The Multi-Session Bridge

G. A. Roediger, W. P. Lidinsky [1,2]

HEP Network Resource Center
Fermi National Accelerator Laboratory, Batavia, IL 60510, USA

The Multicast BackbONE (MBone) is heavily used to carry Internet video conferences. The low cost and ease-of-use have made the MBone and its associated tools such as sdr, vat, vic, wb very popular. But HEP video conferences often requires many-to-many sparse network topologies, an arrangement that cannot easily take advantage of MBone efficiencies. Moreover, many HEP sites do not support multicast and thus researchers at those sites are cut off from important meetings. In addition, the topology of the MBone can cause poor performance or lack of access for some sites. To solve these problems, HEPNRC has developed the Multi-Session Bridge (msb) and its companion tool vtc. The msb allows an arbitrary number of multicast and unicast sessions to be connected in one or more seamless video conferences. Multiple msbs can also be connected to form a conference virtual network. The msb also extends MBone sessions to non-MBone sites through unicast streams. The msb is being extended to provide secure communications.

Key words: Video, conference, MBone, multimedia, collaboration.

1 Introduction

Because of rapid changes in enabling technologies, computing and computer networking seems to go through major shifts in direction and focus every few years. One recent shift has been to the World Wide Web. Another has been toward multimedia. Multimedia looms even more important than the Web,

[1] Supported in part by the US Department of Energy under contract number DE-AC02-76CH03000.

[2] Extensive involvment and consultation were made by the following collaborators: D. Martin and H. Kippenhan of HEPNRC at Fermilab, S. Colson of Fermilab; and T. Thomas, U. of New Mexico.

and has been described as the holy grail of computer networking over which both propeller heads and suits salivate. It has the attention of both technology and money. A major sub-field of the multimedia milieu is video conferencing.

More importantly there is a need for video conferencing. Increasingly people separated geographically need to communicate in a rich manner similar to the way they would communicate if they were in the same room. The high energy physics community is only one discipline that now finds the ability to video conference indispensable.

Today there are two technologically distinguishable types of video conferencing differentiated by the network over which the information flows - circuit or packet. In circuit switched networks calls are set up (i.e., resources reserved) end-to-end, eliminating interference among network traffic competing for the resources while the conference is in progress. Quality of service can be guaranteed at call setup, or alternately the connection can be denied. Performance can be predicted.

In contrast, packet switching such as that of the Internet allows resources to be shared at all times. During a video conference, network traffic unrelated to the conference may use the very same network resources needed by the video conference. A consequence of this can be that packets of audio and video are delivered late and cannot be used. This can give rise to poor quality conferencing. Audio, which is the most time critical, can quickly become unusable.

2 Packet Video Conferencing

For Internet-style packet video conferencing both proprietary and public systems have been developed. In the public domain, there exists an overlay network and a collection of tools. The overlay network, called MBone, is needed to efficiently use network resources by intelligently moving multicast traffic. Recent MBone improvements (e.g., pruning) currently being deployed further improve efficiency by helping eliminate unneeded multicast traffic.

The public tools include sdr for setting up, coordinating, and attaching to conferences; nv and vic for sending and receiving video; vat, rat, and nevot for sending and receiving audio; and ivs for sending and receiving both video and audio. Wb and nte for enable several participants to type and draw concurrently on their own workstation screen, each seeing what the other participants have written or drawn. A recently adopted Internet protocol called RTP2 facilitates the use of real-time software. Modification of the tools to use this protocol has further aided real-time synchronization of packets.

3 Multi-Session Bridge (msb)

3.1 Background

The public domain multimedia tools using the MBone overlay network have been in use since the early 1990s. While amenable to one-to-one and one-to-many densely attended conferences, MBone can be less efficient for certain types of sparse many-to-many sessions characteristic of HEP video conferences. Additionally there are many HEP sites that do not support MBone broadcasts and thus researchers at those sites are cut off from important meetings. Thus the topology of the MBone can also cause poor performance or lack of access for some sites. For these reasons HEPNRC at Fermilab has developed the Multi-Session Bridge (msb).

The network topology of MBone efficiently switches traffic among statically configured domain clusters. HEP conferences, because they are irregular and sparse, cannot easily take advantage of these efficiencies. In contrast, the msb builds conference network topologies for the duration of the conference based on only IP routing. Further, it allows control of conference attendance from a central conference administrator and provides access by conference sites that do not have multicast access.

Version 1 of the msb and vtc (called vtconnect) have been used extensively by major HEP collaborations since late 1995. Recently version 2.01 of msb and vtc have been completed and are being deployed.

3.2 System

The msb, vtc, and the multimedia tools form a client-server architecture with the msb software residing on a server machine or machines. Vtc and the multimedia tools reside on each client. The msb can send and receive both unicast and multicast video, audio, and data streams from client workstations and unite these streams in a video conference. Figure 1 shows a single msb. Here the msb runs on a workstation at Fermilab. Experimenters located at Fermilab, the University of New Mexico, and DESY are participating in a session. One or more experimenters are at each client workstation. Each client "attaches" to the session by running vtc with the appropriate session name, msb-host name and password. Vtc communicates the information to the msb which validates the session name and password against a previously established configuration. It returns to vtc the approval to connect and identifies the multimedia tools that are needed. Vtc then starts the tools, thus including the client workstation in the session. All this is transparent to the user.

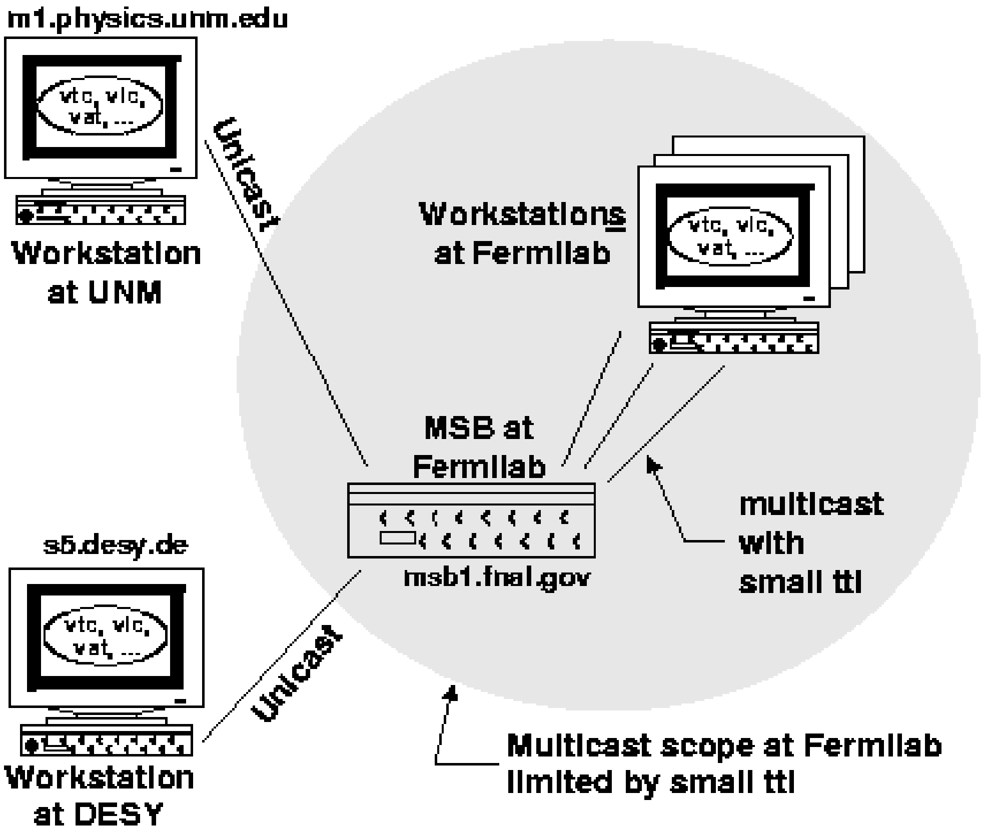

Fig. 1: Single MSB - Both Unicast & Multicast Clients

In Figure 1 the multicast time-to-live (ttl) is set so that workstations at Fermilab are within the multicast scope, thus allowing anyone at these workstations to participate if they know the session name, server name and password. Multicast packets are kept from leaving the site by the ttl scoping. Unicast streams are sent to and received from DESY and UNM. Also note that an msb can run multiple concurrent sessions.

Multiple msbs can be used to further improve the efficiency and cost effectiveness of sparse video conferences. This is especially effective for sessions with clusters of participants in widely separated locations as is shown in Figure 2 where the session of Figure 1 is joined by experimenters at KEK, CERN, Rutgers University, and the University of Tsukuba. Here the msbs work together so that only one information stream is sent across each ocean. Yet everyone participates.

This can be thought of as creating individual virtual networks for each session. Multiple msbs functioning on behalf of a conference create a conference virtual network that is both dynamic and transient. Simplex data streams are set up between end users participating in a meeting possibly using multiple

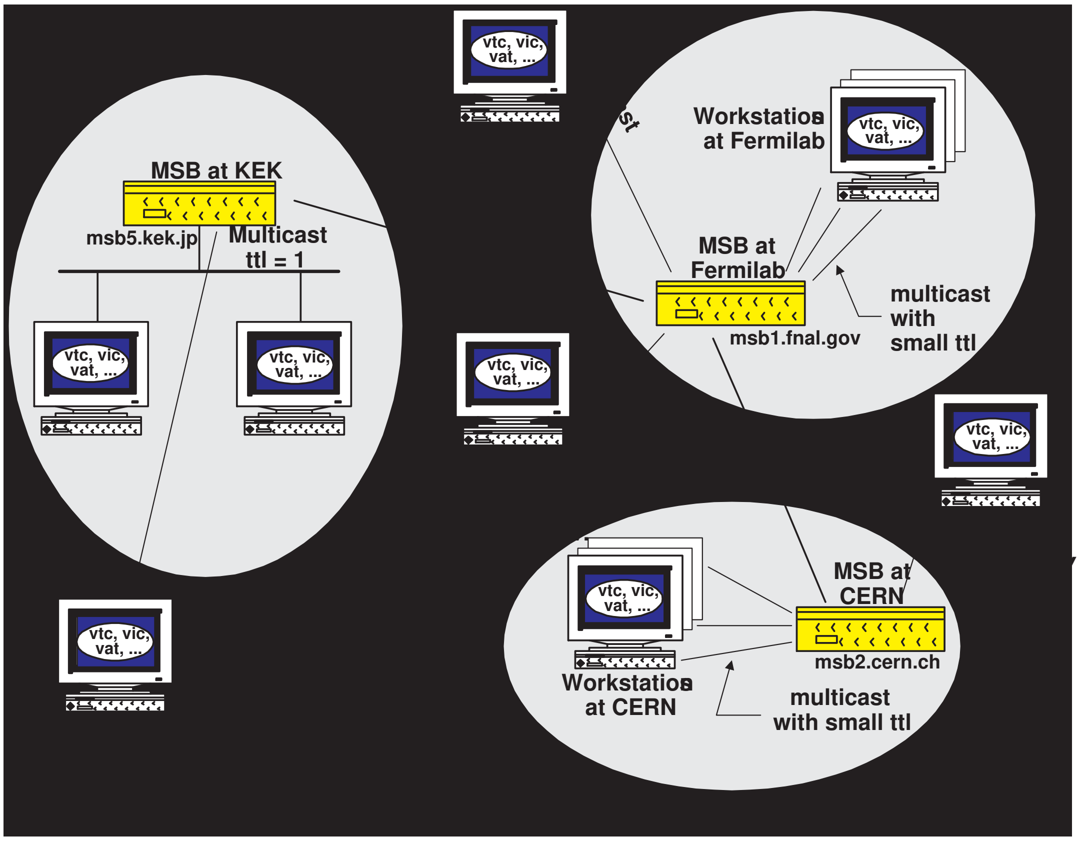

Fig 2: A Multiple MSB Virtual Network

conferencing tools. The bridges set up multiple independent meetings each with its own virtual network and security considerations.

3.3 Managing and Attending Video Conference Sessions

A video conference session is managed from one location. The person (or program) configuring a meeting will create a list of valid attendees, the capabilities of each attendee, and a virtual network of msbs to route the data streams. In Fig. 1 a meeting coordinator would define the meeting in terms of an a msb configuration file. An example file is:

```
meeting name=fnal-lecture user_default_password=x123
    vat 1 meeting=fnal-lecture ip=224.4.111.124 ttl=15
    vat 2 meeting=fnal-lecture ip=m1.physics.unm.edu
    vat 3 meeting=fnal-lecture ip=s5.desy.de

    bridge fnal-lecture 1 to 2 3
    bridge fnal-lecture 2 to 1 3
    bridge fnal-lecture 3 to 1 2

    vic 10 meeting=fnal-lecture ip=224.4.111.124 ttl=15
```

```
vic 20 meeting=fnal-lecture ip=m1.physics.unm.edu
vic 30 meeting=fnal-lecture ip=s5.desy.de

bridge fnal-lecture 10 to 20 30
```

This configures a lecture originating at FNAL that is mbone multicast at the Fermilab site, unicast to the workstation m1.physics.unm.edu at University of New Mexico, and unicast to workstation s5.desy.de at DESY. The first line defines the meeting name and its password. Each tools data stream for a meeting is given an unique number and this number is then used to refer to that stream in the configuration file such as in the bridge statements. The vat audio data streams are defined in lines 2 to 4. The audio streams from all sites are bridged to all other configured sites to allow the lecture to be broadcast and remote sites the ability to ask questions. This is indicated by the bridge statements on lines 5 to 7. The video portion of the meeting is defined by the vic data streams on lines 8 to 10. Since the only source of video is from FNAL only one bridge statement (line 11) is needed to describe its bridging to the two remote sites (video streams 20 and 30). Attendees at the FNAL, UNM, or DESY sites type the following command to participate in the meeting:

```
vtc fnal-lecture msb1.fnal.gov x123
```

Vtc is a client companion program of the msb server code. The first argument on the vtc command line is is the meeting name, second is the ip name or address of the bridge code machine, and last is the meeting password for this user. This command is run from a meeting configured workstation (e.g., ip=m1.physics.unm.edu) and the appropriate commands (vic and vat) with the correct arguments are spawned for the end user. The msb bridges only the data streams to the remote site after vtc is executed and informs the bridge the commands are running. If the remote user terminates one of the video conferencing tools, vtc informs the msb and that simplex data stream is no longer sent. For meeting that do not wish to restrict attendees to a particular workstation the data stream could be configured with ip=0.0.0.0 or INADDR_ANY.

Figure 2 represents the same meeting but with a broader attendance. Physicists at KEK and CERN will also attend along with single participants from University of Tsukuba and Rutgers University. Here it is decided that the meeting will not be restricted to specific workstations as long as the attendee knows the meeting password. The configuration for Figure 2 is then:

```
meeting name=fnal-lecture peer_msbs=msb5.kek.jp,msb2.cern.ch
   vat 1 meeting=fnal-lecture ip=224.4.111.124 ttl=15 password=fnal-123
   vat 2 meeting=fnal-lecture ip=224.4.222.222 ttl=4 password=fnal-xxx
   vat 3 meeting=fnal-lecture ip=0.0.0.0 password=fnal-123
```

```
bridge fnal-lecture 1 to 2 3
bridge fnal-lecture 2 to 1 3
bridge fnal-lecture 3 to 1 2

vic 10 meeting=fnal-lecture ip=224.4.111.124 ttl=15 password=fnal-123
vic 20 meeting=fnal-lecture ip=224.4.222.222 ttl=4 password=fnal-xxx
vic 30 meeting=fnal-lecture ip=0.0.0.0 password=fnal-123

bridge fnal-lecture 10 to 20 30
```

This configuration is given only to the msb running as the root msb for the meeting (msb1.fnal.gov). The two other machines which are running a copy of the msb server at msb5.kek.jp and msb2.cern.ch, receive the fnal-lecture meeting configuration from the root meeting machine. To achieve the meeting topology in Figure 2 the attendees at the sites would type the following commands:

At DESY:	vtc fnal-lecture msb2.cern.ch fnal-123
At CERN:	vtc fnal-lecture msb2.cern.ch fnal-xxx
At FNAL and any US location:	vtc fnal-lecture msb1.fnal.gov fnal-123
At KEK:	vtc fnal-lecture msb5.kek.jp fnal-xxx
At Tsukuba:	vtc fnal-lecture msb5.kek.jp fnal-123

This creates a virtual conference network with msb's at three nodes routing conference streams to end users and connected msb's. Note that both KEK and CERN "see" the meeting as the MBone routed multicast 224.4.222.222 with ttl=4. The msbs provide the tunneling of data between the laboratories. The regional/administrative geometry of the MBone is bypassed. This sparse virtual network is set up for the life of the conference and removed when the meeting ends. Many such conferences can occur simultaneously with arbitrary overlapping geometry.

3.4 Current and Future Development

In future releases of the bridge it is planned that the root msb of the meeting will be the only machine known by the attendee and that msb will tell vtc to re-home to another network closer msb. A WWW interface will also hide the need for attendees to know about different msb machines.

Planned security is composed of two components. The first is user authentication which is being addressed by public/private key authentication schemes. The public keys for the users will be stored in certification authorities (CAs) compatible with Esnet CAs. The information between vtc and msb will be encoded with keys obtained from a secure key exchange. The session itself is the second component of security. A session will define a key to be shared by all participating in the conference. This session key will use the encryption built into the tools (currently 56 bit DES).

Reports and statistics on meetings and attendance will be enhanced so meeting coordinators will be able monitor facility usage and attendance.

4 Conclusions

While MBone is very useful and works well for network-dense meetings, the restrictive domain formed by ttl scoping does not fit well for sparse meetings. This along with the fact that the MBone is not available at a number of HEP locations have driven the development of the msb.

Version 1 of the msb and vtc have been used extensively. In 1996 close to 500 conferences were held which consumed over 1000 hours of meeting time. The use continues unabated and should increase significantly with the migration to version 2.01. The use of the msb has made packet video conferencing a reality in the HEP community.

5 References

H. Schulzrinne, S. Casner, R. Frederick, V. Jacobson, RTP: A Transport Protocol for Real-Time Applications, RFC 1889, January 1996.

S. McCanne, V. Jacobson, vic: A Flexible Framework for Packet Video, ACM Multimedia Systems, pp 1- 11, November 1995.

V. Hardman, M. Handley, M. Sasse, A. Watson, Reliable Audio for Use over the Internet, unpublished paper.

S. Deering, D. Cheriton, A Multicast Extension to the Internet Protocol, RFC 966, December 1985.

S. Armstrong, A Freier, K. Marzullo, Multicast Transport Protocol, RFC 1301, February 1992.

R. Braudes, S. Zabele, Requirements for Multicast Protocols, RFC 1458, May 1993.

C. Semeria, T. Maufer, Introduction the IP Multicast, <draft-ietf-mboned-intro-multicast-00.txt>, Internet-Draft, January 1997.

unidentified authors, CGMP: Enabling IP Multicast in Layer 2 Switching Environments, http://www.cisco.com/warp/public/732/IP/cgmp-wp.htm.

D. Meyer, Administratively Scoped IP Multicast, <draft-ietf-mboned-admin-ip-space-01.txt>,
Internet- Draft, December 1996.

J. Hawkenson, Multicast Pruning a Necessity, <draft-ietf-mboned-pruning-01.txt>, Internet-Draft, September 1996.

S. Deering, D. Estrin, D. Farinacci, M. Handley, A. Helmy, V. Jacobson, C. Liu, P. Sharma, D. Thaler, L. Wei, Protocol Independent Multicast-Sparse Mode (PIM_SM): Motivation and Architecture, <draft- ietf-idmr-arch-04.txt>, Internet-Draft, October 1996.

S. Deering, D. Estrin, D. Farinacci, M. Handley, A. Helmy, V. Jacobson, C. Liu, P. Sharma, D. Thaler, L. Wei, Protocol Independent Multicast-Sparse Mode (PIM_SM): Protocol Specification, <draft-ietf- idmr-PIM-SM-spec-09.txt>, Internet-Draft October 1996.

D. Farinacci, Y. Rekhter, Multicast Tag Binding and Distribution Using PIM, <draft-farinacci-multicast- tagsw-00.txt>, Internet-Draft, November 1996.