Network Support for Realtime Multi-Party Applications

Amit Gupta, Wendy Heffner, Mark Moran, Clemens Szyperski
{amit,wendyh,moran,szyper}@tenet.berkeley.edu

The Tenet Group
Computer Science Division, Department of EECS
University of California at Berkeley

and

International Computer Science Institute, Berkeley Extended Abstract

1. Introduction

Many important applications of the future require predictable performance from the network service. For example, applications that support multimedia interactive communication between multiple remote participants (MIM applications) have several such requirements: guaranteed throughput, bounded delay and delay-jitter. The Tenet group has implemented a suite of realtime protocols to provide a simplex unicast data service with performance guarantees [FBZ92]. In order to minimize the complexity of implementation, the first scheme (Scheme1) was designed to provide minimal support for realtime communication. Therefore, Scheme1 is cumbersome and inefficient for multi-party communication. In this paper we describe extensions to this work to provide support for multi-party realtime communication and a more flexible client-service interface. We refer to the resulting scheme as Scheme2.

Support for multi-party communication falls in two areas: (i) networking abstractions for multi-party communication (e.g., a of destination group id that allows sources to send data to a dynamically-changing group of interested parties); and (ii) optimizations that improve the efficient utilization of network resources for multi-party communication (e.g. using true multicast).

In order to improve the flexibility of the client-service interface, we allow several input parameters to be specified as ranges rather than point values. The client can then specify a desired value and a worst-acceptable value, thus eliminating the need for iterative negotiations.

The next section will examine the requirements of MIM applications in more detail. Section 3 presents a description of those aspects of the Scheme2 most relevant to supporting MIM applications, along with the rationale for some of the design choices. Section 4 illustrates the scheme with a simple example.

2. The Issues

In recent years, most researchers have agreed that the network service needs to provide support for delay- and loss-sensitive applications (see, e.g., [Fer90] and [CSZ92]). In this paper we will not argue for such support, but assume the communication service will provide guaranteed service, as in the Tenet scheme [FBZ92].

We have identified seven goals to guide the design of Scheme2. The Tenet approach to realtime computer communication, described in detail in [FBZ92], consists of three basic principles: (1) the network should offer a general interface (i.e. a parameterized interface abstraction that provides performance guarantees to realtime applications, including continuous-media applications); (2) the solution should be applicable to a wide variety of internetworking environments; (3) the realtime communication service must provide predictable performance to network clients. In designing a new scheme to support multi-party realtime applications, we have augmented these basic principles with two design goals that are specific to multi-party communication: (4) communication should be scalable to allow hundreds of senders and/or receivers; (5) the network must allow dynamic participation of both senders and receivers, i.e. participants must be able to join and leave a communication session at any time during its lifetime. Finally, we have identified two goals from principles of good system design: (6) the design should be simple, including only those features essential for realtime multi-party communication; and (7) efficient to implement, meaning that we seek to minimize the overhead of the communication system, in terms of both latency and usage of network resources. To avoid the dangers of a paper design we have striven for a balance among these goals. Care has been taken to engineer the new scheme to provide an effective service that is reasonable to implement.

3. The Design

To achieve simplicity in the design, we have introduced a few simple abstractions, which can be composed to support complex multi-party communication. To support efficiency, the design allows clients to supply hints that may allow the network to optimize network resource usage. This section will describe the main features of our design, including the communication model and important abstractions.

Communication Model

We have chosen a 1xN simplex realtime multicast channel as the basic communication model, i.e. a channel allows a single sender to transmit the same data to multiple destinations with guaranteed performance. As described in the original Tenet scheme, a parameterized characterization of the sender's traffic (a traffic or flow specification) is supplied at channel establishment time and enforced by the network. In addition, each destination must specify the quality of service (QoS) that it expects from the network. We should note that in our case, the QoS consists of end-to-end performance quantities, (e.g. delay or loss probability) rather than quantities that would change the traffic specification (e.g. image resolution).

We chose the 1xN model over a MxN model for several reasons: (1) the 1xN model easily maps onto lower level multicast services; (2) it allows the network to perform sender-specific functions without requiring the network to demultiplex an MxN stream (e.g. enforcing traffic specifications or controlling data transmission based on sender ID); (3) it allows a more straight-forward implementation for resource reservations and allows senders to have heterogeneous traffic specifications. A more complex MxN model was considered and rejected because of the increased complexity of implementation. As described below, efficiency of resource usage is obtained using sharing groups, rather than depending on statistical multiplexing of multiple senders. The size of forwarding tables and/or the number of local channel ids (lcids) can be optimized by techniques that are not described in this paper.

Parameter Ranges

Scheme2 allows several QoS parameters to be specified by ranges, rather than point values, i.e. the client specifies a desired value and an acceptable range for several QoS parameters. If the network cannot meet the desired service but can provide service within the acceptable range then channel establishment succeeds. In most cases, the use of ranges avoids the need for negotiations between the network and client.

Target Sets

Senders establish 1xN multicast channels to groups of destinations called Target Sets. A target set consists of a set of tuples indicating the address and QoS specification for each receiver. Since multiple channels can address a single target set, MxN semantics can be easily obtained. As in the host group abstraction of IP multicast ([DeC85], [Dee89]) target sets support large-scale communication and dynamic participation by decoupling senders from receivers.

A destination "join" inserts the new destination into the target set and attempts to include that destination in each active channel sending to that target set. Similarly, a new sender joins an active conference by establishing a channel to an existing target set. Target sets have loose semantics: for any given channel addressing the set, i.e. some destinations may be rejected (termed "partial rejection") during channel establishment. Target set membership indicates an upper bound on actual connectivity, which can be determined by querying the network service. The decoupling of senders and receivers also simplifies the requirements to avoid potential race conditions posed by nearly simultaneous attempts to join or leave a conference.

Sharing Groups

Network-level multicast minimizes throughput and processing requirements, enabling the network to support communication to a large number of destinations. Scaling to large numbers of senders can be increased by multiplexing network resources between traffic from multiple senders. Systems that support a MxN model of communication rely on statistical multiplexing between senders to increase scalability in the number of senders. However, this approach is not consistent with the Tenet approach, since predictable performance can no longer be achieved. Scheme2 introduces Sharing Groups to allow clients to communicate application-specific information, such as an upper bound on the number of concurrently active senders in the group, to the network. Using this information, the network can multiplex traffic from different senders on the same network resources without violating performance guarantees. Sharing groups are defined by a set of channels and a sharing specification. This specification indicates the maximum number of concurrently active channels and the aggregate traffic specification. The network can use this information to bound the resource reservation for channels in the group. By combining target sets and sharing groups, the network can perform other optimizations as well, such as reducing the size of forwarding tables and the number of lcids used by a group of channels.

Secondary and On-demand Channels

In order to increase the efficiency and scalability of communication, we have introduced the notion of secondary and on-demand channels.

Requiring all channels to meet the QoS specified for each destination may require the network to use different routes for traffic from different senders, thereby increasing the routing overhead and decreasing the ability to share resources via sharing groups. Therefore, we introduce a distinction between Primary and Secondary Channels. A primary channel adheres to the QoS requirements requested by the destinations in the target set, as described above. A secondary channel, however, is established without regard to the QoS specified in the target set, allowing the network to route a new secondary channel along the route of an existing primary channel. Once established, a secondary channel achieves a guaranteed QoS as an outcome of the establishment procedure, the difference being that the QoS is not pre-specified.

Another relaxation of the semantics of a channel can be made by marking it as On-Demand as opposed to the Immediate Channels described above. Immediate channels are available for data transmission after a successful Establishment call. On-demand channels, however, must be successfully enabled before the sender can begin transmitting. The distinction between establishment and enabling phase allows the network know about a channel's existence without performing the more expensive establishment procedures (such as resource reservation). The motivation for this is that it is possible to optimize routing and establishment by taking on-demand channels into account before they are needed. In particular, by combining on-demand and secondary channels, the network can greatly reduce the overhead of establishing return channels for a distributed lecture.

4. An Example

Consider a distributed panel discussion in which all the panelists and the audience members are equipped with cameras and microphones, along with speakers and video display devices. The panelists are well-behaved and do not usually interrupt each other. The audience members do not interrupt the speakers (they must receive permission before they start speaking). Due to the nature of the conference, the following relationships hold:

1. No more than 2 panel members may speak at a time.
2. At any instant, only the "current speaker" appears on the video screens.

It is important that the panel members should be able to speak during the discussion; audience members will be willing to participate even if they can not speak during the conference (due to non-availability of the required resources). The discussion is set up with two target sets, one each for audio and video. Although the two sets have the same membership (all participants), the destinations can specify different QoS requirements by using separate target sets.

Given the expected interactive discussion of the panelists, each panelist should create primary immediate channel to both audio and video target sets. Primary channels are used to ensure that the QoS requested by the audience is met. The channels are specified as immediate since all panel members are expected to speak. Each audience member should create secondary, on-demand video and audio channels. The secondary channel implies that the QoS requested by each destination is not considered during routing and establishment, allowing the network to exploit optimizations in shared routing of channels. The on-demand qualifier indicates that this channel cannot be used until it is "Enabled". The reason the channels from the audience members are specified as on-demand is that we don't know whether all audience members will speak, and do not want to pay for resources and overhead to fully establish channels for them until they are needed. The advantage of specifying them as on-demand channels at the beginning of the conference is that they can be routed ahead of time, sharing part or all of the multicast tree of the immediate channels. There are two reasons we don't specify all channels as on-demand: (1) allowing a sender to specify a channel as immediate saves a step in the process; (2) the network will probably route on-demand channels together with immediate channels sending to the same target set. Thus it is helpful to know in advance which channels are certain to be required immediately.

Given the dynamics of the discussion, more efficient use of network resources could be achieved by providing a sharing specification. In this particular case, the audio channels may be grouped together with total resource requirement equal to two audio channels, and the video channels grouped together with total resource requirement equivalent to that of one video channel. The network enforces rate control based on the sharing specification independent of the client policy and mechanisms to orchestrate floor control.

5. Summary and Status

The combined consideration of realtime and multi-party communication opens an interesting area of research that is highly relevant for emerging multimedia conferences applications. The Tenet Scheme2 builds on the realtime capabilities of the original scheme and adds efficient and flexible support for multi-party communication.

We are currently designing a protocol suite for Scheme2. This work encompasses designing algorithms, validating them by simulation, and devising the system architecture. Implementation will begin in Fall 1993. In addition, prototypical conferencing tools will be compared on both the Internet and the Tenet protocol suites.

References:

[CSZ92] D. Clark, S. Shenker, L. Zhang, “Supporting Real-time Applications in an Integrated Services Packet Network: Architecture and Mechanism”, Proceeding of ACM SIGCOMM ‘92, Baltimore, Maryland, August 1992.

[DeC85] S. Deering, D. Cheriton, “Host Groups: A Multicast Extension to the Internet Protocol” Request for Comments, RFC966, December 1985.

[Dee89] S. Deering, “Host extensions for IP multicasting” Request for Comments, RFC1112, August 1989.

[Dee91] S. Deering, “Multicast Routing in a Datagram Internetwork,” Ph.D. Thesis, Stanford University, December 1991.

[Fer90] D. Ferrari, “Client Requirements for Real-Time Communication Services”, IEEE Communications Magazine, vol. 28, n. 11, November 1990. [FBZ92] D. Ferrari, A. Banerjea, H. Zhang, “Network Support for Multimedia - A Discussion of the Tenet Approach”, Technical Report TR-92-072, International Computer Science Institute, Berkeley, October 1992.

[PPA92] J. Pasquale, G. Polyzos, E. Anderson, V. Kompella, “The Multicast Multimedia Channel”, Proc. Third International Workshop on Network and Operating Systems Support for Digital Audio and Video, San Diego, California, November 1992.

[ZDE93] L. Zhang, S. Deering, D. Estrin, S. Shenker, D. Zappala, “RSVP: A New Resource ReSerVation Protocol”, Draft available via anonymous ftp from parc.xerox.com, March 1993.