A Scheme for Real-Time Channel Establishment in Wide-Area Networks

Domenico Ferrari, FELLOW, IEEE, and Dinesh C. Verma

Abstract

Multimedia communication involving digital audio and/or digital video has rather strict delay requirements. A real-time channel is defined in this paper as a simplex connection between a source and a destination characterized by parameters representing the performance requirements of the client. A real-time service is capable of creating real-time channels on demand and guaranteeing their performance. These guarantees often take the form of lower bounds on the bandwidth allocated to a channel and upper bounds on the delays to be experienced by a packet on the channel.

In this paper, we study the feasibility of providing real-time services on a packet-switched store-and-forward wide-area network with general topology. We describe a scheme for the establishment of channels with deterministic or statistical delay bounds, and present the results of the simulation experiments we ran to evaluate it. The results are encouraging: our approach satisfies the guarantees even in worst-case situations, uses the network's resources to a fair extent, and efficiently handles channels with a variety of offered load and burstiness characteristics. Also, the packet transmission overhead is quite low, and the channel establishment overhead is small enough to be acceptable in most practical cases.

This research has been supported in part by the Defense Advanced Research Projects Agency (DoD), ARPA Order No. 4871, monitored by Naval Electronic Systems Command under Contract No. N00039-84-C-0089, the University of California with a MICRO Program grant, AT&T Bell Laboratories, Inc., Cray Research, Inc., Hitachi, Ltd., IBM Corporation, Ing. C. Olivetti & C., S.p.A., and the International Computer Science Institute. The views and conclusions contained in this document are those of the authors, and should not be interpreted as representing official policies, either expressed or implied, of any of the sponsoring organizations or of the U.S. Government.

1. Introduction

Real-time computer-computer communication, by which we mean computer communication with guaranteed performance, is expected to become a necessary feature of future networks, as the digital audio and digital video components of their multimedia traffic will require much tighter performance control than is exerted in current networks. Once networks start offering real-time services, we believe that the domain of the applications of these services will rapidly grow. Performance guarantees, along with other guarantees, are going to be an important option among those that will define the various grades of service most observers believe will be offered by the networks of the future.

Such guarantees will have to be provided not only by local-area networks, where solutions like the FDDI ring and its protocols are already being developed [Ross86], but also by wide-area networks: the need for multimedia conferencing services is only one of the many examples that can be invoked to demonstrate the validity of this statement. The available options for providing real-time services in a wide-area context are circuit switching, fast packet switching, and combinations of these two techniques. The former is not particularly convenient when the traffic includes (as we expect it will in most cases) a large data/text component [Harr80]. Fast packet switching (for example, ATM) is better suited to the type of traffic we are envisioning; however, to our knowledge and in its formulations to date, it does not guarantee the low delays it can provide. Hybrid schemes are attractive since they use the more appropriate technique for each of the two basic types of traffic, but they require more complex switches and policies, and do not integrate as completely these basic types as the other two schemes [Chen88].

This paper describes a method for guaranteeing delays in a packet-switching wide-area network, and presents an evaluation of some of its most important characteristics. We believe that the method can be used in all connection-oriented packet-switching environments, including ATM-style fast packet switching. However, we present it in the context of a contemporary connection-oriented packet-switching wide-area network, and we evaluate it by simulation in the same context. We state the problem precisely in Section 2, sketch our approach to it in Section 3, and illustrate the heart of our solution, i.e., the establishment algorithm, in Section 4. Section 5 describes the simulations we ran to evaluate the scheme, and discusses the results we obtained. Finally, conclusions are drawn in Section 6.

2. Problem Statement

In its full generality, the network we consider in this study has an arbitrary topology, and may consist of several networks of various types (LANs and WANs). In it, a message can go from a source host to a destination host passing through a number of intermediate nodes, some of which may be gateways (or routers) connecting one network

to another. The functions of a network node may be implemented by switches, by hosts, or by special purpose communication computers. A source-destination path goes through a number of nodes where the transmitted information can be stored and then forwarded to the next node. Some of the links between adjacent store-and-forward nodes may actually be non-store-and-forward networks, provided their total delay can be bounded. They could be, for example, very high-speed circuit-switched trunks, whose transmission speed is so high that there can be no storing or processing in the switches. The traffic on these trunks will generally include packets from different sources to different destinations, but the delay of each packet will be bounded. Thus, the only assumption we make about the links between store-and-forward nodes is that there be a known and finite bound for the link delay of each packet. Without it, we would not be able to offer hard real-time guarantees, i.e., to guarantee that no packet is delayed more than a given bound. It should be noted that this assumption is not exactly satisfied by links governed by contention-based protocols (e.g., Ethernets, see [Kuro84]), but it is by other types of LANs, such as FDDI rings (see for example [Sevc87], [Dyke88]).

In the type of network we have just described, packet delays are much harder to control, and resources much more difficult to reserve for real-time traffic, if each packet going from a given source to a given destination is in principle allowed to follow any route. For these reasons, we assume that real-time communication will be based on simplex fixed-route connections to be called *real-time channels* or simply *channels*, and that resources will be reserved for them during the connection establishment operation, as well as freed during connection disestablishment. The route of a channel will be chosen at the time of its establishment. A channel then is essentially a virtual circuit with performance guarantees[1].

The connection-oriented approach unfortunately requires explicit establishment, which may be slow. Our solution tries to reduce the establishment time by limiting the connection procedure to a single source-destination round trip and speeding up as much as possible the establishment algorithms. An additional speedup can be obtained by allowing the source to start sending packets immediately following the request message, but will not be described in this paper.

Besides offering real-time service, the network will usually provide other types of service (e.g., datagram, reliable datagram, non-real-time connection-oriented service, and so on). All of these different paradigms will have to coexist within the same network if heterogeneous traffic is to be handled effectively and efficiently.

In order to provide real-time service, we require that the clients declare their traffic

[1] Real-time channels can be regarded as "subcases" of the *parametrized message channels* on which the communication system of DASH [Ande88a] [Ande88b] is based; the only parameters characterizing a real-time channel are the performance-oriented ones. Real-time channels have similarities with the two types of *flows* recently proposed in the literature [Come89] [Zhan87], and can be easily incorporated into a recently proposed framework for a high-speed internetworking environment [Paru89].

characteristics and performance requirements at the time of channel establishment. There are various meaningful ways in which the offered load and the performance bounds of a channel can be specified. For our first study in this area, we have chosen the following parameters:

- for the offered load
  - the minimum packet interarrival time on the channel, x_{min},
  - the minimum value x_{ave} of the average packet interarrival time over an interval of duration I,
  - the maximum packet size s_{max}, and
  - the maximum service time t in the node for the channel's packets;
- for the performance bounds
  - the source-to-destination delay bound (or bounds) for the channel's packets, and
  - the maximum packet loss rate.

Note that x_{ave} is the average interarrival time during the channel's busiest interval of duration I. Note also that specifying x_{min} and x_{ave} together with s_{max} corresponds to requesting that the network provide a certain peak bandwidth and a certain long-term average bandwidth, respectively. If the channel request is accepted, i.e., if the desired bandwidths are allocated to the channel, the client is expected to satisfy the interarrival time bounds specified by the offered load parameters, whereas the delay bounds are to be guaranteed by the provider, i.e., by the network.

Various types of channels can be defined, corresponding to the different types of delay bounds. The algorithm presented in this paper considers the following two types[2]:

- *deterministic*: the bound D is an absolute one; this is necessary in hard real-time applications;
- *statistical*: the bound is expressed in statistical terms; for instance, the probability that the delay of a packet is smaller than the given bound D must be greater than a given value Z.

Channel i will be characterized by its own values of the offered load and performance bounds parameters; we will add the subscript i to the symbols of those parameters (e.g., we will write $x_{min,i}$, D_i, and so on) whenever necessary to avoid ambiguities. Delay bounds are valid only for packets that reach the destination, and are not lost in the network due to buffer overrun or network errors. However, the maximum number of packets lost must be within the loss rate bounds specified at channel establishment time. We also assume that the channels to be established are statistically independent of all the channels sharing totally or partially their route. If there are dependencies between a new

[2] We are not dealing with *best-effort* channels [Ande88b] explicitly here, since these channels do not require any reservation of resources, and since requests for their establishment are always accepted; however, our approach can very easily accommodate them.

channel and some of the existing ones, the client requesting the creation of the new channel should inform the provider. In this paper, we will not deal with the case in which there are dependencies among channels.

3. A Solution

Our approach consists of a scheme (to be described in Sections 3.3 and 4) for real-time channel establishment. The scheme works only in conjunction with specific scheduling and flow control policies. These policies are briefly described in Sections 3.1 and 3.2, respectively.

3.1. Scheduling

Scheduling in the hosts and in the nodes will be deadline-based. More precisely, we adopt a modification of the EDD (Earliest Due Date) [LiuL73] policy in which, in the case of a conflict, priority is given to deterministic over statistical channels. All of the other tasks of a host or node, including sending, forwarding, and receiving datagrams, have a lower priority; local tasks are preemptable by real-time packets.

In each node, the scheduler maintains at least three queues: one for deterministic packets, one for statistical packets, and the third for all other types of packets and all local tasks. The third queue can in turn be replaced by multiple queues, managed by a variety of policies, but we are not concerned with these aspects of node scheduling here.

A packet arriving at the deterministic queue is to be ‘‘aligned’’, that is, its deadline must be appropriately reduced if its service time would overlap with that of another deterministic packet in case they both happened to be shipped at their latest possible times. The ordering of the first two queues is by increasing deadlines.

When the node must choose the next action to be undertaken, the scheduler compares the end time (i.e., the deadline) of the packet at the head of the statistical queue with the beginning time (i.e., the deadline minus the service time) of the head packet in the deterministic queue: if the latter is lower than the former, the deterministic packet is scheduled to be shipped immediately; otherwise, the comparison is repeated between the deadline of the head task in the third queue and the beginning time of the head packet in the statistical queue. If there are more than three queues, the policy, which is a general multiclass version of EDD, can be easily extended.

3.2. Flow control

We argue that flow control in a high-speed network offering real-time service should be rate-based rather than window-based. Rate-based flow control (in which the sender controls its packet sending rate on the basis of its knowledge of the characteristics of the receiver and of those of the channel’s path) does not require flow control acknowledgements; indeed, acknowledgements would either have to use another real-time

channel, and would therefore be quite expensive, or could be sent as datagrams, in which case they might incur a rather large and perhaps highly variable delay. This delay might be too large, especially in long-distance transmissions, to be compatible with the frequencies and regularity of packet generation in most video or audio communication.

Rate-based flow control is feasible because, at channel establishment time, the receiver can check (using the method described in Section 4) whether it will be able to accept packets at the rate declared by the sender.

However, a malicious user could circumvent any rate-based flow control mechanism and from a single-user machine send packets into the network at a much higher rate than the declared maximum value, $1/x_{min}$, or maximum average value, $1/x_{ave}$. The same effect might be caused by a failure in the sending host, even when such a host is a multi-user, protected-kernel system. If we do not take appropriate countermeasures, such malicious or faulty behavior can prevent the satisfaction of the delay bounds guaranteed to other clients of the real-time service, thereby damaging the clients and destroying the credibility of the service.

Our solution to the problem consists of providing distributed rate control by increasing the deadlines of the ‘‘offending’’ packets, so that they will be delayed in heavily loaded nodes, where they would otherwise seriously interfere with the operation of other channels. When buffer space is limited, some of them might even be dropped because of buffer overflow. Of course, a sufficient amount of buffer space will have to be statically allocated to each channel to prevent the offending packets from flooding the buffer space of a heavily loaded node and causing packets from other channels to be dropped. A very simple algorithm for distributed rate control is presented in detail in [Ferr89].

One important consequence of this scheme for the algorithms in Section 4 is that we can assume that each channel satisfies its x_{min} and x_{ave} constraints at every node it traverses.

3.3. The channel establishment procedure

The channel establishment mechanism we propose for the type of network described in Section 2 may be built on top of any procedure that can be used to set up connections. Besides trying to establish a connection, the mechanism will perform several tests and tentatively reserve resources in each node visited by the establishment request message. In order to make channel establishment fast, we impose on our procedure the restriction that it require only one round trip. Thus, the destination host is the last point along the path where the acceptance/rejection decision for a channel request can be made. When a node is revisited by an establishment message during this message’s return trip, the resources previously reserved there must be committed or released; hence a final, irreversible decision must have already been made.

The tests to be done in each node are concerned with the availability of sufficient bandwidth in the links, as well as processing power and buffer space in the node; that is, with determining whether the new channel can go through the node without jeopardizing

the performance guarantees given to the already established channels passing through the same node.

If any test fails at a node, the channel cannot be established along that route; the message will be sent back, either to the sender (which may then decide to wait or try another output link) or to an intermediate node that can try sending the message towards the destination along another path. When an unsuccessful establishment request message revisits a node on its way back to the source, it frees all the resources that were tentatively reserved there during its forward trip.

If all the tests for channel i succeed at all nodes and at the destination host, this host subdivides the delay bound D_i (and the probability Z_i of not exceeding the bound if the channel being established is statistical) among the nodes traversed by the channel, after subtracting the link delays along the route. Let $d_{i,n}$ be the delay bound assigned to the channel in a node n, and $z_{i,n}$ the probability that $d_{i,n}$ will not be exceeded in that node. A packet traveling on that channel and arriving at that node at time t will usually [3] be assigned a node deadline equal to $t+d_{i,n}$. To satisfy the overall delay bound D_i, it is sufficient (though not necessary) to satisfy the bound $d_{i,n}$ (or $d_{i,n}, z_{i,n}$) in each node along the route. For simplicity, we make the very conservative assumption that satisfying this sufficient condition is the goal to be met by the establishment scheme in each node.

If all tests succeed, a reply message is sent back to the source host along the channel's route; this message notifies each node about the delay bound that has been assigned to it for the new channel, and commits the necessary portions of reserved resources. When the reply message reaches the source host, this host learns that the requested channel has been set up, and can start using it.

The next section describes in some detail the algorithms for channel establishment to be executed in each node along the new channel's route and in the destination host.

4. The Algorithms

To simplify our discussion, we shall assume here that the buffer space needed for real-time connections in the network's hosts and nodes is always available, and that the network's error rate is always lower than the acceptable loss rates. These assumptions allow us to ignore the loss rate parameter, and to postpone the treatment of buffer space allocation and management, which would make this paper appreciably longer, to a subsequent paper.

When a node receives an establishment request message, it performs two or all three of the following tests:

---

[3] There may be exceptions due to alignment (see Section 3.1) or distributed rate control (see Section 3.2).

(a) the *deterministic test*, primarily required when the channel to be established is deterministic, and involving the deterministic channels already passing through the node;

(b) the *statistical test*, to be performed for the establishment of both statistical and deterministic channels if at least one statistical channel is already passing through the node or is to be set up; this test involves all deterministic and statistical channels passing through the node;

(c) the *delay bound test*, which is needed in all cases; if successful, it is to be followed by the computation of the minimum feasible delay bound for the new channel in the node.

If the request passes the tests, then the node sends the establishment message on to the next node, after having added to it some of the parameter values resulting from the tests listed above. The destination host performs the final tests, and then, if the request passes these too, subdivides D (and Z) among the nodes on the channel's path. We now briefly describe each one of the tests as well as two simple destination host algorithms.

4.1. The deterministic test

The deterministic test consists of verifying that enough processing power is available in the node to accommodate the additional deterministic channel without impairing the guarantees given to the others. Since we are dealing with deterministic bounds, this condition must be satisfied even in the worst possible case, i.e., even when all deterministic channels are sending packets into the node at their maximum rates. The maximum utilization of a node n by channel j, whose packets have a maximum service time in the node equal to $t_{j,n}$, is $t_{j,n}/x_{min,j}$. The condition to be tested is:

$$\sum_j (t_{j,n}/x_{min,j}) < 1, \qquad (1)$$

where the sum extends to all deterministic channels passing through node n, including the one to be established (channel i).

4.2. The statistical test

The statistical test has two goals:

(i) to determine whether for each statistical channel j passing through node n the probability of a delay higher than the bound $d_{j,n}$ is below its maximum tolerable value, $1 - z_{j,n}$;

(ii) to provide the destination host with the information necessary to compute $z_{i,n}$ for the new channel if this is statistical.

Both these goals can be attained by computing the *probability of deadline overflow* $P_{do,n}$. A packet may be delayed beyond its delay bound in a node because of two reasons: a temporary saturation of the node's processing or transmission capacity (*node saturation*), and impossible scheduling constraints (*scheduler saturation*)[4]. While the latter is avoided, as discussed in Section 4.3, the statistical test deals with the former.

[4] A simple example of scheduler saturation is that of two channels with respective node service times 3

The probability that channel j is *active* (i.e., is carrying packets) at some instant of time during an interval I is

$$p_j = x_{min,j} / x_{ave,j}. \quad (2)$$

Given K independent channels passing through a node, the probability that the members of a subset C of them are simultaneously active is given by

$$Prob(C) = \prod_{i \in C} p_i \prod_{j \in \bar{C}} (1 - p_j). \quad (3)$$

Let at least one of the channels be statistical. To compute $P_{do,n}$, we start by listing all the *overflow combinations*. An overflow combination is a set of channels that, when simultaneously active for a sufficiently long time, may cause packets to miss their deadlines. In other words, an overflow combination is one for which inequality (1) above, with the sum extended to all active channels, is not satisfied. In making this statement, we assume a worst-case situation, i.e., one in which each active channel carries packets at its maximum rate. If the input load were to persist indefinitely, the length of the queue of packets in the node would grow without bounds.

Let H_n be the set of overflow combinations in node n, h be a member of H_n, and $P(h)$ be the probability of occurrence of combination h computed as in (3) above. Then,

$$P_{do,n} = \sum_{h} P(h). \quad (4)$$

We can then check whether the following inequalities are satisfied:

$$P_{do,n} \leq 1 - z_{j,n}, \quad (5)$$

for all statistical channels j existing in the node. To see whether (5) is satisfied, it is sufficient to verify that

$$P_{do,n} \leq \min(1 - z_{j,n}), \text{ or } 1 - P_{do,n} \geq \max z_{j,n}. \quad (6)$$

If test (6) is successful for all statistical channels j in the node, the value of $P_{do,n}$ is sent to the destination host, which will try to assign to channel i a value of $z_{i,n}$ that satisfies inequality (6).

4.3. The delay bound test

The delay bound test determines the minimum delay bound to be assigned to the channel being established so that scheduler saturation can be avoided in the node. We try to eliminate this type of saturation primarily because of the complexity of dealing with it in the general case.

To determine whether scheduler saturation is possible in a node, we divide the K (deterministic or statistical) channels passing through the node into two sets: we call U

---

and 4 units, and respective node delay bounds 5 and 6 units; if two packets from the two channels arrive simultaneously at the node, there is no way to schedule them so that both meet their deadlines.

the set of those channels whose delay bound in the node is lower than the sum of the service times of the K channels, and V the set of those channels whose delay bound in the node is greater than or equal to that sum. With no loss of generality, we assume that the u channels in U are those numbered 1 through u, and that they are numbered according to the order in which their packets would be scheduled by the algorithm described in Section 3.1 if they arrived all at the same instant: channel 1 will be the one whose packet would be shipped first, channel u the one whose packet would be shipped last.

$$\mathrm{U}=\{i \mid i=1,...u;\ d_{i,n} < \sum_{j=1}^{K} t_{j,n}\}, \tag{7}$$

$$\mathrm{V}=\{k \mid k=u+1,...K;\ d_{k,n} \geq \sum_{j=1}^{K} t_{j,n}\}, \tag{8}$$

In the statistical test (see Section 4.2), we took into account node saturation, which delays packets beyond their deadlines whether or not scheduler saturation is present. Thus, in those situations in which the node is saturated, whether we have scheduler saturation or not is immaterial. However, scheduler saturation could occur also in cases in which the node is not saturated. To determine whether this will indeed happen, and to prevent it from happening, we must consider the worst possible combination, i.e., one that maximizes the likelihood of scheduler saturation without saturating the node. If the combination corresponding to all channels through the node being active saturates the node, we will have to consider for the delay bound test a combination in which some of the channels are not active. Details about how such a combination can be arrived at are given in [Ferr89].

For the combination we have selected, let us denote by T the largest of the service times of packets that may traverse the node but do not travel on channels in set U. These include not only those on channels in set V, but also those on other real-time channels (e.g., established channels not included in the combination to avoid node saturation) as well as the non-real-time packets that may be passing through the node. Let us also assume for simplicity of proof that

$$x_{min,i} \geq \sum_{j=1}^{K} t_{j,n}, \quad (i=1,...K). \tag{9}$$

The following theorem is the basis for the delay bound test.

Theorem: Scheduler saturation is impossible if and only if

$$d_{i,n} \geq \sum_{j=1}^{i} t_{j,n} + T, \quad (i=1,...u). \tag{10}$$

A proof of this theorem can be found in the Appendix, where we briefly discuss also the cases in which inequalities (9) do not hold. [Ferr89] contains a few examples of delay bound computations.

The delay bound test consists of verifying that inequality (10) is satisfied for all values of i (i.e., for all channels in U), after the channel to be created has been added to the K already established ones. The arrival of a request for the establishment of a new channel changes $\sum t_j$, hence possibly the memberships of U and V, and may change even

the value of T. While testing the delay bounds of the established channels, if the delay bounds pass the test it is easy to determine the lower bound $d_{i,n}^{l}$ of $d_{i,n}$ such that scheduler saturation will be impossible in the node if we choose $d_{i,n} \geq d_{i,n}^{l}$. The value of $d_{i,n}^{l}$ will then be sent to the destination host in the establishment message.

4.4. The destination host tests and algorithms

The destination host, if and when it receives the establishment request message, has to determine whether the total value of Z_i for the new channel i can be factored into node contributions $z_{i,n}$ (n = 1, 2,...N):

$$Z_i = \prod_{n=1}^{N} z_{i,n}. \tag{11}$$

This question can be answered by checking whether

$$Z_i \leq \prod_{n=1}^{N} (1 - P_{do,n}). \tag{12}$$

This may be called the *Z test*. If (12) does not hold, then the request is rejected. Note that (11) implies the assumption that, once a packet is delayed beyond its deadline in a node, it will not be able to satisfy the channel's overall delay bound. This assumption is justified by the behavior of a saturated node: even if the delay bounds are long, but finite, packets will eventually be delayed beyond them. If, however, the saturation condition is short, some or all of the deadlines might still be met. This, and the assumption, also implied by (11), that each node will delay, if any, only previously undelayed packets, are certainly pessimistic assumptions. Rejecting a request if (12) does not hold is a policy based on worst-case considerations, which guarantee that the desired results will be achieved in all possible circumstances. Of course, there is a cost associated with worst-case design: these and other conservative decisions reduce the maximum number of channels that can be established in a given network with respect to that which the same network could support if the design of the real-time service were less conservative.

The test concerned with the delay bound (the *D test*) consists of verifying that, for channel i with total delay bound D_i to be established, we must have

$$D_i \geq \sum_{n=1}^{N} d_{i,n}^{l}. \tag{13}$$

Two simple expressions for assigning $d_{i,n}$'s and $z_{i,n}$'s to the nodes along the channel's path are

$$d_{i,n} = \frac{1}{N}\left[D_i - \sum_{m=1}^{N} d_{i,m}^{l} \right] + d_{i,n}^{l}, \tag{14}$$

$$z_{i,n} = \left[Z_i / \prod_{m=1}^{N} (1 - P_{do,m}) \right]^{1/N} (1 - P_{do,n}). \tag{15}$$

5. The Simulations

The reader certainly has many questions about the performance of the scheme described in the previous sections. In this section, we provide simulation-based answers to the following ones:

- Is the scheme correct? Can we discover cases when the guarantees of the channels are not met using this scheme?
- How much time is required to establish a connection?
- How much time is required by the scheduler to select the packet to be transmitted next?
- What is the overhead of distributed rate control?
- Providing performance guarantees involves a number of pessimistic assumptions. Is our scheme too conservative? For example, how close are our guaranteed bounds on delays to the actual delays?
- What is the ‘‘cost’’ of a statistical channel as compared to that of a deterministic channel?

We provide answers to these questions by means of a network simulator written in CSIM [Schw87]. We have simulated several different networks but will only present and discuss the results obtained for the 14-node network shown in Figure 1(a) and for a simple network consisting of only two nodes.

In Figure 1(a), nodes 0-8 serve as sources and destinations for the real-time channels. These nine nodes were divided into three groups of three nodes each. Each source generated channel requests at random intervals. The destination group of such a channel request was chosen by the probability matrix in Table 1. The actual destination was picked randomly from the group. One of the four shortest paths from a channel's source to its destination was picked at random as the route of its establishment request.

5.1. Workload

Future multimedia networks will have a wide variety of traffic types. To capture some of this diversity, we considered four types of channels. Table 2 shows the types of channels used in the simulation experiments. Type I channels tax the network's resources heavily, while type IV are very light.

The simulations were run in two phases. In Phase 1, we tried to establish as many channels as we could. The next channel to be established was of any type and either deterministic or statistical with equal probabilities. The exact number and types of channels set up in such a context depended on the order of arrival of channel requests. In some cases, a request for a type I channel arrived first, and many subsequent requests for type IV channels were rejected. On the other hand, if the first few channels were of type IV, a subsequent request for a type I channel might be refused, but many more channels of type III could be established.

Let the *load* of a channel i at node n be defined as $t_{i,n}/x_{min,i}$ (see Section 4.1 for a motivation). The total load of all the channels accepted at a node reflects the maximum instantaneous utilization of that node allowed by our scheme. In that sense, the load serves as a better measure of saturation at a node than the number of accepted channels.

The results of one Phase 1 simulation of the network in Figure 1(a) are shown in Figure 1(b). The two numbers associated with each node are the load of all the channels and of the deterministic channels only, respectively, at that node. The number [1.55,0.5] with node 1 implies that when all the statistical channels are sending packets at the maximum possible rate, the arrival rate of packets exceeds the service rate of the node by a factor of 1.55, while if only all the deterministic channels were sending at peak rate, the arrival rate would be 0.4 times the service rate. The maximum possible utilization of node 1 by the deterministic channels alone is 40%.

We selected this configuration containing about 38 successfully established channels for Phase 2 simulations. Note that routing constraints prevent nodes 10, 11 and 12 from being saturated in this network. In Figure 1(b), we have also highlighted one deterministic and one statistical channel. We shall examine the delays experienced by the packets on