A Distributed Routing and Access Control Scheme for ATM Networks

R. Bolla, F. Davoli, M. Marchese

Department of Communications, Computer and Systems Science (DIST)
University of Genoa
Via Opera Pia, 11/A - 16145 Genova, Italy

ABSTRACT. A proposed model of an ATM access node is extended to describe a network featuring a node by node admission control and routing strategy. A hierarchical dynamic resource allocation scheme is periodically applied at each node to every outgoing link, in order to reassign link capacity partitions among various service classes. Local access control rules compute the maximum number of connections of each class that a link can accept, by using the assigned capacity. A call request packet is generated with every connection request, and it is sent along a path from the source to the destination node. Each node traversed checks the presence of the resources needed to accept a new connection, by using the local access control rule. A connection is accepted only if every node on the path has enough resources to support it (also maintaining the quality of service for connections in progress), and it is refused otherwise. The path of the request packet is chosen by using a distributed routing strategy. At every node traversed, an output link is chosen by minimizing a cost function composed by two terms: a "local" and a "global" one. The local term takes into account the situation of the node by using a simple calculation, whereas the global term holds the aggregate information about the situation of the network, and it is updated and passed along periodically by every node. A complete description of the global scheme and simulation results are reported and discussed.

1. Introduction

The use of the Asynchronous Transfer Mode (ATM) technique [1] requires a considerable amount of real time control in order to guarantee Quality of Service (QoS) requirements to the users. This is mainly due to the presence of non homogeneous traffic flows with different statistical nature and performance characteristics (narrowband, broadband, time- or loss-sensitive, etc.), generating cells that are mixed through statistical multiplexers. The statistical allocation of the resources in this heterogeneous environment gives rise to some challenging network control problems, like bandwidth allocation and access control, which have been widely investigated (see, for instance, [2-8] and [9-19], respectively).

Whereas the initial idea underlying the development of ATM has been that of a complete statistical multiplexing, control architectures oriented towards a more "structured" resource allocation have also been considered, where portions of a resource (e.g., bandwidth) can be temporarily and dynamically assigned to traffic with specific characteristics or requirements. An architecture for a more structured logical partition of network resources among traffic classes with different grade of service has been proposed for instance in [20]; another example of the application of bandwidth allocation in the same sense can be found in [17]. In this respect, the subdivision of the incoming traffic into different classes, characterized by specific properties, as has been also recently recognized by the CCITT in the definition of the BISDN Service Classes [21].

A related approach, somehow connected also to the line of reasoning developed in [22-23] in the context of TDM systems, has been presented by the authors in [24-26], with respect to the control of a single ATM link. In the present paper, our goal is to extend the model for an ATM access node proposed therein, to describe a complete network from the point of view of access control, bandwidth allocation, and routing.

In the model considered in [24-26], several traffic classes share an ATM link. Each traffic class is characterized by statistical parameters (like peak and average bandwidth), as well as by QoS requirements at the cell level (in terms of cell loss probability and cell delay). A separate call admission controller is dedicated to each different traffic class; this device applies a "local" fixed strategy, designed to maintain the required QoS, given the buffer space and bandwidth (percentage of cells) assigned to the class. The bandwidth shares are periodically recomputed on-line by a bandwidth allocation controller that plays the role of a coordinator in a hierarchical dynamic control scheme. This scheme has been considered with respect to a single output link in a node entered only by external incoming traffic.

In general, a real network node has several incoming and outgoing links, and traffic enters the node both from local external sources and from other nodes through the incoming links. Thus, in our network model, every node is assigned a control structure, like the one described above, for each of its outgoing links. The main problem here is related to the fact that a global acceptance rule must be implemented, in order to check if all the nodes to be

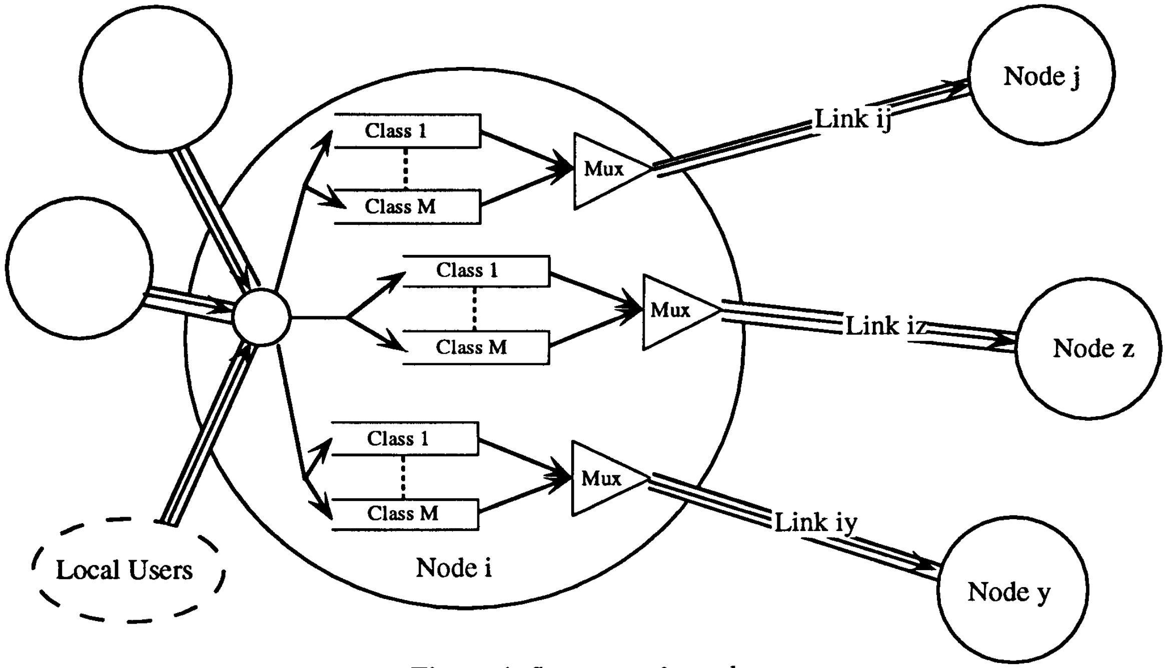

Figure 1 *Structure of a node.*

traversed from source to destination have enough resources to serve an incoming connection, while still guaranteeing QoS to the connections in progress. Two aspects of this problem must be considered. The first one is that we would need to know which path will be followed by the Virtual Circuit (VC) generated by a connection, to find out the nodes where to check the availability of the resources: in this respect, we have a routing problem. The second aspect is related on how to check the different nodes on the path.

We face this second problem by following a technique suggested in [19]. A call request packet, generated by the source node in reaction to a connection request, is transmitted along a path to the destination, and at each node traversed it presents a connection request to the local access control. If the request is accepted, an outgoing link is chosen according to the routing strategy, the packet is forwarded to the next node on the path and repeats the same action; if the request is refused or the destination node has been reached, the packet is sent back to the source node. Globally, a connection is accepted in the network and the corresponding VC is created if every node up to the destination accepts the new connection request.

As regards routing, we propose a distributed strategy, whereby every node decides the output link which the connection request packet must be sent to. The output link is chosen by minimizing a cost function composed by two terms: a "local" and a "global" one. The local term takes into account the situation of the node by using a simple function, whose value depends on the number of connections that every link could accept at the moment the decision is made. The global term holds aggregate information about the situation of the network, and it is updated periodically by every node and passed along to its neighbours.

The paper is organized as follows. In the next Section, we describe the structure of the ATM node. In Section 3, we consider the whole network, and present the global access control rule and the routing strategy. Simulation results are reported and discussed in Section 4 and conclusions are drawn in Section 5.

2. Structure of the ATM node and local access control rule

We suppose the traffic on the network to be divided into M classes of service, each one characterized by statistical parameters like peak and average transmission rate, as well as by QoS requirements, like cell loss probability and cell delay.

The general structure of a node is depicted in Fig. 1. Every node has a certain number of incoming links from other nodes and a local traffic input link. Moreover, every node has a certain number W of outgoing links, each of which receives a part of the global traffic entering the node. A resource allocation scheme and a local control access rule of the type presented in [24-26] are implemented for every outgoing link, and we summarize them briefly in the following.

The overall control structure of every link is shown in Fig. 2. As can be seen, the input traffic is divided into classes and a separate buffer is assigned to each class. The output of each buffer is statistically multiplexed on the outgoing link by a

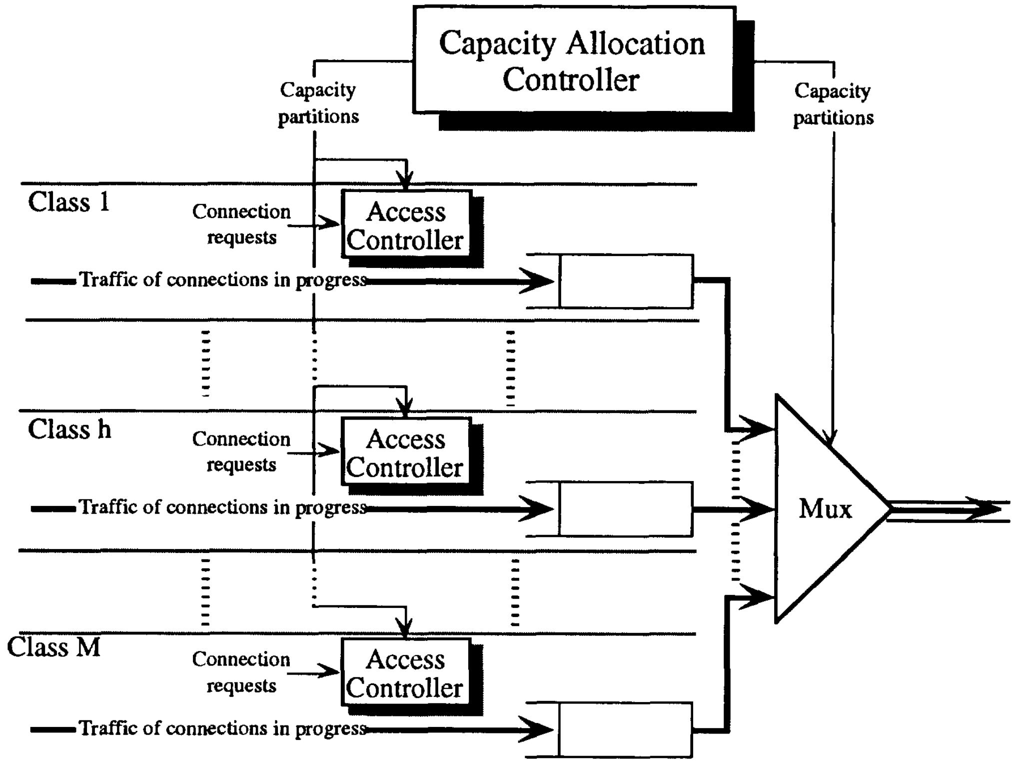

Figure 2 *The overall control structure of a link.*

scheduler that substantially divides the global channel capacity C_T among the classes. The connection requests, which can come from the users directly connected to the node or from other nodes by means of special call request packets (as will be shown in the next Section), are divided into classes, too. Two controls are exerted on the system, by using a two level hierarchical control scheme, one acting on the scheduler and the other on the admission of the connections.

Time is discrete, and a time slot equals a cell transmission time. At the higher level of the hierarchical control scheme, a bandwidth allocation controller periodically reassigns capacity partitions to every class. The scheduler receives the values of the partitions and must assure that every buffer is assigned a percentage of slots equal to the ratio between the total capacity and the capacity assigned to its class. The new capacity partitions $V_m^{(h)}$, $h = 1,\dots, M$, are computed by the controller at discrete time instants $m = 0, K, 2K, \dots$ (K is the length of the intervention period in slots), based on the minimization of a cost function that takes into account the overall cell loss probability.

At the lower level, an access controller for every class decides about the acceptance of a connection request. The decision is taken on the basis of the "virtual" capacity allocated to the class, of the current number of active connections and of the statistical and performance characteristics of the specific traffic.

Class h traffic is supposed to be made up by bursty connections with identical and independent statistical characteristics. We indicate with $b^{(h)}$, $B^{(h)}$, and $P^{(h)}$ the burstiness, the average burst length, and the peak bit rate, respectively, of the h-th class. We model the connection simply as a two state (idle and active, respectively) Markov chain. When a connection is in the active state, it can generate a cell per slot with a certain probability (equal to $P^{(h)}/C_T$) , and it does not generate cells when it is in the idle state.

By assuming the quasi-stationarity of the connections activity process, the packet loss rate for the connections of one class can be computed. The packet loss rate depends on the capacity partitions and on the number of connections in progress on the specific link. The aim of the allocation controller is that of balancing the bandwidth distribution, by setting the values $V_m^{(h)}$ at the beginning of a new K-slot interval, on the basis of a performance measure, which should be capable of reflecting changes in the activity of all traffic classes. We intend as the activity of a traffic class in a certain period of time the total number of service requests, including both accepted and blocked ones. In this respect, we have chosen to use a cost function which reflects the steady-state value of the cell loss rate for the h-th class, due to the connections present in the system, as well as the additional loss that would have been incurred if all calls presented in the previous interval had been accepted. An equality constraint (the sum of the assigned capacity must be equal to the total capacity of the link) and a set of inequality constraints, which assure service quality for the connections already in progress, must be taken into account in the minimization procedure. The latter can be treated as a mathematical programming problem, which is solved by using a gradient projection method.

After every access controller receives its capacity assignment from the bandwidth allocation controller, it computes the maximum number of connections of the h-th class that can be supported on link ij by means of

$$N_{ij}^{(h)}(m) = \min\left\{N_{ij,L}^{(h)}(m); N_{ij,D}^{(h)}(m)\right\} \quad (1)$$

where $N_{ij,L}^{(h)}(m)$ is the maximum number of connections on the link capable of maintaining the cell loss rate below a given upper bound $\varepsilon^{(h)}$, and $N_{ij,D}^{(h)}(m)$ is the maximum number of connections computed by imposing a similar limit on cell delay, namely

$$\Pr\left\{\text{delay of a cell} > D^{(h)}\right\} < \delta^{(h)} \quad (2)$$

where $D^{(h)}$ is a delay requirement for the h-th class (the details of this computation can be found in [26]). Summing up, the local acceptance rule is: a new connection of the h-th class arriving at time slot k, $m+\Delta \leq k \leq m+\Delta+K-1$ (where Δ indicates the number of slots required for the computations), can be accepted on link ij if

$$N_{ij,A}^{(h)}(k) + 1 \leq N_{ij}^{(h)}(m) \quad (3)$$

where $N_{ij,A}^{(h)}(k)$ is the number of connections of the h-th class in progress (i.e., previously accepted on the link and not terminated) at time slot k.

It is worth noting that all computations involved, i.e. the calculation of the capacity partitions and of $N_{ij,L}^{(h)}(m)$, $N_{ij,D}^{(h)}(m)$, are performed only at the beginning of every decision interval, whose duration should be quite long compared to the slot, because the cost function reflects the dynamics of the connection requests, which take place on a much longer time scale. Thus, if the number of time slots Δ to perform the computations after the reallocation instants represents a small fraction of K, we can avoid the use of special approximations to speed up the computations. On the other hand, the acceptance rule is very simple and then it can react very fast, during the decision interval, to the connection requests.

As already mentioned, a detailed description of the allocation and admission control rules has been presented in [24-26]; in particular, we have used in the following the admission rule described and tested in [26].

3. A routing strategy

Let us suppose that a route has already been found; a scheme to verify the availability of this route is the node-by-node control suggested in [19]: we send forward a special call request packet which is just a messenger asking whether a connection can be accepted or not. The packet runs through every node of the route and, at each one, it undergoes the access control rule. Each node must take account of the resources it has allocated; if the call is accepted the messenger runs along the route, otherwise, the packet stops, the call is rejected and a message is sent back, through the same route, so the intermediate nodes can free the allocated resources.

Node-by-node control and routing can be joined and managed together: in our approach the best route is not chosen beforehand and then verified but, at each node, we choose the "best" outgoing channel by means of a cost function associated with each link. The cost function should take into account the link's local traffic and the traffic associated with the subsequent hops to the destination, and it should be considered only for those links that are not congested (in the sense defined by the admission control rule).

Inside the messenger there is a bit stream, which is used to remember the nodes traversed along the route. At each node, a table look-up is performed to find out congested channels, the lowest cost function link among the non-congested ones is chosen, and the packet is sent along that route. If every channel is congested in an intermediate node, the connection is rejected and a message is sent back through the same route, so resources can be released; otherwise, if the messenger reaches its destination, the connection is accepted, a VC is established, and the path that the cells of that connection will follow is fixed.

At instant $\tilde{k}$ (expressed in slots), a generic node i chooses the link to which to forward a call request packet generated by a class h connection request, by minimizing (over all successor nodes j) the quantity

$$c_{ij}^{(h)} = c_{ij,L}^{(h)}(\tilde{k}) + c_{j}^{(h)}(\tilde{s}) \quad (4)$$

where $c_{ij,L}^{(h)}(\tilde{k})$ is a local cost related to link ij and $c_{j}^{(h)}(\tilde{s})$ is an aggregate cost referring to the traffic conditions of node j and its successor at instant $\tilde{s} < \tilde{k}$ (expressed in slots). $c_{ij,L}^{(h)}(\tilde{k})$ should weigh the local congestion of link ij. A possible proposal, which we have used in the simulations in the next Section, can be

$$c_{ij,L}^{(h)}(\tilde{k}) = $$

$$= \begin{cases} \dfrac{1}{N_{ij}^{(h)}(\tilde{m}) - N_{ij,A}^{(h)}(\tilde{k})} & \text{if } N_{ij}^{(h)}(\tilde{m}) > N_{ij,A}^{(h)}(\tilde{k}) \\ Z & \text{if } N_{ij}^{(h)}(\tilde{m}) = N_{ij,A}^{(h)}(\tilde{k}) \end{cases} \tag{5}$$

where $\tilde{m} \le \tilde{k} < \tilde{m}+K$, i.e., $\tilde{m}$ is the reallocation instant when the access control rule parameters, active at the time of the request packet's arrival, have been recomputed. By using this type of function, we have that the link cost increases with decreasing available space on the link, expressed in number of acceptable connections. When the link is saturated, the cost value is Z, which should be high enough to ensure that no saturated link will be chosen if non-congested links are available.

The cost referred to a generic node i is composed by two terms

$$c_i^{(h)}(\tilde{s}) = \alpha_i c_{i,L}^{(h)}(\tilde{s}) + \beta_i c_{i,A}^{(h)}(\tilde{s}) \tag{6}$$

where α_i and β_i are two weighting coefficients. $c_{i,L}^{(h)}(\tilde{s})$ represents the average situation of the node with respect to its congestion state, and $c_{i,A}^{(h)}(\tilde{s})$ is an aggregate information on the average congestion of its successor nodes. More specifically, we have defined

$$c_{i,L}^{(h)}(\tilde{s}) = \frac{1}{L_i} \sum_{j \in Succ(i)} c_{ij,L}^{(h)}(\tilde{s}) \tag{7}$$

$$c_{i,A}^{(h)}(\tilde{s}) = \frac{1}{L_i} \sum_{j \in Succ(i)} c_j^{(h)}(\tilde{s}) \tag{8}$$

As can be seen, the values related to the successor nodes are referred to the instants $\tilde{s}$, where $\tilde{s}$=T, 2T, ..., with T equal to a fixed number of slots. This means that each node i sends its costs $c_i^{(h)}(\tilde{s})$, h=1,...,M, to its predecessors every T slots and then, after receiving the costs from its successors, recomputes its new aggregate information on the congestion of the network. The passage of the parameters is "one step", in the sense that the cost sent back is computed before receiving the new values from the successor nodes.

It is worth noting that the proposed strategies are distributed, based on a mix of local real time (dynamic) and overall delayed aggregate information, and do not require the presence of a real time supervisory controller, which would be questionable in a wide area network.

4. Simulation results

In this Section, we show and briefly comment some results of simulations that were used to test the correctness and the performance of the proposed routing and admission control scheme. The following data have been used:

C_T = 150 Mbit/s; M =3; K = $2 \cdot 10^6$ cells
T_s = slot duration = $4.24 \cdot 10^{-6}$ s (53 bytes/cell)
T = K = $2 \cdot 10^6$ cells
$P^{(1)}$ = 384 kbit/s; $P^{(2)}$ = 1 Mbit/s; $P^{(3)}$ = 10 Mbit/s
$b^{(1)} = 2$; $b^{(2)} = 5$; $b^{(3)} = 10$
$B^{(1)} = 100$; $B^{(2)} = 1000$; $B^{(3)} = 1000$ cells (average burst length)
$1/\mu^{(1)} = 1.5$ s; $1/\mu^{(2)} = 0.8$ s; $1/\mu^{(3)} = 1$ s (average connection duration)
$\varepsilon^{(1)} = \varepsilon^{(2)} = \varepsilon^{(3)} = 1 \cdot 10^{-4}$
$\delta^{(1)} = \delta^{(2)} = \delta^{(3)} = 1 \cdot 10^{-3}$
$D^{(1)} = 10$; $D^{(2)} = 70$; $D^{(3)} = 1000$ slots
$N_a^{(1)}$=200; $N_a^{(2)}$=200; $N_a^{(3)}$=40 Erlangs (average traffic intensities)
$Q^{(1)} = 11$; $Q^{(2)} = 12$; $Q^{(3)} = 12$ cells
$\alpha_i = \beta_i = 1, \ \forall\, i$

We refer to the traffic flow generated by the above data as an offered load of 1; an offered load of "x" corresponds to the same data except for the traffic intensities $N_a^{(h)}$, h=1, 2, 3, which are multiplied by x.

We should remark that some of the above values do not correspond to a real situation, but were chosen especially to limit the length of the simulation runs necessary to obtain a significant number of events. This is true, in particular, for the average duration of the connections. However, one of the main purposes of the scheme is that of coping, to a certain extent, with dynamic variations in the call processes; in our case of relatively short connections, this is achieved by keeping the reallocation interval K also relatively short. In case of longer duration of the connections (with the same values of traffic intensity), the situation would be substantially unchanged by correspondingly enlarging the reallocation interval. We may note, in passing, that a larger value of K renders the computing time for the execution of the reallocation algorithm less critical.

The topology of the network we have used in our simulations is shown in Fig. 3 and is composed by

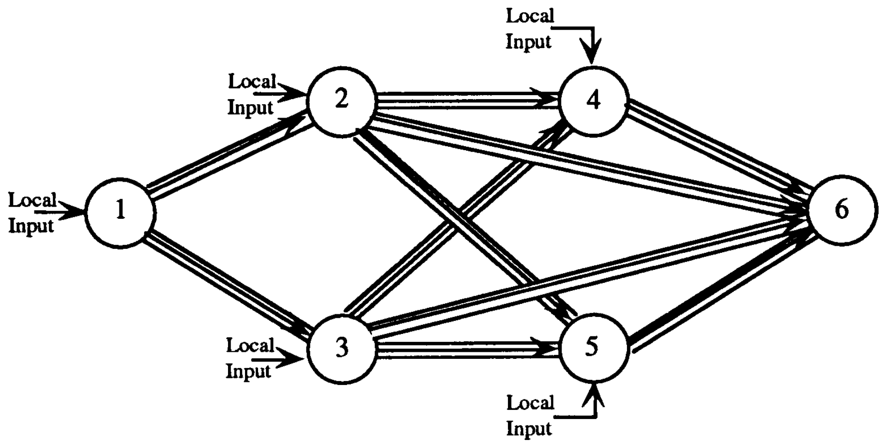

Figure 3 *Topology of the tested network.*

six nodes, only one of which (node 6) is a destination node.

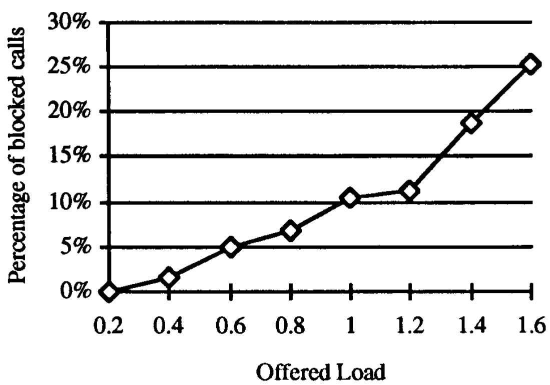

Figure 4 *Total percentage of blocked calls in the network versus the offered load.*

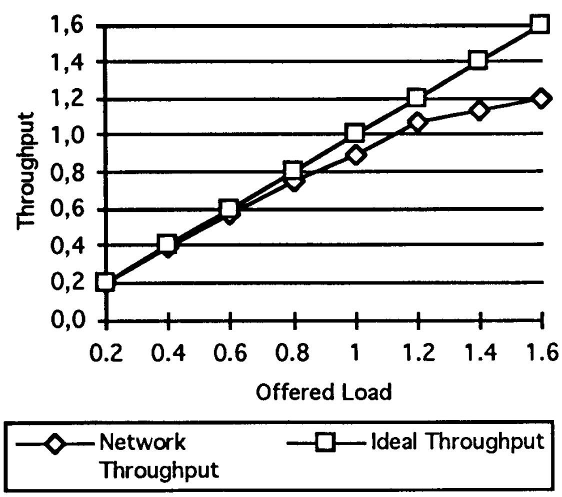

Figure 5 *Total throughput of the network versus the offered load.*

The behaviour of the access control and bandwidth allocation procedure at a node was tested extensively in [26]. Here, we concentrate on the overall network behaviour, and particularly on the performance at the call level.

All the simulations have a duration of 84.8 s, corresponding to 10 reallocation intervals.

Fig. 4 shows the percentage of blocked connection attempts versus the total offered load. The increase is rather limited, until the network becomes heavily loaded. The effect can be better appreciated from Fig. 5, where the normalized network throughput is plotted as a function of the offered load, showing a linear increase very close to the ideal throughput characteristic up to the value 1.6.

5. Conclusions

A global control architecture for access control, bandwidth allocation and routing has been considered, in an ATM network environment. The traffic is organized into service classes, characterized by specific performance requirements, and bandwidth partitions are dynamically allocated among them by controllers assigned specifically to each link. Access control and routing are performed separately for each class on a hop-by-hop basis: the former is exerted by local controllers for each link, whereas the latter stems from a distributed procedure based on local (real time) as well as on aggregated (delayed) information. The routing strategy has been explicitly defined in the paper. Simulation results have been reported, which show a low call rejection rate as an overall effect of the control structure over a large range of network load values.

Acknowledgment

This work was supported by the Italian Ministry of University and Scientific and Technological Research (M.U.R.S.T.).

References

[1] CCITT Recommendation I.121 - Broadband Aspects of ISDN, Blue Book, vol. 3, fasc. III. 7, 1988.

[2] H.Ohnishi, T.Okada, K.Noguchi, "Flow control schemes and delay/loss tradeoff in ATM networks", *IEEE J. Select. Areas Commun.*, vol. 6, n. 9, 1988.

[3] J.Y.Hui, "Resource allocation for broadband networks", *IEEE J. Select. Areas Commun.*, vol. 6, n. 9, pp. 1598-1608, Dec. 1988.

[4] G.Gallassi, G.Rigolio, L.Fratta, "ATM: bandwidth assignment and bandwidth enforcement policies", *Proc. Globecom '89*, Dallas, TX, Nov. 1989, pp. 1788-1793.

[5] G.M.Woodruff, R.Kositpaiboon, "Multimedia traffic management principles for guaranteed ATM network performance", *IEEE J. Select. Areas Commun.*, vol. 8, n. 3, April 1990.

[6] M.Dècina, T.Toniatti, "On bandwidth allocation to bursty virtual connections in ATM networks", *Proc. ICC '90*, Atlanta, GA, April 1990, pp. 844-851.

[7] J.A.Suruagy Monteiro, M.Gerla, L.Fratta, "Statistical multiplexing in ATM networks", *Proc. 4th Internat. Conf. on Data Commun. Syst. and their Performance*, Barcelona, Spain, June 1990.

[8] S.Erfani, M.Malek, "Dynamic access capacity management in a multiservice packet-mode environment", *Proc. Workshop on Network Control and Management*, New York, 1990, pp. 277-285.

[9] R.Guérin, H.Ahmadi, M.Naghshineh, "Equivalent capacity and its application to bandwidth allocation in high speed networks", *IEEE J. Select. Areas Commun.*, vol. 9, n. 7, pp. 968-981, Sept. 1991.

[10] F.P.Kelly, "Effective bandwidths at multi-class queues", *Queueing Systems*, vol. 9, pp. 5-15, 1991.

[11] R.J.Gibbens, P.J.Hunt, "Effective bandwidths for the multi-type UAS channel", *Queueing Systems*, vol. 9, pp. 17-27, 1991.

[12] F.C.Schoute, "Simple decision rules for acceptance of mixed traffic streams", *Proc. ITC 12*, Torino, Italy, June 1988.

[13] T.Kamitake, T.Suda, "Evaluation of an admission control scheme for an ATM network considering fluctuations in cell loss rate", *Proc. Globecom '89*, Dallas, TX, Nov. 1989, pp. 1774-1780.

[14] A.Lombardo, S.Palazzo, D.Panno, "Admission control over mixed traffic in ATM networks", *Internat. J. of Digital and Analog Commun. Systems*, vol. 3, Aug. 1990.

[15] Z.Dziong, J.Choquette, K.Q.Liao, L.Mason, "Admission control and routing in ATM networks", *Comput. Networks and ISDN Syst.*, vol. 20, pp. 189-196, Dec. 1990.

[16] C.Rasmussen, J.H.Sørensen, K.S.Kvols, S.B.Jacobsen, "Source independent call acceptance procedures in ATM networks", *IEEE J. Select. Areas Commun.*, vol. 9, n. 4, pp. 351-358, April 1991.

[17] Y.Takagi, S.Hino, T.Takahashi, "Priority assignment control of ATM line buffers with multiple QOS classes", *IEEE J. Select. Areas Commun.*, vol. 9, pp. 1078-1092, Sept. 1991.

[18] T.Murase, H.Suzuki, S.Sato, T.Takeuchi, "A call admission control scheme for ATM networks using a simple quality estimate", *IEEE J. Select. Areas Commun.*, vol. 9, pp. 1461-1470, Dec. 1991.

[19] D.Hong, T.Suda, "Congestion control and prevention in ATM networks", *IEEE Network Mag.*, vol. 5, no. 4, pp. 10-16, July 1991.

[20] A.A.Lazar, A.Temple, R.Gidron, "An architecture for integrated networks that guarantees quality of service", *Internat. J. of Digital and Analog Commun. Systems*, vol. 3, pp. 229-238, July 1990.

[21] CCITT Recommendation I.211 - BISDN Service Aspects, 1990.

[22] M.Aicardi, R.Bolla, F.Davoli, R.Minciardi, "A parametric optimization approach to admission control and bandwidth assignment in hybrid TDM networks", *Internat. J. of Digital and Analog Commun. Systems*, vol. 6, pp. 15-27, Jan. 1993.

[23] R.Bolla, F.Davoli, "Dynamic hierarchical control of resource allocation in a broadband integrated services network", *Computer Networks and ISDN Systems*, vol. 25, no. 10, May 1993.

[24] R.Bolla, F.Davoli, A.Lombardo, S.Palazzo, D.Panno, "Adaptive access control of multiple traffic classes in ATM networks", *Proc. Globecom '91*, Phoenix, AZ, Dec. 1991, vol. 2, pp. 331-338.

[25] R.Bolla, F.Davoli, A.Lombardo, S.Palazzo, D.Panno, "Adaptive bandwidth allocation by hierarchical control of multiple ATM traffic classes", *Proc. Infocom '92*, Florence, Italy, May 1992, pp. 30-38.

[26] R.Bolla, F.Danovaro, F.Davoli, M.Marchese "An integrated dynamic resource allocation scheme for ATM networks", *Proc. Infocom '93*, San Francisco, CA, March 1993, vol. 3, pp. 1289-1297.