3] F. Bonomi and K.W. Fendick (1995) The Rate-Based Flow Control Framework for the Available Bit Rate ATM Service, *IEEE Network*, **March/April**, pp. 25–56.
 - [4] K.-Q. Liao and L.G. Mason (1994) A Congestion Control Framework for Broadband ISDN Using Selective Window Control, *IFIP Conf. on Broadband Communications (Broadband'94)*, **Paris**.
 - [5] D.W. Davies (1972) The Control of Congestion in Packet Switching Networks, *IEEE Trans. Commun.*, pp. 546–550.
 - [6] L.G. Mason and X.D. Gu (1986) Learning Automata Models for Adaptive Flow Control in Packet Switching Networks, in *Adaptive and Learning Systems – Theory and Applications*, **K.S. Narendra, New-York: Plenum Press**, pp. 213–227.
 - [7] F.J. Vazquez-Abad and L.G. Mason (1992) Decentralized Adaptive Isarithmic Flow Control for Packet-switched Networks, *Second ORSA Telecom. Conf., Boca Raton*.
 - [8] M. Cotton and L.G. Mason (1995) Adaptive Isarithmic Flow Control in Fast Packet Switching Networks, *IEEE Trans. Commun.*, **April**.
 - [9] E. Létourneau and L. G. Mason (1995) Comparison Study of Credit-Based and Rate-Based ABR Control Scheme, in *Proc. of the First IFIP Workshop on ATM Traffic Management, Paris, December*.
 - [10] T. Toniatti and F. Trombetta (1992) Performance simulation of end-to-end windowing in ATM networks, in *Proc. of INFOCOM*, pp. 495–502.
 - [11] L. Kleinrock (1979) Power and Deterministic Rules of Thumb for Probabilistic Problems in Computer Communications, in *Proc. Int. Conf. on Commun.*, pp. 43.1.1–43.1.10.
 - [12] K. Bharath-Kumar and J.M. Jaffe (1981) A New Approach to Performance Oriented Flow Control, *IEEE Trans. Commun.*, pp. 427–435.
 - [13] R. Mazumdar, L.G. Mason and C. Douligeris (1991) Fairness in Network Optimal Flow Control: Optimality of Product Forms, *IEEE Trans. Commun.*, pp. 775–782.
 - [14] L.G. Mason (1973) An Optimal Learning Algorithm for S-Model Environments, *IEEE Trans. on Automatic Control*, pp. 493–496.
 - [15] S. Pingalu, D. Tipper and J. Hammond (1990) The performance of adaptive window flow controls in a dynamic load environment, in *Proc. of INFOCOM*, pp. 55–62.
 - [16] D. Mitra and J.B. Seery (1991) Dynamic Adaptive Windows for High Speed Data Networks with Multiple Paths and Propagation Delays, in *Proc. of INFOCOM*, pp. 39–48.
 - [17] L.G. Mason and K.-Q. Liao (1993) Self-Optimizing Window Flow Control in High Speed Data Networks, *Proc. of IEEE Globecom 1992*, also *Computer Communications*, **Vol. 16, No. 11**, pp. 706–716.

- [18] A. Pelletier, M. Cotton and L.G. Mason (1993) Combined Adaptive Routing and Flow Control in Fast Packet Switching Networks, *in Proc. of COMCON 4*, pp. 557–569
- [19] Z. Dziong, Y. Xiong, and L.G. Mason (1996) Virtual Network Concept and its Application for Resource Management in ATM Based Networks, *in Broadband Communications, Chapman & Hall, Lorne Mason and Augusto Casaca, editors*, pp. 223–234.