MBONE, the Multicast BackbONE

Mike Macedonia and Don Brutzman
Naval Postgraduate School

(draft article accepted for publication in *IEEE Computer*)

Introduction. The joy of science is in the discovery. It was a year ago when we heard that the JASON Project, an underwater exploration educational program supported by Woods Hole Oceanographic Institution, was showing live video over the Internet from an underwater robot in waters off Baja, Mexico. Our group here at the Naval Postgraduate School (NPS) worked furiously to figure out how to receive that video signal. We labored for several days to gather the right equipment, contact the appropriate network managers, and get hardware permissions from local bureaucrats, only to find that an antenna uplink on the JASON support ship had flooded a few hours before we became operational. Despite this disappointment we remained enthusiastic because we had discovered how to use the Internet's most unique network, MBONE.

MBONE stands for Multicast BackbONE, a virtual network that has been in existence since early 1992. Named by Steve Casner of Information Sciences Institute (ISI), the network originated from an effort to multicast audio and video from the Internet Engineering Task Force (IETF) meetings. MBONE today is used by hundreds of researchers for developing protocols and applications for group communication. Multicast is used because it provides one-to-many and many-to-many network delivery services for applications such as videoconferencing and audio that need to communicate with several other hosts simultaneously.

This article describes the network concepts underlying MBONE, the importance of bandwidth considerations, various application tools and MBONE events. Sidebars describe interesting MBONE uses. We also provide guidance on how to connect your Internet site to the MBONE.

Multicast networks. Multicasting has existed for several years on local area networks such as Ethernet and Fiber Data Distributed Interface (FDDI). However, with Internet Protocol (IP) multicast addressing at the network layer group communication can be established across the Internet. IP multicast addressing is an Internet standard (Request For Comment RFC-1112) developed by Steve Deering of Xerox Palo Alto Research Center (PARC) and is supported by numerous workstation vendors including Sun, Silicon Graphics Inc., Digital Equipment Corporation and Hewlett Packard. Categorized officially as an IP Class D address, an IP multicast address is mapped to the underlying hardware multicast services of a local area network. Two things make multicasting feasible on a world-wide scale: the installation of high bandwidth Internet backbone connections, and the widespread availability of workstations with adequate processing power and built-in audio capability.

The reason that MBONE is a virtual network is that it shares the same physical media as the Internet. It uses a network of routers that can support multicast (mrouters). These mrouters are either upgraded commercial routers, or dedicated workstations running with modified kernels in parallel with standard routers. MBONE is also augmented with "tunnels". Tunneling is a scheme to forward multicast packets among the islands of MBONE subnets through Internet IP routers which (typically) do not support IP multicast. This is done by encapsulating the multicast packets inside regular IP packets. As installed commercial hardware is upgraded to support multicast traffic, this mixed system of specially-dedicated mrouters and tunnels will no longer be necessary. We expect that most commercial routers will support multicast in the near future, eliminating the inefficiencies and management headaches of duplicate routers and tunnels.

Bandwidth constraints. The key to understanding the constraints of MBONE is thinking about bandwidth. The reason why a multicast stream is bandwidth-efficient is that one packet can touch all workstations on a network. Thus a 128 Kbps (kilobits per second) video stream (typically 1-4 frames/second) uses the same bandwidth whether it is received by one workstation or twenty. That is good. However that same multicast packet is ordinarily prevented from crossing network boundaries such as routers. The reasons for this current restriction are religious and obvious: if a multicast stream which can touch every workstation could jump from network to network without controls, then the entire Internet would quickly become saturated by such streams. That would be very bad! Therefore some controls are necessary.

MBONE has two means to control multicast packet distribution across the Internet. The first is to limit the lifetime of multicast packets. The second (now in testing) is to use sophisticated pruning algorithms to adaptively restrict multicast transmission. Responsible daily use of the MBONE network consists merely of making sure you don't overload your local or regional bandwidth capacity.

MBONE protocol developers are currently experimenting with automatically pruning and grafting subtrees, but for the most part MBONE uses thresholds to truncate broadcasts to the leaf routers. The truncation is based on the setting for the time-to-live (*ttl*) field in a packet which is decremented each time the packet passes though an mrouter. A *ttl* value of 16 would limit multicast to a campus, as opposed to values of 128 or 255 which might send a multicast stream to every subnet on the MBONE (currently about thirteen countries). *ttl* is sometimes decremented by large values under a global thresholding scheme provided to limit multicasts to sites and regions if desired.

These issues can have a major impact on network performance. For example, a default video stream consumes about 128 Kbps of bandwidth, which is almost 10 percent of a T1 line (a common site-to-site link on the Internet). Several simultaneous high-bandwidth sessions might easily saturate network links and routers. This problem is compounded by the fact that general purpose workstation routers typically used by the MBONE are normally not as fast or as robust as the dedicated hardware routers used in most of the Internet.

Networking details. When a host on an MBONE-equipped subnet establishes or joins a common shared session, it announces that event via the Internet Group Management Protocol (IGMP). The multicast router on the subnet forwards that announcement to the other mrouters in the network. Groups are disestablished when everyone leaves, freeing up the IP multicast address for reuse. The routers occasionally poll hosts on the subnets to determine if any are still group members. If there is no reply by a host, the router stops advertising that host's group membership to the other multicast routers.

The routing protocols for MBONE are still immature and their ongoing design is a central part of this network experiment. Most MBONE routers employ the Distance Vector Multicast Routing Protocol (DVMRP), which is commonly considered inadequate for rapidly changing network topologies because routing information propagates too slowly. A Multicast extension to the Open Shortest Path (MOSPF) link-state protocol has been proposed by the MOSPF Working Group that addresses this issue, using an algorithm developed by Steve Deering. With both protocols, mrouters must dynamically compute a source tree for each participant in a multicast group. MBONE is currently small enough that this technique is not a problem. However, some researchers speculate that for a larger network with frequently changing group memberships these routing techniques will be computationally inefficient. Research efforts on these issues are ongoing, since every bottleneck conquered reveals a new bottleneck.

Topology and event scheduling. The MBONE topology and the scheduling of multicast sessions must be actively managed by the MBONE community to minimize congestion. By the beginning of 1994, 750 subnets worldwide were already connected to MBONE. Topology changes for new nodes are added by consensus: a new site announces itself to the MBONE mail list, and the nearest potential providers decide who can establish the most logical connection path to minimize regional Internet loading.

Scheduling MBONE events is handled similarly. Special programs are announced in advance on an MBONE event mail list (*rem-conf@es.net* for messages, *rem-conf-request@es.net* for subscription requests). Advance announcements usually prevent overloaded scheduling of Internet-wide events and also alert potential participants. Cooperation is key. Many people are surprised to learn that no single person or entity is "in charge" of either local topology changes or event scheduling.

Protocols. The magic of MBONE is that teleconferencing can be done in the hostile world of the Internet where variable packet delivery delays and limited bandwidth play havoc with applications that require some real-time guarantees. Limited experiments demonstrated the feasibility of audio over the ARPANet as early as 1973. However it is worth noting that only a few years ago putting video across the Internet was considered impossible. Development of effective multicast protocols disproved that widespread opinion. In this respect MBONE is like the proverbial talking dog: it's not so much what the dog has to say, it's more that the dog can talk at all that is amazing.

The key network concepts that make MBONE possible are IP multicast and real-time stream delivery via adaptive receivers. For example, many MBONE applications are using the draft Real Time Protocol (RTP) on top of User Datagram Protocol (UDP) and Internet Protocol (IP) in addition to the multicast protocols. RTP is being developed by the Audio-Video Transport Working Group within the IETF. RTP provides timing and sequencing services, permitting the application to adapt and smooth out network-induced latencies and errors. Related real-time delivery schemes are also being evaluated. The end result is that even with a time-critical application such as an audio tool, participants normally perceive conversations as if they are in real-time. This is because there is actually a small buffering delay to synchronize and resequence the arriving voice packets. Protocol development continues. Although operation is usually acceptable in practice, many aspects of MBONE are still considered experimental.

Data compression. Other aspects of this research include the related needs to compress a variety of media and optionally provide privacy through encryption. Several techniques to reduce bandwidth include Joint Photographic Experts Group (JPEG) compression, wavelet-based encoding and the ISO standard H.261 for video. Visually this translates to "velocity compression:" rapidly changing screen blocks are updated much more frequently than slowly changing blocks. Encodings for audio include Pulse Coded Modulation (PCM) and Group Speciale Mobile (GSM, name of the standardization group for the European digital cellular telephony standard). Outside of the concerns for real-time delivery, audio is a difficult media for both MBONE and teleconferencing in general. This is because of the need to balance signal levels for all parties, who may have different audio processing hardware (e.g. different microphones and amplifiers). Audio also generates lots of relatively small packets, which are the bane of network routers.

Application tools. Besides basic networking technology, MBONE researchers are developing new applications that typify many of the goals associated with an "information superhighway." MBONE session availability is dynamically announced using a tool called *sd* (session directory) which displays active multicast groups. *sd* is also used for launching multicast applications and for automatically selecting an unused address for any new groups. *sd* was developed by Steve McCanne and Van Jacobson of Lawrence Berkeley Laboratory, University of California Berkeley. Video, audio, and a shared drawing whiteboard are the principal MBONE applications, provided by software packages called *nv* (net video), *vat* (visual audio tool) and *wb* (whiteboard). The principal authors of these tools are Ron Frederick of Xerox Palo Alto Research Center (PARC) for *nv*, and Steve McCanne and Van Jacobson for *vat* and *wb*. Each of these programs is available in executable form without charge from various anonymous ftp sites on the Internet. Working versions are available for Sun, SGI, DEC and HP architectures with ports in progress for Macintosh. No DOS, OS-2, Amiga or Windows versions are currently available although ported tools can be found for 386 boxes running the (free) 386BSD Unix. Pointers to all public application tools are included in the Frequently Asked Questions (FAQ). Mirror ftp sites are available overseas.

sd: session directory of MBONE events

Additional tools are also available or under development. Winston Dang of the University of Hawaii has created *imm* (Image Multicaster Client), a low-bandwidth image server. It is typically used to provide live images of planet Earth from various geostationary satellites at half-hour intervals in either visible or infrared (IR) spectra. Henning Schulzrinne of AT&T/Bell Laboratories developed *nevot*, a network voice terminal providing multiple party conferences with a choice of transport protocols. Eve Schooler of University of Southern California (USC)/Information Sciences Institute (ISI) is part of a team developing *mmcc*, a session orchestration tool and multimedia conference control program. Article author Mike Macedonia has created a multicast version of *NPSNET*, which is a 3D distributed virtual environment that uses the IEEE Distributed Interactive Simulation (DIS) application protocol. Other researchers are experimenting with using graphics workstation windows as image drivers. Kurt Lidl of UUNET Technologies, Falls Church Virginia is working on a network news distribution application which uses multicast to reduce overall Internet loading and speed news delivery. (The target goal is 120 ms total propagation coast to coast, which is pretty amazing since light takes about 16 ms to make that trip.)

Events. Many of the most exciting events on the Internet appear first on MBONE. Perhaps the most popular is *NASA Select*, the NASA in-house cable channel broadcast during space shuttle missions. Be warned that this is a work stopper! It is hard to describe the excitement of seeing one astronaut positioning another astronaut by the boots to repair a satellite, all live on your desktop from 150 miles above the earth. Conferences on supercomputing, Internet Engineering Task Force, scientific visualization and many other

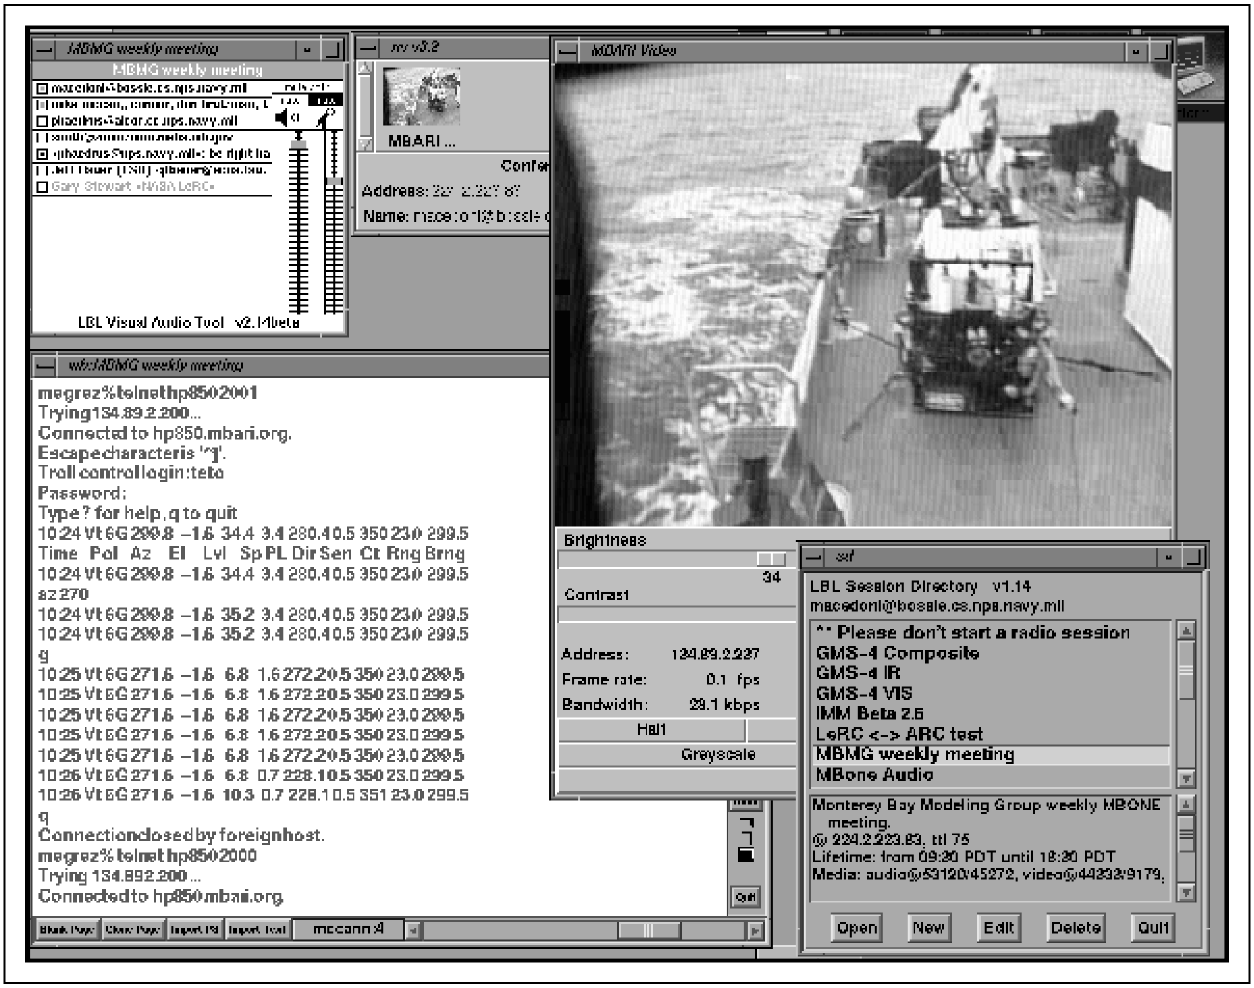

MBONE session with MBARI showing application tools *nv* network video, *vat* visual audio tool, *wb* whiteboard and *sd* session directory

topics have appeared, often accompanied by directions on how to download PostScript copies of presented papers and slides from anonymous ftp sites. *Radio Free VAT* is a community radio station whose DJ's sign up for air time via an automated server (vat-radio-request@elxr.jpl.nasa.gov). Xerox PARC occasionally broadcasts lectures by distinguished speakers. Internet Talk Radio (Carl Malamud, info@radio.com) has presented talks by Vice President Al Gore, Larry King and others. Another new area is remote learning, which can use the MBONE to bring expertise over long distances and multiply training benefits. Finally, default MBONE audio and video channels are provided for experimentation by new users and advice from more experienced users.

Groupwork on groupware. The MBONE community is active and open. Work on tools, protocols, standards, applications and events is very much a cooperative and international effort. Feedback and suggestions are often relayed to the entire MBONE

MBONE and Distance Learning at NPS
Mike McCann, NPS Visualization Lab

In March 1993 the W.R. Church Computer Center at NPS dedicated a Sun SPARCstation 2 to act as a multicast backbone (MBONE) router for the NPS campus and the Monterey Bay research community. This router and an IP-encapsulated tunnel from Stanford University provides the campus backbone with real-time audio, video, and other MBONE data feeds.

The MBONE is an excellent tool for those doing research in networks and video teleconferencing technology. Though it is not generally thought of as "ready for prime-time" (audio dropouts may be frequent and video, at best, is only 3 frames per second over the Internet) we successfully used it to provide training in Cray Fortran optimization from the National Center for Atmospheric Research (NCAR) in Boulder, Colorado. Five people who would not have been able to afford travel to Boulder remotely "attended" this course at the NPS Computer Center's Visualization Laboratory. We had 2-way audio and video between the classroom at NCAR and the Lab at NPS. We could ask questions of the instructor and had local hard copies of all the overheads he used. Advanced preparation, good audio, and a camera operator in the classroom gave a real feeling of presence. One attendee said "it was just like being there."

Paul Hyder of NCAR was instrumental in helping set up a direct "backup" tunnel between NPS and NCAR where the slowest link is the T-1 line between NPS and Stanford. In the 3-day period of the course there was only one 30 minute period of broken-up audio. We later determined that this interruption was caused by excess congestion on NCAR's ethernet Local Area Network.

For much of 1993 the Visualization Lab loaned a Sun SPARCstation 10 to the Monterey Bay Aquarium Research Institute (MBARI) for testing and incorporation into the live audio/video link to the research vessel *Point Lobos* and the remote operated vehicle *Ventana* that daily explore the Monterey submarine canyon. Local researchers in oceanography, virtual reality, and autonomous underwater vehicles continue to take advantage of the collaboration opportunities enabled by this technology. We should all look forward to videoconferencing to a classroom or a colleague half way around the world directly from our desktop workstations.

mailing list. (As an example, this article was reviewed by that group.) Cooperation is essential due to the limited bandwidth of many networks, in particular transoceanic links. So far no hierarchical scheme has been necessary for resolving potentially contentious issues such as topology changes or event scheduling. Interestingly, distributed problem solving and decision making has worked on a human level just as successfully as on the network protocol level. Hopefully this decentralized approach will continue to be successful even in the face of rapid addition of new users.

Cost of admission. The first thing you need to get on MBONE is the willingness to study and learn how to use these new and fast-moving tools. The second thing you need is bandwidth. Lastly you need some hardware, but the cost of equipment is often relatively low. Here at NPS we run MBONE tools on workstations connected via Ethernet (10 Mbps). Off-campus links are via T1 lines (1.5 Mbps). We have found that bandwidth capacities lower than T1 are generally unsuitable for MBONE video, though it is worth noting that some users (sometimes entire countries!) on specially-configured networks have managed to make the tools work at 56 and 64 Kbps.

Remote Science over the MBONE during the 1993 JASON Project

Andy Maffei, Deep Submergence Laboratory
Woods Hole Oceanographic Institution, Woods Hole Massachusetts 02543

JASON/MEDEA is a remotely-operated, dual-vehicle system developed by the Woods Hole Oceanographic Institution (WHOI) for underwater science and exploration. The JASON Foundation for Education uses this system as part of an annual "JASON Project" expedition, which actively involved more than 600,000 K-12 students during the live satellite transmissions in 1993. While prime broadcast hours of the expedition focused on the JASON Project's educational mission, an intense science program was conducted on a concurrent round-the-clock basis.

A 56 kbps data circuit was provided by the EDS Corporation and the University of Texas at Dallas. MorningStar PPP software, running on a Sun Workstation, established the Internet connection with research vessel *Laney Chouest* from which the vehicles were deployed. This Internet connection gave a transparent link with the Multicast-IP-based MBONE. Although we experimented with multicast videoconferencing tools such as *nv* and *vat*, our primary interest in using the MBONE was to transmit experimental data and to support shore-based models that depicted the positions and movement of the *Laney Chouest* and the two JASON/MEDEA vehicles. This technology was used by several investigators collaborating on JASON science projects at different locations throughout the country.

Both 2-dimensional and 3-dimensional models were developed at WHOI's Deep Submergence Laboratory for use on Sun and SGI workstations. Software packages to access these models along with real-time MBONE data were made available to anyone on the MBONE who wanted to try them out. As sonar surveys progressed during the expedition, data were transmitted back to shore and the detailed models were updated then distributed over the Internet.

A workstation on board the *Laney Chouest* generated multicast packets containing navigation and attitude information for the three vehicles. These packets were distributed in real time over the MBONE so users running the modeling software could watch a graphic display of JASON prowling real seafloor features as scientists investigated seamounts, thermal plumes, and the area's unique ecology. In addition to vehicle information, experimental data variables such as temperature were also multicast on the MBONE. Scientists and other interested users could write programs to read these experimental values, watch the models evolve, and get immediate feedback on progress being made during different experiments.

By the accounts of participating researchers, MBONE use enhanced the science carried out during the cruise. However, since we spent so much time in support of specific experiments, we were unable to spend much time helping other interested MBONE users get the models up and running at their own sites. As this was the first time we had used multicast IP during an experiment we learned a great deal. Such unique experiments demonstrate the value of other science-based tools in addition to more generic videoconferencing applications. Next time we hope to involve even more people on the MBONE in our project.

Given adequate network bandwidth, you now need a designated MBONE network administrator. It typically takes one to three weeks for a network-knowledgeable person working part-time to establish MBONE at a new site. Setup is not for the faint of heart, but all of the tools are documented and help is available from the MBONE list. Read the FAQ a few times and ensure that software tools and multicast-compatible kernels are available for your target workstations. Subscribe to the mail list in advance so that you will be able to ask questions and receive answers. A sidebar shows the various worldwide MBONE list subscription request addresses. After subscribing, read the FAQ again.

Mail addresses for requesting addition to MBONE mail list		
mbone-eu:	mbone-eu-request@sics.se	Europe
mbone-jp:	mbone-jp-request@wide.ad.jp	Japan
mbone-korea	mbone-korea-request@mani.kaist.ac.kr	Korea
mbone-na:	mbone-na-request@isi.edu	North America
mbone-nz:	mbone-nz-request@waikato.ac.nz	New Zealand
mbone-oz:	mbone-oz-request@internode.com.au	Australia
mbone-sg:	mbone-sg-request@lincoln.technet.sg	Singapore
mbone:	mbone-request@isi.edu	other

Note that these tools can also work in isolation between workstations on a single local area network without any mrouter. We recommend that you test the application tools locally in advance (before going through the dedicated mrouter effort) to see if they are compatible with your system and match your expectations.

To receive multicast packets on your local area network (LAN), you will need to configure a multicast router (mrouter). This can be either a single workstation on a LAN, or a host dedicated as a parallel mrouter. A non-dedicated single workstation can receive and pass multicast to its LAN neighbors, but this arrangement has the disadvantage of placing double MBONE traffic on that LAN. A more practical approach is to dedicate an old unused workstation as an mrouter and equip it with two Ethernet cards. The two network cards are needed for this mrouter to act independently and in parallel with your standard IP router. (Both approaches are used at NPS.) After deciding on your mrouter configuration, obtain and load the application software tools. You are now ready to put multicast on your local area network.

Once you are connected, pass along any lessons learned to the tool authors or the MBONE list. Later show your overall network site administrator something spectacular on MBONE (such as a live spacewalk) and then make sure that your site is programming funds to increase your network bandwidth. Demands on network bandwidth are significant and getting more critical. You might consider Tengdin's First (and Only) Law of Telecommunications: "The jump from zero to whatever baud rate is the most important jump you can make. After that everyone always wants to go straight to the speed of light."

Caveats aplenty. Some problems still exist and a lot of work is in progress. The audio interface takes coaching and practice. Leaving your microphone on by mistake may disrupt a session since typically only one person can be understood at a time. You will need a video capture board in your workstation to transmit video, but no special hardware is needed to receive video. One-to-four frames/second video seems pretty slow (standard video is 30 frames/second), but in practice it is surprisingly effective when combined with phone-quality voice. A big danger: one user blasting a high-bandwidth video signal (greater than

128 Kbps) can cause severe and widespread network problems. Controls on access to tools are rudimentary and security is minimal; for example, a local user might figure out how to listen through your workstation mike (unless you unplug it). Audio broadcast preparations are often just as involved as video broadcast preparations. Network monitoring tools are not yet convenient to use. There is no guaranteed delivery: lost packets stay lost. Internet bandwidth is still inadequate for MBONE in many countries. On one occasion an unusual topology change at our school caused a feedback loop that overrode the NASA Select audio track. Although plenty of people were willing to point out the symptoms of our error (!) it was not possible for the rest of the network to cut off the offending workstation cleanly. More situations will undoubtedly occur as the MBONE developers and users learn more. Surprises usually trigger a flurry of discussion and a corresponding improvement in the tools.

Expect to spend some time if you want to be an MBONE user. It is time consuming because learning and fixing are involved, and also because it is a lot of fun!

The future. It is not every day that someone says to you "Here is a multimedia television station that you can use to broadcast from your desktop to the world." These are powerful concepts and powerful tools that tremendously extend our ability to communicate and collaborate. They have already changed the way people work and interact on the net.

See you later!

Further reading

1. Casner, Steve. "Frequently Asked Questions (FAQ) on the Multicast Backbone," 6 May 1993, **file://venera.isi.edu/mbone/faq.txt**

2. Schulzrinne, Henning and Casner, Steve. "RTP: A Transport Protocol for Real-Time Applications," IETF Draft, 20 October 1993, **file://nic.ddn.mil/internet-drafts/draft-ietf-avt-rtp-04.ps**

3. Casner, Stephen and Deering, Stephen. "First IETF Internet Audiocast," ACM SIGCOMM Computer Communication Review, July 1992, pp. 92-97; also as **file://venera.isi.edu/pub/ietf-audiocast.article.ps**

4. Comer, Douglas E. *Internetworking with TCP/IP*, volume I, Prentice-Hall, New Jersey, 1991.

5. Deering, Stephen. "Host Extensions for IP Multicasting", RFC 1112, August 1989, **file://nic.ddn.mil/rfc/rfc1112.txt**

6. Deering, Stephen. "MBONE-The Multicast Backbone," CERFnet Seminar, 3 March 1993, **file://parcftp.xerox.com/pub/net-research/cerfnet-seminar-slides.ps**

7. Moy, John. "Multicast Extensions to OSPF," IETF Draft, July 1993, **file://nic.ddn.mil/internet-drafts/draft-ietf-mospf-multicast-o4.txt**

8. Perlman, Radia. *Interconnections: Bridges and Routers*, Addison-Wesley, New York, 1993, p. 258.

9. Curtis, Pavel, MBONE map, available via anonymous ftp from **file://parcftp.xerox.com/pub/net-research/mbone-map-big.ps**

10. Baker, Steve. "Multicasting for Sound and Video," *Unix Review*, February 1994, pp. 23-29.

11. *IEEE Standard for Information Technology - Protocols for Distributed Interactive Simulation (DIS) Applications*, version 2.0 third draft, University of Central Florida, Institute for Simulation and Training, 12424 Research Parkway, Suite 300, Orlando Florida 32826, May 28, 1993.

12. Pratt, David R., Zyda, Michael J., Falby, John and Kelleher, Kristen. *NPSNET Annual Lab Review*, December 1993, available at **file://taurus.cs.nps.navy.mil/pub/npsnet/NPSNET.annual.report.ps**

The authors thank the originators of the MBONE tools and dozens of MBONE users who provided essential contributions to this article.

This article is available electronically via anonymous ftp to **taurus.cs.nps.navy.mil** in subdirectory **pub/mbmg** as **mbone.hottopic.ps**; in Uniform Resource Locator (URL) format used above, this anonymous ftp request corresponds to **file://taurus.cs.nps.navy.mil/pub/mbmg/mbone.hottopic.ps**

Major Michael R. Macedonia USA is a Ph.D. student. Lieutenant Commander Donald P. Brutzman USN is an instructor and Ph.D. candidate. Both can be reached at Computer Science Department, Naval Postgraduate School, Monterey California USA 93943-5000; e-mail addresses are macedonia@cs.nps.navy.mil and brutzman@nps.navy.mil.