Stardust Technologies, Inc
1901 S. Bascom Ave, #333
Campbell, CA 95008 USA
Tel: 408-879-8080
Fax: 408-879-8081
Web: www.ipmulticast.com

IP Multicast Initiative (IPMI)

Higher Level Protocols used with IP Multicast

An IP Multicast Initiative White Paper

An introduction to selected protocols used with IP Multicast to support multimedia and reliable data transmission

Rev: 021197-02 - HighProt.doc

Stardust Technologies, Inc., 1901 S. Bascom Ave., Suite 333, Campbell, CA 95008

Higher Level Protocols used with IP Multicast

An introduction to selected protocols used with IP Multicast to support multimedia and reliable data transmission

Scope Of This Document

This document provides a brief technical overview of several higher level protocols used with IP Multicast. First, protocols that have been developed to support real-time multimedia delivery and Quality-of-Service (QoS) specifiers for multicast and unicast network services are presented. These include the Real-time Transport Protocol (RTP), the control protocol (RTCP) that works in conjunction with RTP, the resource reservation protocol (RSVP), and the real-time streaming protocol (RTSP). Quality-of-service routing is briefly considered. Next, reliable IP Multicast protocols, an emerging standards area, are discussed. Lastly, group setup protocols are introduced. The protocols described in this paper are at different levels of maturity.

If you are an engineer interested in evaluating or implementing IP Multicast for your organization, product or service, this document will help you understand how these protocols enhance the usefulness of IP Multicast. It assumes you have a conceptual understanding of IP Multicast addressing, group management and routing. If you are unfamiliar with these, we recommend you first read the IP Multicast Initiative white papers “How IP Multicast Works” and “Introduction to IP Multicast Routing”.

Need for additional protocols

The Internet has been used primarily for the reliable transmission of data with minimal or no delay constraints. The TCP/IP protocols were designed for this type of traffic and work very well in this context. However, multimedia traffic, which comprises a significant portion of potential multicast traffic, possesses different characteristics and hence requires the use of different protocols to provide the necessary services. For example, if a receiver has to wait for a TCP retransmission, there can be a noticeable and unacceptable gap in playout of the real-time data, whether audio, video, or some other delay-sensitive data. In addition, the "slow start" TCP congestion-control mechanism can interfere with the audio and video "natural" playout rate. Since there is no fixed path for datagrams to flow across the Internet, there is no mechanism for ensuring that the bandwidth needed for multimedia is available between the sender and receiver(s), so quality of service cannot be guaranteed. In addition, TCP doesn't provide timing information, a critical requirement for multimedia support.

Multimedia applications can generally forego the complexity of TCP and use instead a simpler transport framework. Most playback algorithms can tolerate missing data much better than lengthy delays caused by retransmissions, and they do not require guaranteed in-sequence delivery. A number of protocols have been developed to enhance the Internet architecture and improve support of applications like audio, video and interactive multimedia conferencing. Introduced below are the RTP, RTCP, RSVP and RTSP protocols. These real-time oriented protocols are designed to be used over both multicast or unicast network services. Since many real-time applications can conserve network and server resources by using IP Multicast, the special requirements and characteristics of IP Multicast have been considered in the design of these protocols, such as scalability, multicast routing, and accommodation of large numbers of receivers and heterogeneous receivers.

Following the discussion of these protocols for real-time multimedia, protocols for reliable IP Multicast are considered. Reliable delivery is required by many real time and non-real time applications. For unicast IP services, error correction and detection in the TCP layer provides reliability. For reliable multicast, new approaches to tracking acknowledgments and detecting and correcting errors have been developed, since an IP Multicast sender may have thousands of recipients.

Real-time Transport Protocol (RTP)

RTP (version 2) is a real-time transport protocol that provides end-to-end delivery services to support applications transmitting real-time data, e.g., interactive audio and video, over unicast and multicast network services. RTP is defined in IETF RFC 1889, along with a profile for carrying audio and video over RTP in RFC 1890. Both are IETF Proposed Standards, and are expected to be finalized in early 1997. RTP is used on the MBONE by *vat*, the video/audio tool. In addition commercial implementations of RTP and applications that use RTP are currently available for a number of platforms.

RTP services include payload type identification, sequence numbering, and time stamping. Delivery is monitored by means of a closely integrated control protocol called RTCP (see next section).

RTP provides end-to-end delivery services, but it does not provide all of the functionality that is typically provided by a transport protocol. In fact, RTP typically runs on top of UDP to utilize its multiplexing and checksum services. Other transport protocols besides UDP can carry RTP as well.

The RTP header provides the timing information necessary to synchronize and display audio and video data and to determine whether packets have been lost or have arrived out of order. In addition, the header specifies the payload type, thus allowing multiple data and compression types. RTP is tailored to a specific application via auxiliary profile and payload format specifications. As an example, a payload format might specify what type of audio or video encoding is carried in the RTP packet. Encoded data can be compressed before delivery.

To set up an RTP session, the application defines a particular pair of destination transport addresses (one network address plus a pair of ports for RTP and RTCP). In a multimedia session, each medium is carried in a separate RTP session, with its own RTCP packets reporting the reception quality for that session. For example, audio and video would travel on separate RTP sessions, enabling a recipient to select whether or not to receive a particular medium.

An audio-conferencing scenario presented in RFC 1889 illustrates the use of RTP. Suppose each participant sends audio data in segments of 20 ms duration. Each segment of audio data is preceded by an RTP header, and then the resulting RTP message is placed in a UDP packet. The RTP header indicates the type of audio encoding that is used, e.g., PCM. Users can opt to change the encoding during a conference in reaction to network congestion or, for example, to accommodate low-bandwidth requirements of a new conference participant. Timing information and a sequence number in the RTP header are used by the receivers to reconstruct the timing produced by the source, so that in this example, audio segments are contiguously played out at the receiver every 20 ms.

RTP does not provide any mechanisms to ensure timely delivery or provide quality-of-service guarantees. It does not guarantee delivery or prevent out-of-order delivery, nor does it assume that the underlying network is reliable. Some adaptive applications do not require such guarantees, but for those that do, RTP must be accompanied by other mechanisms to support resource reservation and to provide reliable service.

Real-time Control Protocol (RTCP)

RTCP is the control protocol that works in conjunction with RTP. RTCP control packets are periodically transmitted by each participant in an RTP session to all other participants. Feedback of information to the application can be used to control performance and for diagnostic purposes.

RTCP performs the following four functions.

1) Provide information to application:

The primary function is to provide information to an application regarding the quality of data distribution. Experiments with IP multicasting have established the importance of user feedback from RTCP to diagnose distribution faults. Each RTCP packet contains sender and/or receiver reports that report statistics useful to the application. These statistics include number of packets sent, number of packets lost, interarrival jitter, etc. This reception quality feedback will be useful for the sender, receivers, and third-party monitors. For example, the sender may modify its transmissions based on the feedback; receivers can determine whether problems are local, regional or global; network managers may use information in the RTCP packets to evaluate the performance of their networks for multicast distribution.

2) Identify RTP source:

RTCP carries a transport-level identifier for an RTP source, called the canonical name (CNAME). This CNAME is used to keep track of the participants in an RTP session. Receivers use the CNAME to associate multiple data streams from a given participant in a set of related RTP sessions, e.g., to synchronize audio and video.

3) Control RTCP transmission interval:

To prevent control traffic from overwhelming network resources and to allow RTP to scale up to a large number of session participants, control traffic is limited to at most 5 percent of the overall session traffic. This limit is enforced by adjusting the rate at which RTCP packets are transmitted as a function of the number of participants. Since each participant sends control packets to everyone else, each can keep track of the total number of participants and use this number to calculate the rate at which to send RTCP packets.

4) Convey minimal session control information:

As an optional function, RTCP can be used as a convenient method for conveying a minimal amount of information to all session participants. For example, RTCP might carry a personal name to identify a participant on the user’s display. This function might be useful in loosely controlled sessions where participants informally enter and leave the session.

Real-Time Streaming Protocol (RTSP)

The application-level Real Time Streaming Protocol, RTSP, aims to provide a robust protocol for streaming multimedia in one-to-many applications over unicast and multicast, and to support interoperability between clients and servers from different vendors. Its draft specification is in the very early stages of submission to the IETF. Products using RTSP are available today (even though RTSP is likely to undergo significant change as it goes through the IETF process).

“Streaming” breaks data into many packets sized appropriately for the bandwidth available between the client and server. When the client has received enough packets, the user software can be playing one packet, decompressing another and receiving a third. The user can begin listening almost immediately without having to download the entire media file. Sources of data for streaming can include both live data feeds and stored clips.

RTSP is considered more of a framework than a protocol. It is intended to control multiple data delivery sessions, provide a means for choosing delivery channels such as UDP, TCP, IP Multicast, and delivery mechanisms based on RTP. Control mechanisms such as session establishment and licensing issues are being addressed. RTSP is being designed to work on top of RTP to both control and deliver real-time content. Thus RTSP implementations will be able to take advantage of RTP improvements, such as the new standard for RTP header compression. Although RTSP can be used with unicast in the near future, its use may help smooth the transition for environments transitioning from unicast to IP multicasting with RTP.

RTSP can be used with RSVP to set up and manage reserved-bandwidth streaming sessions.

Resource Reservation Protocol (RSVP)

While RTCP can provide feedback on reception quality, other protocol mechanisms are needed to request timely delivery and guarantee a specific QoS for a session between a sender(s) and receiver(s).

The Reservation Protocol (RSVP) is a resource reservation setup protocol designed for an integrated services internetwork. An application invokes RSVP to request a specific end-to-end QoS for a data stream. RSVP aims to efficiently set up guaranteed QoS resource reservations which will support unicast and multicast routing protocols and scale well for large multicast delivery groups. RSVP is currently being defined by the IETF. The number of products incorporating RSVP is growing, and at the time of this writing RSVP is being offered by at least one major ISP in conjunction with its multicast services, with interest expressed by others.

A host receiver uses RSVP to request a specific QoS from the network for a particular data stream(s) from a data source. An elementary RSVP reservation request consists of a specification for an end-to-end desired QoS (e.g. peak/average bandwidth and delay bounds) and a definition of the set of data packets to receive the QoS. RSVP is useful for environments where QoS reservations can be supported by reallocating rather than adding resources. For multicast, a host sends IGMP messages to join a host group and then sends RSVP messages to reserve resources along the delivery path(s) of that group.

RSVP provides access to internetwork integrated services, where hosts and networks work in concert to achieve a guaranteed quality of end-to-end transmission. All the hosts, routers and other network infrastructure elements between the receiver and sender must support RSVP. They each reserve system resources such as bandwidth, CPU and memory buffers in order to satisfy the QoS request. It is expected, therefore, that premiums will be charged by ISP's for RSVP QoS reservations. Approaches to handling bandwidth reservation and billing across multiple carrier networks still need more definition.

RSVP QoS control request messages are sent to reserve resources along all the nodes (routers and hosts) along the (reverse) delivery path to the sender. Note that RSVP is receiver-initiated -- RSVP requests resources in only one direction. For multicast, the reservation request need only travel to a point where it merges with another reservation for the same source stream. This receiver-oriented design is intended to accommodate large multicast groups and dynamic group membership.

Figure 1: Merging IP Multicast RSVP requests

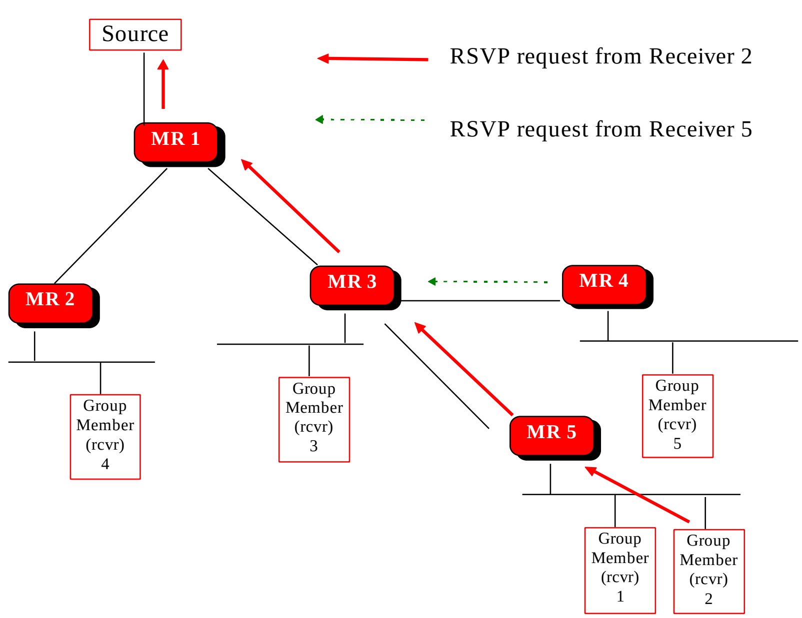

The RSVP request for Receiver 5 is merged at Multicast Router 3 (MR 3) with the earlier RSVP request made by Receiver 2.

The request is not propagated on to MR 1.

At each node along the receiver-sender path, RSVP attempts to make a resource reservation for the requested stream. At each intermediate node, two general actions are taken on a request:

1) Make a reservation

> The request is granted or rejected according to admission and policy controls. These controls utilize information from underlying integrated services mechanisms which are not part of RSVP. The admission control determines whether the node has sufficient resources available, and the policy control determines whether the user has authorization to make a reservation. If the reservation is rejected, RSVP returns an error message to the appropriate receiver(s). If the reservation can be accommodated, the node configures a packet classifier to select the appropriate incoming data packets and a packet scheduler to achieve the desired QoS on the outgoing interface. One RSVP session might have priority over others.

2) Forward the request upstream

> The reservation request is propagated to nodes upstream towards the appropriate senders.

RSVP also defines teardown messages that systems use to relinquish resources.

RSVP operates on top of IP (either IPv4 or IPv6), occupying the place of a transport protocol in the protocol stack, but provides session layer services (it does not transport any data). RSVP is an Internet control protocol (like IGMP or ICMP), not a routing protocol. It uses underlying routing protocols to determine where it should carry reservation requests. As routing paths change, RSVP adapts its reservation to new paths if reservations are in place. The RSVP protocol is used by routers to deliver QoS control requests to all nodes along the path(s) of the flows and to establish and maintain state to provide the requested service. After a reservation has been made as described above in item 1, routers supporting RSVP determine the route and the QoS class for each incoming packet and the scheduler makes forwarding decisions for every outgoing packet.

One drawback of RSVP is the computational resources required for routers to inspect and handle these packets in a priority order. Approaches such as tag switching are being developed to alleviate this. Another area of research is enhancing RSVP to use routing services that provide alternate and fixed paths.

RTP can complement RSVP by allowing applications to respond to the underlying network performance.

Figure 2. Use of RSVP and RTP for a multimedia IP Multicast application

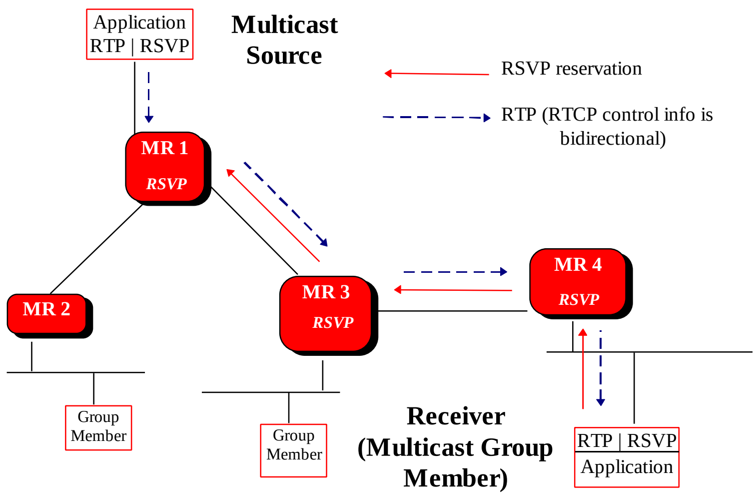

For multimedia, the audio and video are carried in a separate RTP session with RTCP packets controlling the quality of the session. Routers communicate via RSVP to set up and manage reserved-bandwidth sessions.

Vendors have implemented RSVP both above and below Winsock. RSVP-aware applications can be developed with Winsock 2 which has a QoS-sensitive API. As Windows NT 4.0 and future Winsock 2 versions of Windows 95 are deployed this approach will become practical. Another approach is to use an RSVP proxy that runs independently of the real application, making RSVP reservations.

Routing and QoS

Multicast routing protocols were presented in the IP Multicast Initiative white paper “Introduction to IP Multicast Routing”. These protocols use various approaches to construct multicast delivery trees for efficient multicast transmission. The resultant delivery trees provide a route from a multicast sender to all the receivers; however, without additional mechanisms, such routing is not guaranteed to provide a specified quality of service. In this section we discuss the necessity of quality-of-service (QoS) based routing and current work to address this issue.

QoS considerations for a multicast application include tolerance to jitter, delay, and lost packets. In order for a network to satisfy an application’s QoS requirements, the application must be able to reserve and control network resources. This implies that finding the shortest path to a destination when constructing a delivery tree is no longer sufficient; the route must satisfy other QoS requirements as well. This adds considerable complexity to the routing task.

A common metric used by routing algorithms is a count of the number of hops in the path from the source to the destination. In order to satisfy QoS requirements when setting up a path, parameters such as delay, bandwidth, loss probability, cost, and jitter must be considered as well. Since it is difficult to satisfy multiple constraints simultaneously, a possible approach would be to define an order in which the various parameters are evaluated and to evaluate the second parameter, for example, only for those paths considered best with respect to the first parameter.

The issue of incorporating QoS routing with the various multicast routing protocols is not yet resolved. Given sufficient bandwidth, e.g., within a local intranet, quality of service is not likely to be a problem and so QoS routing would not be required.

Reliable IP Multicast

IP Multicast enables applications to significantly reduce the load on network resources and to scale to higher levels. For example, an application which broadcasts information to hundreds of recipients in a corporate network can employ reliable multicast services to reduce network load while maintaining reliable delivery. Similarly, a corporate trainer can use a multicast enabled shared whiteboard application to deliver a course to employees around the world.

Reliable services ensure the sender that all packets are received by all of the recipients. Reliable delivery is required by many real time and non-real time applications. In the real time area, data conferencing, web services and data broadcast applications use reliable services. Non real-time applications such as information and software distribution, and file transfer also need reliable services.

For unicast IP services, error correction and detection in the TCP layer provides reliability. Such traditional techniques for error detection and correction in a large scale multicast environment might result in an overload of acknowledgments to the sender, which would increase network congestion. Reliable multicast protocols provide error correction schemes which are designed to overcome the limitations of unreliable multicast datagram delivery without burdening the network. There are many approaches to reducing the number of acknowledgments in a reliable multicast service. The error correction mechanisms can vary, depending on application requirements and the multicast environment. Application and network characteristics which may impact the design and operation of reliable multicast protocols include:

- Real-time requirements
- Single or multiple senders
- Number of multicast groups and recipients per group

- Scalability as the number of senders, groups and recipients grows
- Ordered or unordered packet delivery
- Delay tolerance
- Average bandwidth requirements
- Continuous or bursty bandwidth usage
- LAN or WAN operation
- Network infrastructure, such as satellite or multi-hop terrestrial links. Some protocols are designed to handle asymmetrical data channels, others do not.

Mechanisms employed by reliable multicast services include multicasting ACKs to reduce duplicate requests, using separate polling services and holding information for retransmission at intermediate points. Some solutions use a pre-defined error correction protocol, while others support the flexibility to choose from a set of reliable protocols.

Applications supporting both simultaneous reliable unicast and reliable multicast transmission can be useful, enabling network managers to gradually transition from a unicast to a multicast network.

Several Internet drafts have been submitted related to reliable multicasting (see references), and an IRTF (Internet Research Task Force) has recently been formed to advance reliable multicast standards efforts. No IETF standards have yet been established for reliable IP Multicast. Several vendors offer proprietary reliable IP Multicast solutions designed for different applications and network environments. Developers and network engineers should select a reliable multicast protocol which optimizes benefits for their applications.

Surveys of reliable IP Multicast protocols have been conducted as part of the DARPA Multicast Implementation Study (MIST) Project and a recent SIGCOMM workshop (see reference section). Vendors can help you identify suitable reliable multicast products for your needs. See also the IP Multicast Initiative vendor product and services directory.

Group Setup Protocols

IGMP, the Internet Group Management Protocol described in RFC 1112, is used by hosts to join IP Multicast sessions and by routers to communicate information about group members on their directly attached subnets. But how does a user or application learn about forthcoming IP Multicast sessions? Out of band mechanisms such as e-mail or web sites can be used for announcements, but there is clearly a need for mechanisms for announcing sessions, determining (temporary) multicast addresses and ports for sessions, issuing invitations (e.g. for conferences), negotiating parameters such as membership, media encodings and encryption keys, for adding or deleting members during the session, and other control functions.

The IETF Multiparty Multimedia Session Control (mmusic) Working Group is chartered to design protocols for the management and coordination of multiple sessions, and their multiple users, in multiple media (e.g., audio, video). Their efforts have been primarily directed at conferencing over the MBONE. They have defined a number of protocols described in Internet Drafts, such as the Session Description Protocol, Session Directory Announcement Protocol and the Simple Conference Control Protocol (see references). None are yet standardized.

Parameters for a session can include the name and topic of the session and its multicast address; date, time and duration; media types (e.g. audio), media encodings and media ports; security parameters; and so forth. There are different ways to set up and tear down groups, and different models for doing it. A conference setup protocol, for example, uses a many to many model. A one to many model is used by permanent guides (analogous to TV Guides): a listener can join the “guide” group to determine IP Multicast sessions of interest.

On the MBONE, for example, the tool *SD* (Session Directory) can be used by MBONE users to reserve and allocate media channels, and to view advertised channels. SD makes announcements periodically over a well-known multicast address and port.

Some IP Multicast application protocols, such as for reliable IP Multicast, might include a group setup protocol, while others rely on external mechanisms or protocols. In general, some of the group setup functions described above may be needed as an adjunct to IP Multicast applications and should be considered during application design. This is an emerging area in protocol development efforts.

More information

IP Multicast Initiative

There are many other technical aspects of IP Multicast and business benefits that were not discussed here. The IP Multicast Initiative web site at www.ipmulticast.com has a technical resource center that provides more background in-depth information including white papers and relevant RFC’s. The web site also offers a vendor product and services directory and lists members of the IP Multicast Initiative who can be contacted for information and assistance.

The IP Multicast Initiative provides marketing and educational services to promote the creation, use and deployment of multicast products and solutions. Supported by a growing number of the most important vendors in the IP Multicast arena, the Initiative and its services are managed and provided by Stardust Technologies, Inc.

For more information about the Initiative and membership, or to see other white papers, contact Stardust Technologies at 408-879-8080 or visit the Initiative web site.

Useful References

The following references are recommended; many were used in the preparation of this document. The authors are gratefully acknowledged.

BOOKS

MBONE: Interactive Media on the Internet, Vinay Kumar, New Riders, 1996

Gigabit Networking, Craig Partridge, Addison-Wesley, 1994

IPng and the TCP/IP Protocols, Stephan Thomas, Wiley, 1996

WEB SITES

http://www.ipmulticast.com IP Multicast Initiative

http://www.cs.columbia.edu/~hgs/rtp/ RTP information

http://www.isi.edu/div7/rsvp/rsvp-home.html RSVP information

http://www.realaudio.com/prognet/rt/ RTSP information

http://www.tascnets.com/mist/doc/mcpCompare.html
MIST Reliable Multicast Protocols Survey

http://gaia.cs.umass.edu/sigcomm_mcast/talk1.html
Overview of Reliable Multicast Protocols from the 8/96 ACM SIGCOMM Multicast Workshop

IETF RFCs [http://ds.internic.net/rfc/rfcnnnn.txt, nnnn is the RFC number]

RFC 1889 RTP: A Transport Protocol for Real-Time Applications
RFC 1890 RTP Profile for Audio and Video Conferences with Minimal Control
RFC 1458 Requirements for Multicast Protocols
RFC 2090 TFTP Multicast Option
RFC 1301 Multicast Transport Protocol

IETF Internet Drafts [ftp://ietf.org/internet-drafts/name-of-file]
(works in progress)

Resource ReSerVation Protocol (RSVP) Version 1 Functional Specification
[draft-ietf-rsvp-spec-14.txt, .ps] (see also other *rsvp* drafts)

Real Time Streaming Protocol (RTSP) [draft-ietf-mmusic-rtsp-00.txt]

IETF Criteria For Evaluating Reliable Multicast Transport and Application Protocols [draft-mankin-reliable-multicast-00.txt]

RTP extension for Scalable Reliable Multicast
[draft-parnes-rtp-ext-srm-01.txt]

SDP: Session Description Protocol [draft-ietf-mmusic-sdp-02.txt, .ps]

Simple Conference Control Protocol [draft-ietf-mmusic-sccp-00.txt]

SAP: Session Announcement Protocol [draft-ietf-mmusic-sap-00.txt, .ps]

SIP: Session Initiation Protocol [draft-ietf-mmusic-sip-01.txt, .ps]

StarBurst Multicast File Transfer Protocol (MFTP) Specification
[draft-miller-mftp-spec-00.txt]

ARTICLES

"Multimedia over IP: Specs Show the Way, " Judy Estrin and Stephen Casner, Data Communications, August 1996, pp. 93-98.