DESIGN AND APPLICATIONS OF A DELAY JITTER CONTROL SCHEME FOR PACKET-SWITCHING INTERNETWORKS

Domenico Ferrari
The Tenet Group
University of California
and International Computer Science Institute
Berkeley, California, U.S.A.

Abstract

Delay jitter is the variation of the delays with which packets traveling on a network connection reach their destination. For good quality of reception, continuous-media (video, audio, image) streams require that jitter be kept below a sufficiently small upper bound. This paper proposes a distributed mechanism for controlling delay jitter in a packet-switching network. The mechanism can be applied to an internetwork that satisfies the conditions detailed in the paper, and can coexist with other schemes (including the absence of any scheme) for jitter control within the same network, the same node, and even the same real-time channel. The mechanism makes the distribution of buffer space requirements more uniform over a channel's route, and reduces by a non-negligible amount the total buffer space needed by a channel. The paper argues that, if these advantages are sufficient to justify the higher costs of the distributed jitter control mechanism with respect to a non-distributed one, it would be useful to offer to the network's users a jitter control service based on the mechanism proposed here.

1. Introduction

Digital audio, video, and image-sequence transmission, both of the interactive and of the non-interactive kind, need a high degree of regularity in the delivery of successive stream-frames[1] to the destination. Ideally,[2] the end-to-end connections that carry those continuous-media streams should delay each element (e.g., a frame) of a stream by the same amount of time; in other terms, each connection should behave as a perfect delay line. This should normally be the case even when the elements of a stream are not transmitted at a constant rate (for example, in image browsing applications, where the user needs to be able to do pausing, fast and slow forwarding, fast and slow rewinding, and so on).

This research was supported by the National Science Foundation and the Defense Advanced Research Projects Agency (DARPA) under Cooperative Agreement NCR-8919038 with the Corporation for National Research Initiatives, by AT&T Bell Laboratories, Hitachi, Ltd., Hitachi America, Ltd., the University of California under a MICRO grant, and the International Computer Science Institute. The views and conclusions contained in this document are those of the author, and should not be interpreted as representing official policies, either expressed or implied, of the U.S. Government or any of the sponsoring organizations.

[1] We call *stream-frames* the data elements (video frames, audio samples, image frames) that constitute a continuous-media stream.

[2] We ignore here the case of stored streams that are shipped in batches and are to be played at the destination upon reception, as well as all the other cases in which input patterns are not to be faithfully reproduced at the output of the network. These either must be regarded as file transfers (i.e., not as continuous-

An ideal constant-delay network could be characterized also as a *zero-jitter* one, where jitter (more precisely, *delay jitter*) is any variation of stream-frame delays. A more realistic requirement for continuous-media transmission is that jitter have a small upper bound. Given an upper bound for delays, which is necessary for interactive continuous media and often desirable even in non-interactive cases, assigning a jitter bound is equivalent to specifying also a lower delay bound fairly close to the upper bound. A jitter bound alone will not guarantee the faithful reproduction of the input pattern at the output. To achieve this result, such a bound will have to be accompanied by a delay bound guarantee.

The computers and the (packet-switching) networks that connect users involved in continuous-media communications may introduce large amounts of jitter, due to the fluctuations of their loads. End-to-end delay bounds can be guaranteed by suitable design of the various components of a connection, including the host computers at the two ends; for instance, a scheme for establishing real-time connections with deterministic or statistical delay bounds, as well as throughput and loss probability bounds, in a packet-switching network is described in [FeVe90]. Is it possible to reduce (below a small upper bound) end-to-end delay jitter as well?

The answer is affirmative, as long as (*i*) a bound on delay can be guaranteed, (*ii*) a sufficiently large amount of buffer space is available at the destination or close to it, and (*iii*) the clocks of the source and the destination are kept in synchrony with each other [Part91]. This jitter control scheme is very simple: the source timestamps each stream-frame, and the destination, knowing the timestamp and the ideal constant delay (which must be greater than or equal to the delay bound guaranteed by the network), buffers each stream-frame until the time it has to be delivered to the receiving user.

The scheme can be implemented at the application level or at a lower level (e.g., at the transport or even the network layer in the communication protocol hierarchy); it is easy to see that the existence of a jitter control mechanism at a lower level, while not automatically providing low jitter to the application, greatly facilitates the achievement of this objective. Essentially, a bounded-jitter packet stream is not a necessary, but may be regarded as a sufficient, condition for obtaining a bounded-jitter sequence of stream-frames, as long as stream-frame reassembly, higher-level protocol processing, and stream-frame delivery to the application do not introduce appreciable amounts of jitter.

An alternative approach to jitter control has been presented in [VeZF91]. While the approach described above (see for example [Part91]) operates primarily at the destination, the one proposed in [VeZF91] is distributed; thus, it can only be implemented by the network layer, not by the upper protocol layers or by the application. In the kinds of networks in which it can be used, the distributed approach has been found to require a substantially smaller total amount of buffer space, and to yield a more uniform buffer allocation along the path of a bounded-jitter connection. In this paper, we introduce a distributed scheme that can be applied to much more complicated networks (in fact, to internetworks), and argue that, if a network already implements distributed jitter control for

---

media transmissions) from a storage device to a remote one, or require a receiver that will time their presentation following a given (e.g., periodic) time law. Some form of jitter control may be needed even in these cases.

its internal purposes, it could, at almost no extra cost, offer a jitter control service to its users; indeed, if the receiving systems can be designed to inject only negligible amounts of jitter into the incoming streams, this service would provide the end-to-end jitter bounds needed by continuous-media streams.

A more detailed discussion of the two approaches (distributed and non-distributed) to jitter control can be found in Section 4. That discussion is preceded by a description of the network environment in Section 2, and by a presentation of our scheme in Section 3.

The rest of this section summarizes the approach to real-time communication presented in [FeVe90], which, like the jitter control scheme in [VeZF91], applies to what we will in the sequel call a *simple network*. Such a (packet-switching) network is characterized by hosts and switches interconnected by physical links in an arbitrary topology. The switches operate according to the store-and-forward principle, the links introduce only propagation delays, and the output queues of a switch (one per output link) are managed by a multiple-class version of the Earliest Due Date (EDD) [LiLa73] discipline.

A path from source to destination in such a network may be depicted as in Figure 1. Nodes 2, 3,... N-1 represent the switches traversed by the path. Node 1 and node N implement the protocol layers below the transport layer on the host systems. For these systems to offer end-to-end delay guarantees, they must be scheduled by a policy (such as Multi-Class EDD) that gives higher (and preemptive, when necessary) priority to the real-time activities of the upper protocol layers, as well as to the output queues of the "resident network node" reserved for real-time packets.

The basic abstraction the scheme in [FeVe90] is capable of implementing is the *real-time channel*. A real-time channel is a simplex connection between a sender and a receiver[3] characterized by performance guarantees. The performance guarantees a client may request include a delay bound *D* for each packet; on a *deterministic channel*, packets are guaranteed to arrive at the destination with a delay not larger than *D*.

Besides specifying performance requirements such as *D*, a client will have to provide a simple description of the maximum traffic on the channel to be established. One of the parameters this description consists of in the approach we are summarizing here is $x_{\min}$, the minimum time between successive packets at the entrance to the channel.[4] This will also automatically specify the client's maximum bandwidth requirements.

The creation of a real-time channel may be requested by a client running on the host system that should become the channel's source. An establishment message carrying the performance requirements and the traffic specifications given by the client leaves the source and, while building the connection, tests the ability of each switch to handle the new channel without violating the bounds promised to the other clients; it also collects information, to be used by the destination, on the loading of each switch. If the message reaches the destination without having been rejected, the destination runs the final tests; if also these tests are successful, the destination subdivides the client-given bounds into per-switch bounds, and sends them to the switches with the channel acceptance message.

---

[3] The notion of real-time channel can be extended to the case of multiple receivers. For the sake of simplicity, we will not consider such an extension in this paper.

[4] The client will actually provide the value of the minimum time between successive stream-frames; this will be translated into the minimum interpacket time by the channel establishment protocol.

Thus, switch *n* receives, for example, the value of its local delay bound $d_{i,n}$ for the new channel *i*, and will use it to compute the local deadline of each packet transmitted on that channel.

2. The Network

In this section, we describe the environment for which our jitter control scheme has been designed; in particular, the class of networks (to be more precise, internetworks) that can guarantee delay jitter bounds, as well as delay bounds, using the approach to be presented in Section 3. The scheme described in [FeVe90] for simple networks can be extended to a class of internetworks [Ferr91b], which is the one we consider here. The rest of this section gives a brief description of the characteristics of that class.

We consider an internetwork in which it is possible to create real-time channels. We assume that the internetwork offers channels with deterministic (i.e., absolute) delay bounds, and we wish to devise a way to obtain guaranteed jitter bounds for the same channels when the clients need such guarantees.

The path of a channel from the sending host to the receiving host will generally traverse a number of different networks, interconnected by gateways.[5] We consider a hierarchical model of the internetwork in which the level 1 nodes are the gateways. The level 1 path of a channel is a path consisting of gateways interconnected by *logical* links; each one of these links may represent a path through the network of which the two nodes at its extremes are gateways. This path may be very complicated; for example, the network in question might in turn be an internetwork, whose gateways constitute a level 2 network, and so on recursively.

In this hierarchical model, a path through a network may be abstracted by transforming it into a simple logical link at the immediately higher level (note that in this paper higher levels are labeled with lower numbers). This process of abstraction ends with the level 0 path, which is a single logical link connecting the sender and the receiver. This is the logical link whose delay and delay jitter are to be bounded. What are the minimum necessary properties of the underlying internetwork that allow the bounds selected by the client to be guaranteed, if possible?[6] What must we assume about the various networks constituting the hierarchy?

If all the constituent networks were capable of guaranteeing delay and jitter bounds, the internetwork could provide similar guarantees at the end-to-end level. However, such a restrictive condition is by no means necessary. All we need (see [Ferr91b]) is the assumption that the delay on each level 1 logical link be bounded or at least boundable (with a sufficiently small bound), and that those bounds be either known *a priori* or obtainable by a recursive application (at the various levels in the hierarchy) of an extension of the scheme described in [FeVe90].

We also assume that level 1 nodes are store-and-forward, and schedule packets for transmission on each of their output links in such a way that, if the client's request for the channel is accepted, it is possible to guarantee the end-to-end delay bound specified by

---

[5] We use this term in the generic sense of "a computer system interconnecting two distinct networks at the network or higher layer;" thus our "gateways" may be routers as well as gateways proper.

[6] If it is not possible to provide the requested guarantees, the client's request will be rejected.

the client.[7] Many scheduling algorithms satisfy the latter assumption (see [Ferr91b]); for example, round-robin and its extensions and modifications (Fair Queueing [DeKS89], Hierarchical Round-Robin [KaKK90]), deadline-based disciplines (Earliest Due Date or EDD, Multi-Class EDD [FeVe90]), FCFS, and others.

If these assumptions, which are summarized in Figure 2, are satisfied, the internetwork can provide guaranteed delay bounds [Ferr91b], and we assume that the internetworks in the class we are considering indeed do provide them.

For convenience, we break the delay bound $dl_{i,l}$ of a level 1 logical link l for channel i into the sum of two terms, a fixed delay $df_{i,l}$ and a maximum variable delay $dv_{i,l}$, where $df_{i,l}$ is the delay encountered by the smallest possible packet travelling on the logical link under zero-load conditions, and $dv_{i,l}$ is an upper bound to the variable components of packet delays along the same link. Thus, $df_{i,l}$ includes propagation, switching, and transmission delays for the minimum-size packet, while $dv_{i,l}$ primarily consists of queueing delays and the additional transmission delays that a maximum-size packet would incur. The actual delay of a packet on that channel along that level 1 logical link will be between $df_{i,l}$ and $dl_{i,l}$.

As in a simple network, each level 1 node traversed by channel i has a delay bound $d_{i,n}$ assigned to it by the destination at channel establishment time. This bound is the maximum residence time of a packet in node n. The minimum residence time is the packet's processing time in the node (which includes also its transmission time on the output link); we call $dp_{i,n}$ the smallest of these times, corresponding to the idle-node residence time of the smallest packet the channel will carry; by $dq_{i,n}$ we denote the maximum queueing time in the node, which can only be reached by the smallest packet, since for all packets the sum of processing time and queueing time will be at most $d_{i,n}$. Thus, we have $d_{i,n} = dp_{i,n} + dq_{i,n}$.

3. Jitter Control with Perfectly Synchronized Clocks

In an internetwork of the class defined in Section 2, let us assume that the clocks of all level 1 nodes are in perfect synchrony with each other. That is, simultaneous readings of these clocks are always identical. A reasonably good approximation to this ideal situation can be achieved by having the gateways run a clock synchronization protocol (e.g., NTP [Mill89]), or listen to the times broadcast by a satellite and adjust their clocks accordingly. In the level 1 path of a real-time channel, all components of the end-to-end delay bound are known after the establishment of the channel. Each logical link l of channel i has a total delay bound equal to $df_{i,l} + dv_{i,l}$, and each node n has a local bound $d_{i,n}$ assigned to it by the destination (see Section 1 and [FeVe90]). Thus, it is easy to calculate the maximum delay between a node and the next, and to compensate for the jitter due to the variable parts of the delay, which may make the actual delay shorter (sometimes much shorter) than the maximum.

The mechanism is very simple: it requires timestamping each packet's header by the source host, reading the clock at a packet's arrival at each level 1 node, and keeping the packet until the *canonical* (i.e., zero jitter) arrival time at that node before giving it to the

---

[7] Note that no assumptions are made about gateways or routers at lower levels in the hierarchy. The only constraint on them is that they allow level 1 delays to be boundable.

node's scheduler. Each node is informed, when a real-time channel is created, about the maximum delay that separates it from the source for the packets on that channel; this is the *canonical delay* $D^*_{i,n}$ for channel i at node n. Assuming that the clock is read in each node immediately at packet arrival time, the maximum delay with respect to the canonical arrival time at the previous node (node $n-1$) equals $df_{i,l} + dv_{i,l} + d_{i,n-1}$.

Given the time of departure from the source, which is carried by the header, the node calculates the canonical arrival time and compares it with the actual one (which is guaranteed not to exceed the canonical time). The packet will be kept in storage in the node for a time interval equal to the difference between the canonical and the actual arrival times. When that interval expires, the packet becomes *schedulable*; the jitter introduced by the scheduler in node n will be corrected exactly by the same mechanism in node $n+1$. The only jitter that cannot be corrected in this way is therefore the one introduced by the last node on the path. This will have to be corrected by the receiving host, which will need to know the maximum delay assigned to the last node and keep the packet until that bound is reached before giving it to the application. If this is done, a zero-jitter (perfect delay line) transmission is obtained; if it is not done, the only end-to-end jitter is that generated by the last node, since those produced by all previous nodes and (logical) links are perfectly compensated for by the simple mechanism we have just illustrated.

What is the maximum value of jitter when no jitter control is used, and how much is that value reduced by the above scheme? To answer these questions, we refer to the channel whose path is shown in Figure 1.[8] We therefore omit the channel index i in the symbols we will use, and assign to each link the same number as the node immediately downstream.

Without jitter control, the link-node pair numbered n contributes to the total maximum jitter the quantity

$$j_n = dl_n + d_n - dp_n - df_n = dv_n + dq_n, \tag{1}$$

Equation (1) can be explained by observing that the maximum delay through the link-node pair is $dl_n + d_n$, whereas the minimum delay is $dp_n + df_n$.

Thus, the maximum jitter for the packets on the channel is

$$J_{\max} = \sum_{n=1}^{N} j_n = \sum_{n=1}^{N} (dv_n + dq_n). \tag{2}$$

From the above discussion, we have that the maximum jitter permitted by the jitter control mechanism described in this section is[9]

$$J'_{\max} = d_N - dp_N = dq_N. \tag{3}$$

Note that the maximum jitter at the input to node n is

---

[8] The nodes in the figure are level 1 nodes, which run an internetwork protocol. Node 1 represents the network layer and lower layers on the source computer, and shares that computer with source S, which represents the application and the protocols at and above the transport layer.

[9] We use primed symbols to refer to cases in which our distributed jitter control mechanism is present.

$$j'_{in,n} = dv_n + dq_{n-1}, \tag{4}$$

while the maximum jitter at the same point when there is no distributed jitter control is

$$j_{in,n} = dv_n + \sum_{m=1}^{n-1} j_m. \tag{5}$$

The value of $j'_{in,n}$ equals the maximum amount of time a packet may have to wait in the node before becoming schedulable. Thus, the maximum residence time of a packet in node n with distributed jitter control is

$$r'_n = dv_n + dq_{n-1} + d_n, \tag{6}$$

and only

$$r_n = d_n \tag{7}$$

without distributed jitter control; in the last node, however, if jitter is to be absorbed by buffering (a network-layer implementation of the non-distributed approach to jitter control described in Section 1), the maximum residence time must equal the total maximum jitter instead of d_N as dictated by (7):

$$r_N = \sum_{n=1}^{N} j_n \tag{8}$$

To calculate the number of packet-sized buffers needed by each channel in node n to avoid packet losses due to buffer overflows, we observe that the worst case corresponds to a stream of packets entering the channel at node 1 with interpacket distances equal to $x_{\min}$. The latest time a packet from the stream can leave node n equals the sum of its departure time from the source, the canonical delay at node n, and the local delay bound in the same node. Similarly, the earliest time a packet can arrive at node n is given by its departure time from the source plus the canonical delay at node n minus the maximum jitter at the input to the same node. If we number the packets in the stream 1, 2, ..., and we assume that node n has allocated B_n buffers to the channel, the packet numbered 1 will have to have left the node before the arrival of the packet numbered $B_n + 1$ even in the worst case. This will be the case in which packet 1 departs at the latest possible time and packet $B_n + 1$ arrives at the earliest possible time. Thus,

$$d_n \le B_n x_{\min} - j_{in,n}, \tag{9}$$

and

$$d_n > (B_n - 1)\, x_{\min} - j_{in,n}. \tag{10}$$

Combining (9) and (10), we obtain[10]

$$B_n = \left\lceil \frac{d_n + j_{in,n}}{x_{\min}} \right\rceil. \tag{11}$$

[10] To be more precise, we should add the maximum transmission time on the output link to the numerator of the fraction in equations (11) and (12); however, this time is often negligible with respect to the terms that appear in that numerator, and we will therefore omit it.

If jitter is controlled with the distributed method described above, our treatment is still valid, as long as we replace $j_{in,n}$ with $j'_{in,n}$. Thus,

$$B'_n = \left\lceil \frac{d_n + j'_{in,n}}{x_{\min}} \right\rceil = \left\lceil \frac{r'_n}{x_{\min}} \right\rceil . \qquad (12)$$

Since $j'_{in,n} \le j_{in,n}$, we usually have $B'_n \le B_n$. In fact, while $j_{in,n}$ (hence also B_n) grows monotonically with the number of nodes traversed along a channel's route (see equation (5)), $j'_{in,n}$ does not (see equation (4)). These simple worst-case calculations show what we believe to be the main advantages of distributed jitter control when no packet loss can be tolerated, i.e., the smaller amount of buffer space to be allocated, and the much more uniform distribution of that space along channel routes. Because of the better balance among per-node space requirements provided by our distributed jitter control scheme, this greater uniformity may be expected to increase the maximum number of real-time channels a given network can support.

It is interesting to observe that the value of B_N is given by equation (11) even when jitter is to be eliminated by buffering in the last node (in this case, however, buffer space utilization in that node will be much higher). This rather counterintuitive result can be explained by noting that, even when the last node is expected to absorb jitter, a maximally delayed packet will not be made to wait in the last node's buffer, and that a minimally delayed packet will be delayed in the last node only until it reaches the maximum delay. The size of the last node's buffer will have to be B_N (if no packet losses can be tolerated) whether or not that buffer is used to absorb jitter. If jitter absorption is done by a separate buffer at the transport or higher layer or in the user's application, that buffer must have the same size as the one in the last node. Thus, unless the two buffers can be implemented in the same physical memory space (which is at least in principle possible, since the last node and the "host" coexist within the same computer system), jitter absorption at the network layer requires half the space needed by any solution that does it at a higher layer.

4. Discussion

Good-quality reception of digital video streams, audio streams, and image sequences over computer networks normally requires bounded delay jitter as well as bounded delay connections. Acceptable jitter bounds can be obtained, in a system with delay guarantees, sufficient buffer space at the destination, and sufficiently accurate synchronization between the source and the destination clocks, by a simple scheme that absorbs the jitter by buffering data at the destination. The scheme can be implemented at one of several different levels; a network-layer implementation may even buffer data in the last switch or the last gateway of a connection rather than in the destination host.[11]

In [VeZF91] and in this paper, two distributed jitter control schemes (which, being distributed, can only be implemented at the network layer) have been introduced. The one described in [VeZF91] only applies to simple networks, but, unlike the one based on buffering at the destination, does not need to make any assumptions about clock

[11] This will provide end-to-end controlled jitter only if jitter is kept under control also between the place where data are buffered and the destination.

synchronization. The scheme presented in Section 3 requires clocks to be synchronized (at least loosely [Ferr91a]), but can be used in hierarchical internetworks.

The main advantage of the distributed approach is that it does not require increasing amounts of buffer space as the distance from the source grows; it spreads buffer space requirements much more evenly over the nodes on a channel's path. This will in general concern network designers and managers more than clients, though some users may actually find it inconvenient to have to allocate large buffers in their hosts for each of the several continuous-media streams those hosts may have to receive simultaneously.

The main disadvantage of a distributed scheme is that it complicates the network switches and switching operations (which in high-speed networks should be as simple as possible); however, this problem is not so crucial if the scheme is only implemented in an internetwork's gateways.

If the advantage mentioned above is deemed useful enough, and if the disadvantage just discussed seems tolerable enough, that designers decide to implement jitter control inside the internetwork anyway, the maximum jitter perceived by the clients will decrease very substantially, and it should be much easier and less expensive to reduce the end-to-end jitter below the bounds desired by the clients by buffering data at the destination; furthermore, a jitter control service based on the distributed scheme could be offered to the internetwork's clients, and the clients might be able to rely on it for most of their jitter bound requirements, without any need for further end-to-end jitter reductions. This latter proposal might be regarded as being in conflict with the end-to-end principle, but we are actually assuming that controlling jitter in the network would come so close to achieving end-to-end jitter control that, at least for most applications, no additional jitter control mechanism would be necessary.

If we now compare an application-level jitter control scheme (which can only be a non-distributed one) with a system-level one (which may be distributed or non-distributed), we see that with the former, as already mentioned above, the bound is really an end-to-end one; on the other hand, the latter allows the application programmer to avoid worrying about timestamping, reading the clocks, allocating buffers for jitter control, recovering from synchronization losses, and so on, simply by using the service offered by the system. Of course, this service must be able to provide clients with what they need, and whether or not a generally useful jitter control service can be implemented can only be determined by experimentation.

This advantage of a system-level solution is actually even more important than it looks, since a jitter control service can be used to solve many stream synchronization problems, which would add substantial complexity to the programming of multiple-stream applications (see for example [Ravi91], [ShSa90]). Such applications can be categorized according to their synchronization requirements. Brief descriptions of three of these classes follow.

(a) *Synchronization of multiple periodic streams.*

The client requests the creation of a real-time bounded-jitter channel for each stream, and specifies for each channel the same delay bound and the same jitter bound. In the simplest case (e.g., that of a video stream and of its separate but co-stored sound track), all streams have the same source and the same destination. For (approximately) simultaneous arrivals of stream-frames from different streams at the destination, it is then

sufficient that the stream-frames be input into the source node at the same time. The maximum difference between delivery times will be equal to the common jitter bound. Any phase difference between stream-frames at the input will be reproduced (with a maximum error equal to the jitter bound) at the output.

(b) *Synchronization of multiple aperiodic streams.*

This is a simple extension of class (a). If the rates of the streams are not constant but follow a more complicated time law, this law will be (approximately) reproduced at the output. This is, for example, what a scientist may need when comparing, for the purposes of validating and tuning a numerical model of a natural phenomenon, a video sequence recording the phenomenon and an image sequence produced by a simulator of the model [John91]. Initially, the speeds of the two streams will have to be matched, and this may require changing the rate of one or both of the streams. Then, the rates of both may have to be modified together according to a given law. Since our stream synchronization approach deals with each stream independently, all of these operations are straightforward: they only entail changing the input rates of those streams whose speeds have to be modified.

(c) *Simultaneous arrivals at different destinations.*

The method works (by definition of jitter) even when the streams to be synchronized have different destinations. This may be required in those applications in which (approximate) simultaneity is important; for example, in distributed music rehearsing, which has very strict delay and jitter requirements.

5. Conclusions

This paper has presented a distributed method for jitter control of wider applicability than the one proposed in [VeZF91]. The method requires synchronized clocks, but can tolerate small divergences between clock readings [Ferr91a]. Both methods can be used together, and may coexist with the absence of jitter control and with destination-based methods, in the same internetwork, the same node, and even the same channel [Ferr91a].

Distributed jitter control reduces the total buffer space requirement of a real-time bounded-delay channel, and spreads those total requirements much more evenly over the nodes of the channel's path. The presence of such a scheme will substantially reduce the maximum jitter perceived by the clients, and will facilitate further destination-based efforts to reduce end-to-end jitter. Careful design of the upper protocol layers and of operating system support mechanisms for real-time applications may make a jitter control service based on the distributed scheme sufficiently accurate for most of those applications.

Only experimentation with protocols based on these schemes in internetworking testbeds will allow us to confirm, or force us to modify, these conjectures. This is the obvious main goal of our future work in this area.

Acknowledgments

The author wishes to express his gratitude to all the members of the Tenet Group at the University of California at Berkeley and the International Computer Science Institute for the extensive discussions of the ideas and results presented in this paper. He is

particularly indebted to Anindo Banerjea, Srinivasan Keshav, Massimo Maresca, Mark Moran, Dinesh Verma, and Hui Zhang for their suggestions, and to Professor Bernd Wolfinger of the University of Hamburg for his extremely careful reading of an earlier version of the manuscript and the many improvements that came from it.

References

[DeKS89] Demers, A., S. Keshav, and S. Shenker, "Analysis and Simulation of a Fair Queueing Algorithm," *Proc. ACM SIGCOMM Conf.*, Sept. 1989, 1-12.

[Ferr91a] Ferrari, D., "Distributed Delay Jitter Control in Packet-Switching Internetworks," Tech. Rept., International Computer Science Institute, Berkeley, October 1991.

[Ferr91b] Ferrari, D., "Real-Time Communication in an Internetwork," in preparation.

[FeVe90] Ferrari, D., and D. C. Verma, "A Scheme for Real-Time Channel Establishment in Wide-Area Networks," *IEEE J. Selected Areas in Communications*, SAC-8, April 1990, 368-379.

[John91] Johnston, W. E., personal communication, May 1991.

[KaKK90] Kalmanek, C. R., H. Kanakia, and S. Keshav, "Rate Controlled Servers for Very High-Speed Networks," *Proc. GlobeCom '90 Conf.*, San Diego, CA, Dec. 1990, 300.3.1 - 300.3.9.

[LiLa73] Liu, C. L., and J. W. Layland, "Scheduling Algorithms for Multiprogramming in a Hard Real Time Environment," *J. ACM*, 20, January 1973, 46-61.

[Mill89] Mills, D., "Measured Performance of the Network Time Protocol in the Internet System," Network Working Group, Request for Comments 1128, October 1989.

[Part91] Partridge, C., "Isochronous Applications Do Not Require Jitter-Controlled Networks," Network Working Group, Request for Comments 1257, September 1991.

[Ravi91] Ravindran, K., "Real-time Synchronization of Multimedia Data Streams in High Speed Packet Switching Networks," Tech. Rept., Dept. of Computing and Information Sciences, Kansas State Univ., February 1991.

[ShSa90] Shepherd, D., and M. Salmony, "Extending OSI to Support Synchronization Required by Multimedia Applications," *Computer Communications*, 13, 7, September 1990, 399-406.

[VeZF91] Verma, D. C., H. Zhang, and D. Ferrari, "Delay Jitter Control for Real-Time Communication in a Packet Switching Network," *Proc. TriComm '91 Conf.*, April 1991.

Fig.1. Model of a channel's path. Packet delays are measured between A and B. Since node 1 (node N) shares a host system with the source S (destination D), times at A (B) may be measured at the input (output) of node 1 (node N).

(1) Deterministic delay bounds on all level 1 links exist and are known or can be obtained by recursive application of an extension of the scheme in [FeVe90].

(2) Level 1 nodes are store-and-forward.

(3) Each output link of a level 1 node and each host system is scheduled by a suitable discipline.

Fig.2. Sufficient conditions for guaranteeing end-to-end deterministic delay bounds in a hierarchical internetwork.