NEW! MEW! MULTICASTING TOMORROW'S INTERNET

Learn About Cutting-Edge Internet Multimedia Technology & Tools!

Take Advantage of the MBONE for Live Audio over the Internet, Videoconferencing, and Other Real-Time Events!

Get a Handle on the Technology That Takes Your Research, Presentations, and Meetings to the Extreme

KEVIN SAVETZ NEIL RANDALL YVES LEPAGE





MBONE: Multicasting Tomorrow's Internet Published by IDG Books Worldwide, Inc.

An International Data Group Company 919 E. Hillsdale Blvd. Suite 400 Foster City, CA 94404

Text and art copyright ©1996 by IDG Books Worldwide, Inc. All rights reserved. No part of this book, including interior design, cover design, and icons, may be reproduced or transmitted in any form, by any means (electronic, photocopying, recording, or otherwise) without the prior written permission of the publisher.

Library of Congress Catalog Card No.: 95-81949

ISBN: 1-56884-723-8

Printed in the United States of America

1B/ST/QS/ZW/IN

Distributed by Macmillan Canada for Canada; by Computer and Technical Books for the Caribbean Basin; by Contemporanea de Ediciones for Venezuela; by Distribuidora Cuspide for Argentina; by CITEC for Brazil; by Ediciones ZETA S.C.R. Ltda. for Peru; by Editorial Limusa SA for Mexico; by Transworld Publishers Limited in the United Kingdom and Europe; by Al-Maiman Publishers & Distributors for Saudi Arabia; by Simron Pty. Ltd. for South Africa; by IDG Communications (HK) Ltd. for Hong Kong; by Toppan Company Ltd. for Japan; by Addison Wesley Publishing Company for Korea; by Longman Singapore Publishers Ltd. for Singapore, Malaysia, Thailand, and Indonesia; by Unalis Corporation for Taiwan; by WS Computer Publishing Company, Inc. for the Philippines; by WoodsLane Pty. Ltd. for Australia; by WoodsLane Enterprises Ltd. for New Zealand.

For general information on IDG Books Worldwide's books in the U.S., please call our Consumer Customer Service department at 800-762-2974. For reseller information, including discounts and premium sales, please call our Reseller Customer Service department at 800-434-3422.

For information on where to purchase IDG Books Worldwide's books outside the U.S., contact IDG Books Worldwide at 415-655-3021 or fax 415-655-3295.

For information on translations, contact Marc Jeffrey Mikulich, Director, Foreign & Subsidiary Rights, at IDG Books Worldwide, 415-655-3018 or fax 415-655-3295.

For sales inquiries and special prices for bulk quantities, write to the address above or call IDG Books Worldwide at 415-655-3200.

For information on using IDG Books Worldwide's books in the classroom, or ordering examination copies, contact the Education Department at 800-434-2086 or fax 817-251-8174.

For authorization to photocopy items for corporate, personal, or educational use, please contact Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, or fax 508-750-4470.

Limit of Liability/Disclaimer of Warranty: Author and Publisher have used their best efforts in preparing this book. IDG Books Worldwide, Inc., and Author make no representation or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for any particular purpose and shall in no event be liable for any loss of profit or any other commercial damage, including but not limited to special, incidental, consequential, or other damages.

Trademarks: All brand names and product names used in this book are trademarks, registered trademarks, or trade names of their respective holders. IDG Books Worldwide is not associated with any product or vendor mentioned in this book.



is a trademark under exclusive license to IDG Books Worldwide, Inc., from International Data Group, Inc.

СНА

Today's Technology

Ithough some great new multimedia tools and toys will certainly be heading down the pike in the coming months and years, an impressive batch of programs that work multimedia wonders on the Internet is available today — now — this minute! The current lineup of audio and video tools works for ordinary users who dial into the Internet from home, using plain old phone lines and 14.4 Kbps modems.

This chapter examines some of today's multimedia programs. It isn't meant to be a complete roundup of what's available today, just a sampling of what's out there and what's in store. Multimedia tools for the Internet are becoming more common. Last year, you could count on one hand the multimedia Internet programs for Macs and PCs, but things are speeding up. In the first half of 1995, the number of these tools easily doubled, and the number will likely double again before the middle of 1996. It's a great time to be an Internaut, even one with a slow modem.

The Internet has been accessible in one form or another via dialup links for many years, so why are multimedia tools for PCs just beginning to blossom? It may be that software developers are just beginning to see the light of multimedia. A bigger reason, however, is that the SLIP and PPP connections that are necessary to use these tools have become commonplace enough and (thanks to products such as Internet in a Box and Internet Chameleon) easy enough to use that the user base is large enough to warrant creating these tools.

Conferencing Tools

The Internet has always been a superb tool for conferencing, but not, unfortunately, for *real-time* conferencing. Electronic mail enables you to conference with other users, and conferencing is the very basis of Usenet newsgroups; but you can't have a real-time discussion by using e-mail or a Usenet newsgroup because the point-counter-point exchanges are separated by minutes, hours, and, sometimes, even days.

Of course, real-time communication over the Net has been available as well. Early *chat* programs, such as UNIX's *Talk* (see Figure 2-1), enable you to engage in real-time typing with another Internaut, while Internet Relay Chat, or IRC (see Figure 2-2), expands on this idea to include multiple typists all doing their thing at the same time. As anyone who has participated in *Talk* or IRC sessions can tell you, however, both experiences can be extremely disorienting. IRC sessions typically consist of reading lines of text and then trying to type something in fast enough to see it before it scrolls away. And Talk sessions, which seem like a good idea, quickly degenerate into a competition over who can type faster or better. The point is, typing was never intended as a real-time communication activity, and no amount of fancy programming can make it one.

northcoast.com 3	<u>ព</u>
[Connection established^G^G^G] HI there. Whats up?	4
Yeah, good stuff. I forwarded it to Mitch for perusal.	
OKwhere?	
L	
Not much. Did you receive the copy of the article i sent you? OK, fine. How 'bout lunch?	
<u>ئ</u>	

Figure 2-1: A Unix talk session.

That's why today's conferencing tools show such promise. Instead of typing, you spend your on-line time watching and listening, two activities users are far more familiar with in face-to-face or voice-to-voice communications. Significant hurdles remain, particularly in the area of access speed and bandwidth, but even now the potential for conferencing is clear.

18



Figure 2-2: An IRC session.

Audioconferencing tools

I am not a fan of long-distance telephone companies. Their simpering television ads and pathetic bickering just grate my soul. You've probably never heard of my longdistance company. It doesn't advertise on television, its salespeople don't call you up during dinner to ask you to switch to its service, and its prices are better than prices of the big three companies. Who needs 'em? There's an even better way to annoy the phone companies — I don't use the phone at all.

If the newest crop of Internet tools is any indication, fewer and fewer Internet users may be making traditional phone calls. The early months of 1995 saw the advent of a new type of application: live voice conferencing over the Internet. Armed with the right hardware, a common modem, and special software, you can chat over the Internet in your own voice. Although the technology is relatively new, audioconferencing software has the potential to change the way people use the Internet, just as the World Wide Web did. Dare I say it? One of these programs could be the Internet's next "killer app."

Imagine doing your regular on-line thing — maybe reading your e-mail or checking out some Web pages — when the computer signals an incoming call. You switch to your audioconferencing software, click Answer, and you're chatting in real time with your boss, your mom, or anyone else who has an Internet connection, anywhere in the world. This scenario is not a fantasy. A variety of programs exist today that can put that magic in your machine. All you need are a Mac or a PC that is running Windows or OS/2; a microphone and sound card; a SLIP or PPP Net connection; and a modem or a LAN. Amazingly, for most of these programs, a 14.4 Kbps modem can transmit intelligible voices, and 9600 bps is sometimes enough bandwidth.

9

I'm not suggesting that your PC and modem will replace Ma Bell anytime soon. The obvious drawback of on-line audio is that you can talk only to associates who have similarly-equipped hardware. Although millions of folks have Internet accounts, it will probably be many years before all of your friends and relatives are on the Net.

How does audioconferencing work?

Audioconferencing programs work by digitizing your speech as you talk and sending the digital data over the Internet. But there's a problem. A typical modem connection has limited bandwidth: 14.4 Kbps modems can send and receive a maximum of 1,800 bytes of noncompressible data each second. Telephone-quality speech needs 8,000 bytes per second of bandwidth. The two solutions to the problem are to get a faster modem or to compress the sound information before transmitting it. Most programs compress the audio. For example, one audio application, NetPhone, finesses the problem by compressing each second of audio into fewer than 1,300 bytes before transmitting it. At the receiving end, NetPhone uncompresses the audio and plays it out of the computer's speaker.

A variety of methods for encoding and compressing sound data are available, and as you might expect, the standards aren't necessarily very standard, yet. The quality of the audio that you can send and receive depends on the application that you're using, the speed of your computer, and the compression method that's used. In my tests, audio is understandable, albeit less clear than a phone call. Still, talking across the country or around the world for the cost of an Internet connection is kind of amazing, and easy on the wallet.

Audio quality and compression

Depending on the software that you use, talking over the Internet may be more like using CB radio than like using a telephone. In a phone conversation, you can interrupt or inject "uh-huhs," but many of today's programs allow only one person to speak at a time. This type of communication is called *half-duplex* conversation, and it means that you have to wait for the speaker to release the line before you can respond.

Even if your software provides full-duplex, real-time conversation, a short but sometimes noticeable gap occurs between when you speak and when the person to whom you're speaking hears your words. The modems themselves introduce about ¹/₄ second of "latency" in the sound transmission because of the buffering, error correcting, and data compression that they do. The result is a bit like talking over an intercontinental satellite telephone link. If the time lag bothers you, you may be able to reduce this effect by turning off your modem's compression bells and whistles.

The audio quality can vary, but generally, no matter what program you use, voice quality is acceptable. Surprising to most people, if you have a good Internet connection, then the sound quality is comparable to a regular phone call. However the sound quality varies dramatically between applications. In nearly all cases, the limiting factor

21

will not be the speed and capacity of the Internet but will depend on the local work at each end. If both parties have fast network connections, some applications provide sound quality significantly better than a regular telephone call because they use 16-bit signals rather than 8-bit signals in regular phone lines.

Audio compression introduces distortions into the sound that are made worse by the lousy speakers that are built into most computers. If the speaker in your computer doesn't sound good enough, try using headphones. They eliminate the worst of the distortion by bringing the source of the sound closer to your ear. Compression causes audio to lose intensity in important frequencies, but the loss is less noticeable when the source of the sound is closer to your ear.

CVSD (continuous variable slope deltamodulation) is a compression method that isn't very CPU intensive, but it provides a respectable 8:1 compression ratio. CVSD-encoded voices are understandable but not terribly crisp. Another encoding type, *GSM* (which stands for *Global System for Mobile telecommunications*) requires faster computers on the sending and receiving ends, but the quality is quite a bit better than with CVSD. In fact, GSM at a sampling rate of 6 kHz (6,000 sound samples per second) is clearer than CVSD at 12 kHz (12,000 samples per second).

Other encoding methods include *raw* or *linear* encoding, which uses one byte per sample without compression, and *Intel DVI*, which is a common standard that does some compression (about 2:1).

The protocol that everyone is waiting for is *RTP (real time protocol)*. RTP is a standard that is being created by the people who created vat, which is the UNIX audioconferencing tool that started it all. So far, the developers of audioconferencing programs have been working in isolation. RTP is emerging as an elegant protocol that can provide better interoperability to bridge the gap between the many audio tools that don't currently talk to one another.

Maven

Maven (see Figure 2-3) was the first Internet audioconferencing tool for Macintosh. Maven is free software, but it is a bandwidth hog. Maven requires a minimum of 13 Kbps. Maven can talk to other Macs that are running Maven, as well as to the UNIX vat program.

Maven was the first real-time Internet audio program to work on an affordable home computer. According to Charley Kline, the programmer who started the Maven project, "Maven started out as a challenge to me when I heard at the Audio-Video Transport (AVT) working group at the Internet Engineering Task Force (IETF) meeting a couple years ago that the 'only serious platform' for multimedia network conferencing was a high-end UNIX workstation. I knew that the Macintosh had better audio processing ability than a Sun (true!) and set about to write a little 'toy' that would interoperate with the existing UNIX conferencing tools. Within a few weeks, I had something demonstrable."

Part I + Tomorrow's Technology Today



Figure 2-3: Maven.

According to Kline, Maven "sort of escaped to the Internet" when a friend passed the program around on the *Global Schoolhouse*, a network of K–12 schools that uses Internet videoconferencing tools. "Before I knew it, I had a phone call from the National Science Foundation asking if I would be interested in continuing work on Maven under a grant." The rest, as they say, is history.

Maven is a free program that you can get via FTP in k12.cnidr.org :/pub/Mac/.

For more information, join the Maven e-mail discussion list. To subscribe, send e-mail to listserv@cnidr.org with a message body of subscribe maven Your Name.

NetPhone

NetPhone, like Maven, gives an Internet-connected Macintosh the ability to do audioconferencing (see Figure 2-4). But NetPhone has the upper hand over Maven because it works over slower connections. It works reliably using only a 14.4 Kbps modem. At a sampling rate of 10,000 samples per second, the sound quality isn't wonderful, but the voices of my NetPhone compatriots are certainly understandable. NetPhone shines because it can handle various sound encoding methods, and sampling rates from 4,000 to 16,000 samples a second. All you need is a Mac, a microphone, and a SLIP or PPP connection. You can use NetPhone to talk to users of NetPhone and Maven, as well as to users of any program that uses the UNIX vat protocol. Sadly, Netphone is not available for Windows, so you're limited to talking to Mac and UNIX people. (A version for Windows may be available by the time you read these words.)

Note

22





Figure 2-4: NetPhone.

NetPhone requires a Mac LC or faster computer. It is a commercial program that costs \$75 per copy or \$125 for two copies. NetPhone certainly won't be the only program on the Mac to follow in Maven's footsteps, but it's a great program nonetheless, especially if you're not lucky enough to be connected to a high-speed network.

You can try a free demo version of NetPhone that limits calls to 90 seconds. For information, send e-mail to netphone-orders@emagic.com or visit NetPhone's Web page at http://www.emagic.com/.

Speek Freely

Speak Freely, available for Windows and UNIX systems, is free and one of the best applications available. It offers many advanced features such as voice mail, multi-casting, encryption, and usually offers excellent sound quality. Its GSM compression routine requires a high-end 486 or Pentium processor. It will work for 14.4 Kbps modems with GSM, but sound quality is degraded. The current version is not compatible with other software except for Speak Freely for UNIX. Complete source code is available.

For information, visit the Web page http://www.fourmilab.ch/netfone/
windows/speak_freely.html.

Note

Internet Phone

Internet Phone is a \$69 voice communications program for Windows 3.1 or higher. It requires a 33 MHz 486 or faster machine, Winsock 1.1, and a 14.4 Kbps modem over SLIP or PPP. The Internet Phone software uses a proprietary compression algorithm, so it won't talk to other programs. Internet Phone runs on top of Internet Relay Chat, providing you with a list of on-line users and topics of conversation; so should you be inclined, you can easily find perfect strangers to talk with.

In mid-1995, a new version of Internet Phone was released that added full-duplex capabilities. What full-duplexing means, from the standpoint of audioconferencing, is that both users can speak at the same time. Because telephone conversations obviously work this way, full-duplexing software is important if audioconferencing is to catch on; and Internet Phone (which has become extremely popular among Windows users) will almost certainly help it to do so. In order to use the full-duplex version, you need either an audio board that will support this technology or two separate audio boards, one for receiving sound and one for transmitting it. Given the pure hell of configuring peripherals for PCs, most of us will experience the joys of full-duplexed conversations only when audio boards that support full-duplexing become commonplace.

For more information, see the Vocaltec Web page at http://www.vocaltec.com/ or send e-mail to info@vocaltec.com. The software is available via FTP from ftp.vocaltec.com in the /pub directory.

Internet Global Phone

Unlike the other audioconferencing programs mentioned in this chapter, Internet Global Phone is not a clean, ready-to-use application. In fact, it is very much a work in progress. Internet Global Phone (IGP) is a technology demonstration project that provides a code platform for two-way, real-time voice exchange over the Internet. The details of the project are documented in an article in *Dr. Dobb's Journal* in December 1994. If you like hacking in Microsoft Visual C++, you too can hack on Internet Global Phone.

"The state of the software is currently an engine without a user interface. The voice digitalization, compression/decompression, and transmission mechanisms are built-in and functional, but the user interface is simplistic," according to Sing Li of microWonders, codeveloper of the program. "Users and programmers are encouraged to add their own user interface and functions to this base code; essentially constructing their own phone," he said.

Internet Global Phone is based on the GSM compression standard. To run IGP (after you compile it), you need an Internet connection, a Windows-supported PC sound-board, and a microphone. A 20 MHz 386 or faster machine is recommended for best performance, as is a 14.4 Kbps modem link.

Note

74.



Note

You can get the source code via FTP from ftp.cs.tu-berlin.

de:/pub/local/kbs/tubmik/gsm/ddj. You can find a newer, ready-to-run version on CICA or any mirror under the win3/demos directory. Get the file IGP8_102.ZIP for 8-bit sound cards, IG16_102.ZIP for 16-bit sound cards, and IGS_102.ZIP for the source code. The version numbers will certainly change by the time you get there. That's the price of progress.

More to come? Certainly

In the coming months, more companies will help Internauts to reach out and touch someone. The future will certainly bring more audio tools and features, limited only by imagination. Some audioconferencing tools already provide a sort of digital answering machine to answer your digital telephone when you're not online. Others provide encryption. What about the ability to have a conference call of three or more people? How long will it take for anonymous Internet phone sex services to be available? Will people really use Internet telephones for long-distance calling on a regular basis? Will the long-distance companies wither away, or will they revel in the increased bandwidth that they'll be able to sell to the Internet Service Providers?

For updated information on Internet audioconferencing software, get the FAQ "How can I use the Internet as a telephone?" from the World Wide Web at http://www.northcoast.com/~savetz/voice-faq.html.Or you can receive it via electronic mail by sending a message to voice-faq-request@northcoast.com with a subject of archive and a body of send voice-faq.

A videoconferencing tool

If you really want to push the limit of your modem, you might try playing with CU-SeeMe, a videoconferencing program. CU-SeeMe sends and receives video in real time, enabling you to see the people you're chatting with onscreen. It doesn't work well for two-way communication using anything less than 56 Kbps, but those of us without a fast link can still be voyeurs by watching well-connected netizens wave to each other in blocky, slowly-moving windows. OK, so it's not perfect, but it is an interesting start.

Currently, video images over CU-SeeMe are black and white, using a very low frame rate (1 - 2 frames per second is typical). As faster links to the Internet become more commonplace, performance and video quality will improve.

CU-SeeMe consists of a client program (available for the Mac and Microsoft Windows) and a server-like component called a *reflector*. CU-SeeMe's person-to-person connections for video phone calls provide real-time interactive video and voice communication over the Internet. When several users connect to a reflector rather than directly to another CU-SeeMe user, group videoconferencing is possible. Figure 2-5 shows the basics of reflector technology.



Figure 2-5: How a reflector works.

Reflectors simulate multicasting by sending audio and video signals to multiple Internet sites at the same time. A reflector acts as a mirror. CU-SeeMe users can connect to the reflector and bounce their AV signal off the reflector; other users can connect to the reflector and receive the signal or signals being reflected. This system works well as long as only a few users are sending video signals and only a few are receiving them. Unlike with multicasting, reflectors must send out a stream of IP packets for each user who wants to receive the signal; so running a reflector requires a tremendous amount of bandwidth. Similarly, although you may be able to receive one low-quality video signal with a 14.4 modem, a five-person videoconference requires five times the bandwidth, too much for most of us. Sending your own picture into the ether requires yet more bandwidth.

CU-SeeMe is available as free software and will work with just about any video capturing hardware. According to Tim Dorcey, a CU-SeeMe senior developer, "one of the central objectives in developing CU-SeeMe was to make it available to the widest possible range of users, under the belief that the value of a communication technology is in large part determined by the number of people with access to it."

When the time comes that regular users have enough bandwidth to send *and* receive video, you'll still need a video camera to beam your mug into Cyberspace. Luckily, cheap hardware is already available. A company named Connectix has a gadget called the QuickCam, a \$99 black-and-white snowball-size job that works with CU-SeeMe. Word has it that a color version will be available by the time you read this. QuickCams are available for Windows and Macintosh.

Another product, LogiTech's VideoMan, is even cooler (although a bit more pricey at \$399). It's a sound and full-color video capture card for Windows. Besides working with CU-SeeMe, the VideoMan can do CD-quality audio and 30 frames-per-second video, and it includes its own digital camera.

CU-SeeMe was developed at Cornell University, where it has grown from a simple experiment into the de facto standard for Internet videoconferencing. The Cu-SeeMe development team began work on desktop videoconferencing in 1993. In May 1995, Cornell selected White Pine Software to enhance CU-SeeMe and create a commercial version of the program.

The CU-SeeMe development team is preparing a new version with a slide projector and a plug-in interface to support externally developed function modules, such as onscreen chat windows.

CU-SeeMe has been used for more than simple talking-heads style videoconferencing. Other "broadcast" events with Cu-SeeMe have included speeches, lectures, and even a kitschy low-budget comedy flick. On June 3, 1995, a film called *Party Girl* made Internet history as the first full-length movie broadcast on the Internet with CU-SeeMe. The one-time Net showing of the black-and-white flick coincided with its premiere at the Seattle International Film Festival.

Information about the film and its Internet debut is available at http://www.polis. com/firstlook/party/default.html.Information about the Seattle International Film Festival is available at http://www.film.com.

Unlike with audioconferencing software, CU-SeeMe is the only choice available for Internet videoconferencing. When the masses have faster modems and higherbandwidth connections, other software options will certainly appear. For now, though, CU-SeeMe is in a class of its own and does a fine job.

For more information about CU-SeeMe, point your Web browser to http://www. wpine.com/cuseeme.html or subscribe to one of the CU-SeeMe mailing lists. CU-SeeMe-L is for discussion of the program. To subscribe, send e-mail to listserv@cornell.edu with a message body of subscribe cu-seeme-L <first name> <last name>. CU-SeeMe-Announce-L is a low-volume mailing list dedicated to announcements about the CU-SeeMe software. To subscribe, send e-mail to listserv@cornell.edu with a message body of subscribe cu-seemeannounce-L <first name> <last name>.

Note

Note

Note

Part I * Tomorrow's Technology Today

Note

28

To obtain CU-SeeMe for Windows, use anonymous FTP to cu-seeme.cornell.edu in the directory /pub/video/PC.CU-SeeMe WO.xxx. Download the ReadMe file from that directory in text mode and the CUSEEME.ZIP file in binary mode.

Note

To obtain CU-SeeMe for the Macintosh, use anonymous FTP to cuseeme.cornell. edu in the directory pub/video/Mac.CU-SeeMe0.xxx. Download the ReadMe file from that directory in text mode. Read it and obtain the other files you need.

Collaboration Tools

Today, the Internet can provide more than the equivalent of the telephone and videophone. Using so-called collaboration tools, two or more people separated by any distance (for example, a department, a task force, or just two authors working on a book) can share ideas, outlines, art, text, and other information in real time.

Collaboration tools often work with the metaphor of a *shared whiteboard*, an area on each user's computer screen that can be written to, erased, drawn upon, or scribbled upon. Imagine a meeting in a boardroom with a large chalkboard or one of those nifty whiteboards that you write on with colored pens. Collaboration tools put the whiteboard on each participant's screen, enabling all of the participants to create outlines, diagram ideas — and even play Pictionary — without schlepping to the same physical location.

Some collaboration tools provide public and private chat windows and shared text editing fields, and some even enable participants to vote on issues. Some collaboration tools provide a *WYSIWIS* (what-you-see-is-what-I-see) feature, which allows a moderator to run other applications on his computer and have the moderator's screen be mirrored on each of the participants' computers.

NCSA (the National Center for Supercomputing Applications, the same folks who brought you Mosaic) has created a program called *Collage* for the analysis and sharing of scientific data. Collage's functions include image display and analysis, color table editing, and spreadsheet display of floating-point numbers.

Note

Note

For more information on Collage, see http://www.ncsa.uiuc. edu/SDG/Software/Brochure/Overview/MacCollage.overview.html.

Another collaboration tool, Co-Motion Lite, focuses on group idea generation, annotation, evaluation, and issue resolution. Co-Motion Lite is commercial software. For more information, contact Bittco at bittco@ccinet.ab.ca or by phone at 403-922-5514 or fax at 403-922-2859.

Because most collaboration tools transmit text and line art rather than bandwidthintensive data, you can use these tools reliably with Internet connections of 9600 bps or even slower. Coupled with a traditional conference call for voice communication, collaboration tools can be a useful way to share information and ideas, even if you can't share the same office.

Radio

Because audio works better than video on modem connections, when Internauts try broadcasting on the Internet, the most-successful experiments involve audio broadcasting — Internet radio.

Note

As I write this section, my computer is blaring a song by Sunfish, a Canadian band that I've never heard of before. I downloaded the tune from Virtual Radio (http://www. microserve.net/vradio), an online warehouse of, well, bands that I've never heard of. The song, "Difference," is online in it's entirety. It's about 2MB, and it took several minutes to download, but it was definitely worth the wait. It has a beat; you can dance to it.

Yes indeed, there's music on the Net. Another Web site offers news and commentary from National Public Radio's "Morning Edition" show. Yet another site houses an impressive collection of music from California garage bands.

Sound on the Internet is in its infancy. Standards are still being hashed out, and no "killer apps" are available for audio (yet). But as the Internet takes its slow course toward being the source for 500 channels of multimedia superhighway mayhem, the ability to transmit sound is becoming essential. Today's experiments with audio on the Net are a great first step. They provide interesting content, and they work on almost any computer that can play audio. However, the technology still has hurdles to overcome. Most notably, the sound quality of most of these services is low, scratchy, and low fidelity. Despite the emerging press releases and corporate propaganda, don't believe that the current uses of sound on the Net are anything more than experiments.

Still, it's fun to be on the cutting edge, listening to songs from bands that I didn't know existed.

Today's audio on the Internet loosely falls into two categories:

- Traditional broadcasters (such as National Public Radio and the American Broadcasting Company), who are looking to the Internet for added exposure
- A new kind of Internet-only broadcaster (such as the Internet Underground Music Archive), who is looking to the Net for the only realistic chance at a widespread audience

Both categories are interesting for different reasons. The traditional broadcasters know very well that the Net threatens to grab the attention of their listeners (and is, therefore, a threat), and they're actively looking for ways to use the new medium to enhance the old. A wide range of radio stations now offer Web sites that advertise their programming, make programs available as downloadable files, and point to the home pages of local and regional musicians whose music they feature. For the Internet-only broadcasters, the challenge lies in getting people to tune in beyond the first try. They depend completely on listeners' bookmarking their sites and returning to them, and thus they have a much harder task than the traditional broadcasters. On the other hand, Internet-only radio is ultracool, and Internet-savvy bands are beginning to make their music available on the Net. Similarly, Internet-savvy listeners are tuning in, and as access becomes faster, broadcasting on the Internet could very well change the way that people experience radio.

Internet Talk Radio

Internet Talk Radio (ITR) was one of the first well-known experiments in "broadcasting" on the Internet. Actually, ITR doesn't do broadcasting — or even multicasting. It publishes audio files that you can download and then listen to on your computer. A typical half-hour radio show consumes a whopping 15MB of disk space.

Carl Malamud is the founder of ITR. "The idea for ITR came from my frustration with the trade press. I knew they weren't providing the information I wanted, and [I] was looking for an alternative." He noted that the trade press focuses on marketing and reviews, leaving a gap for a general-interest, technically-oriented publication for Internet users. "I couldn't start a magazine, because it takes money to print and distribute a magazine," he said. Malamud turned to the Internet as a general-purpose distribution method.

Some Net users have criticized the talk radio concept as a grandiose waste of network bandwidth, given the fact that the same information in text format could fit into only a few kilobytes. "The reason that you get audio information from a \$3,000 (or \$30,000) computer," Malamud said, "is because ultimately this gives you a very new medium. We're not trying to replace radio, just as the trucks didn't replace the railroads and the telephone didn't replace the telegraph. There are things we can do that you can't do on a radio, like go interactive or add WAIS databases to support a program, or make an audio-on-demand server."

Internet Talk Radio began in 1993, and although it was innovative, it seems to have fallen by the wayside. It has been surpassed by newer technologies (some of which we examine later in this chapter), and Malamud has moved on to newer avenues of Internet multimedia.

Note

For more information on Internet Talk radio, visit http://www.town.hall. org/radio/index.html.

RealAudio

Fast forward to 1995, when a company called Progressive Networks created a program called the RealAudio player, software that deals with the biggest shortcoming of Internet Talk Radio — those huge time-sucking, hard-disk-eating audio files. Instead of making you download a complete audio file prior to listening, RealAudio uses audio

30

streaming: It plays the sound as it downloads. Its audio-on-demand technology reduces the wait to hear a program to just a few seconds. Because RealAudio uses strong audio compression, you can listen to audio by using only a 14.4 Kbps modem. So you don't need to wait to download 15MB to hear a radio program. With RealAudio, you can (theoretically) listen to "live" broadcasts, such as sports events and news as it breaks.

The RealAudio software (see Figure 2-6) runs on Windows, Macintosh, and UNIX. The RealAudio Player is free, and it works from within your Web Browser. When you click on a RealAudio file, the Player launches and begins to play. Progressive Networks is lining up traditional broadcasters, including National Public Radio and ABC, to provide material for RealAudio. Net notables including former MTV veejay Adam Curry. Radio Canada and C-SPAN are using RealAudio to publishing information, and certainly more broadcasters will be announced.

	RealAudio Play	er 🔤
► / II	III.,,	I.
Title:	Morning Edition - January 19, 199	6 = =
Author ;	National Public Radio]@[
Copyright	: 1996	T
Playing ne	twork stream	0.0/36:15.1

Figure 2-6: RealAudio player.

The RealAudio player presents the listener with a set of stop, play, rewind, and fastforward buttons. Unlike with "real" radio, you can stop the music if you need to step away from your desk, and then pick up where you left off when you return. Even better, you can index radio shows so that you can jump straight to the material that interests you. For example, if you want to hear Radio Canada's news story on why huge chunks of ice are breaking off Antarctica, you can go directly there, skipping over the program's introduction and never hearing the feature about how sex may control the cockroach population.

RealAudio's sound quality isn't wonderful — it's something akin to AM radio — but speech is certainly understandable. Unfortunately, whether you're connected to the Internet with a 14.4 Kbps modem or a T-1 line, RealAudio sounds the same. It's too bad that higher bandwidth doesn't bring better-quality sound. Then again, better sound is coming, according to Progressive Networks, although almost certainly not in the range of CD-quality audio files. RealAudio's true future appears to be in Internet-based talk radio; and because the Net spans the globe, this alone is reason to be excited.

Note

About five minutes before this book went to press, RealAudio 2.0 was released — the new version is said to support better-quality audio for users with faster connections.



Figure 2-7: The RealAudio Web site.

Note

For more information check out RealAudio's Web site: http://www.realaudio.com. (See Figure 2-7.)

National Public Radio

For years, National Pubic Radio's primary outlet for its news and features has been hundreds of public radio stations throughout the United States. Now, NPR is experimenting with using RealAudio to publish its material on the Internet. (See figure 2-8.)

Richard Dean, producer of New Media Services at National Public Radio, is excited about the benefits of publishing audio on the Internet. He notes that audio-on-demand has its benefits and drawbacks. "As a listener, it's a great thing. As a radio programmer, maybe it's not such a great thing — it disturbs the continuity of the program," he says. "Instead of pushing info at someone — that's what radio broadcasters do — you allow listeners to come get it from you."

"The biggest problem with radio is that your work is really ephemeral; it hits the ether, and it's gone. No one can hear it again, unless they go buy the tape or the transcript." Because a transcript can't capture the nuances of a news story, and audio-tapes take several days or weeks to arrive, the Internet gives NPR listeners the ability to hear — or hear again — a story at any time, when it's convenient. "This is where radio broadcasting — information providing — is going," Dean said.





Figure 2-8: National Public Radio on the Web.

"Here's my problem with the MBONE: Nobody's on it. I think MBONE is a really neat technology, but it doesn't reach who I want to reach; it reaches a bunch of guys with UNIX boxes. With NPR, our goal is to get our content to the end user. That end user is using a dial-up Internet account," he said. RealAudio technology may not be the ultimate in Internet multimedia, but for NPR, it's the best of the technologies available today. "I look at things two ways — one is what's now, what can I do today, and then, where is the technology going?"

As faster connections become available to the masses, "the MBONE will become something more important than it is right now. But I need to reach listeners now, not in years." He equates today's audio technology to the hand-crank phone: "It may not be the right solution in the long run, but in the short run it is." People don't want to wait. When the technology reaches the next level, you jump to the next level.

For more information about NPR and its Internet broadcasts, point your Web browser to http://www.npr.org/NPR.

Dean says that one of the drawbacks of RealAudio is that the sound quality is the same no matter the speed of your Internet connection. RealAudio sounds like an AM radio whether you're using a 14.4 Kbps or a T-1 connection to your desk, unlike with an emerging protocol, MPEG-2 streaming. In contrast, if you have the hardware and software necessary to handle MPEG-2 streaming and you have plenty of bandwidth to spare, you can hear audio with excellent quality. If you don't have the speed, you can increase the compression — dialing up less quality for less bandwidth.

Note

34. Part I + Tomorrow's Technology Today

Another issue for National Public Radio has to do not with technology but with copyright law. Currently, NPR needs to edit the musical "bumpers" from the Internet broadcasts. Although NPR can legally play songs (or portions of songs) on the air by virtue of paying ASCAP and BMI fees, the music organizations don't yet dole out online broadcast rights. It's still unclear how these organizations, who traditionally get paid based on the broadcast range of radio stations, will charge on-line broadcasters. Should a radio station pay royalties based on a potential on-line audience of 30 million? Or can tracking software be developed so that stations can tell exactly who is tuning in to e-broadcasts and when they are tuning into them?

Small college radio stations are experimenting by broadcasting their signals live on the Internet (Figure 2-9). What they're doing may be illegal (at least according to the music clearinghouses ASCAP and BMI), but these stations are pushing the envelope of what can be done and pushing us to contemplate these issues sooner rather than later.

Dean says this reflects a fundamental change in how radio broadcasters must think about their function. His advice to radio stations entering the Internet realm is to "Think of yourself as an information provider, not a radio station." To avoid obsolescence, today's broadcasters need to care about the information they're providing rather than about the medium itself. If the railroads realized that they were in the transportation business, they would be airlines today.



Figure 2-9: The WXYC Web site.

Like everyone else who is hawking wares on the infobahn, NPR needs to make money. NPR uses underwriting sponsorships and sells products (T-shirts, mugs, and the like) online.

Traditionally, NPR gets money from member stations, who in turn solicit funds from their listeners. If NPR and other broadcasters-turned-Internet-multimedia-info-pushers want to stay in the black, the rules of where the money comes from must shift. Today, NPR's on-line audio is free to the listener. "In several years, it might not be free to the end user," Dean said. "But there is no model yet. The price will have to be something incredibly low. The on-demand audience is skyrocketing. The money may not be there, but the audience is, and the money will follow."

The same technology that enables NPR to publish information on the Net enables anyone else — J. Random Hacker — to create an on-line cyberstation. Could these basement stations create competition for traditional broadcasters who are trying to find a home on the Internet? Certainly. "Competition is a good thing. It only makes what we are going to be doing better," Dean said.

Voice of America

For more than 50 years, the Voice of America has earned the reputation of providing up-to-the-minute, accurate, and balanced news and features, as well as music and other entertainment, to its international audience. Competing with nearly 125 broad-cast services worldwide, VOA is one of the top three international broadcasters in today's vast global media market, along with BBC World Service and Radio Moscow. An estimated 92 million listeners around the world tune in to VOA programs in English and 46 other languages via direct medium wave (AM) and shortwave broadcasts. Millions more listen to VOA programs on local AM and FM stations around the world, giving VOA a vast and almost unequaled global reach.

And, as you might expect, VOA broadcasts on the Internet as well. Its Gopher site, gopher.voa.gov, offers news broadcasts in English, Spanish, Cantonese, Hungarian, Slovak, Urdu, and other languages. (See Figure 2-10.)

Voice of America doesn't use audio streaming, so as with Internet Talk Radio, you need to download a large audio file before listening in. According to the service, "These audio files are necessarily large. We have done what we could to keep them manageable, even sacrificing sound quality to some extent in an attempt to minimize their size and, consequently, the time and network bandwidth required to collect them." Downloading the audio for a ten-minute newscast through a modem can take an hour or more.

VOA's on-line audio files are encoded with Sun Microsystems' 8-bit u-law encoding at a rate of 8,000 samples per second. This format, the de facto standard for cross-platform audio files, is relatively compact and offers good sound quality. VOA's on-line news offerings are one of the most reliable and useful venues for Internet multimedia to date.

6

Gopher Menu
About the Voice of America
About Worldnet
TOA News and English Broadcasts Wire Service
Chinese Radio Scripts (Guo Biao Encoding)
TOA Internet Audio
TOA Program and Frequency Schedules
Talk to America
Worldnet Television Schedules and Satellite Downlink Information
Public Electronic Mail Addresses for VOA and Worldnet
All Files Available from WOA and Worldnet
Radio and TV Marti (Broadcasting to Cuba)
Recent Additions to the WOA and Worldnet Public Internet Server
Other Public Gopher Servers

Figure 2-10: Voice of America.

IUMA

You can't think of the cutting edge of music on the Internet without thinking of *IUMA* — the *Internet Underground Music Archive*. Combining a unique retro WWW interface and a flair for pushing the limits of the technology, IUMA is helping to bridge the gap between tomorrow's technology and today's. (See Figures 2-11 and 2-12.)

IUMA has helped pioneer much of the multimedia that's common on the Internet today, from getting college radio stations on the Net, to live audio and video broadcasts. First and foremost, IUMA is an archive of music — it provides a forum for bands and record labels to exhibit their wares online. It also sponsors live performances on the MBONE and helps college radio stations put their signals online. Besides the increasingly common feat of putting traditional radio stations on the Internet, IUMA has done the opposite: It helped a radio station program a radio show entirely from IUMA's digital archives.

With IUMA's help, San Jose College radio station KSCU-FM broadcast a full three-hour show live and direct from the IUMA computer archive. On September 14, 1994, KSCU and IUMA played 28 songs during the show. Raji Rai, the DJ for the show, commented that the new system would revolutionize the way DJs format their shows. "Having a resource of over 300 bands, easily searchable in a graphic user interface that is constantly updated from the IUMA site headquarters and complete with full biographical information about the artists is a DJ's dream come true," she said. "It means that we can format an entire show in a fraction of the time it used to take and readily provide listeners with all the band information they could ever want." According to IUMA, "despite some last-minute technical chaos and a quick run to Circuit City, the show came off very well."

37



Figure 2-11: Internet Underground Music Archive.

IUMA has plans to make its archives, along with the technology of the Internet, available to college radio stations around the country. Implementation will be through a combination of CD-ROMs that contain the IUMA site and special hardware and software.

See IUMA for yourself: http://www.iuma.com.

Virtual Radio

Virtual Radio went online in August 1994. Its goals are to provide a forum for exposure for bands and to sell music. On its face, it appears to be a lot like IUMA, providing information on bands as well as downloadable songs and song excerpts. (See Figure 2-13.)

But when it comes to making music available on the Internet, the folks at Virtual Radio take a different approach than IUMA's. "IUMA is way ahead trying to pioneer the new technology — they are really big on going for high-quality CD sound," said Brent Marcus of Virtual Radio. Of course, high-quality sound means big files to download. "We cut back on that for two reasons," Marcus said. "First is file size, and second is copyability. We don't want to give away free CD-quality music. We want people to sample the music and then buy CDs."

Note



Figure 2-12: Crusing through IUMA.



Figure 2-13: Virtual Radio.

Virtual Radio sees its role "as an on-line promoter of bands — versus IUMA, which is a music archive," Marcus said. "To outside people, it appears much the same. Our goal is to put the radio station concept online."

Because Virtual Radio can use its World Wide Web server to track exactly how many times each song has been downloaded, it can gauge the bands and music genres that are most popular.

The administrators of Virtual Radio want it to be heard by as many people as possible, so they post music in a variety of audio formats. Since each platform — Mac, Windows, and Sun, to name a few — uses a different audio format, Virtual Radio's music is posted in a format that's native to each platform. That way, users can download and hear music quickly, without going through a time-consuming (and possibly confusing) audio conversion process. "We want to reach the most people with the least work for them," Marcus said, even at the cost of disk space and staff time to prepare the sound files.

Marcus's vision of the middle future? "In five years, it would make sense to put a whole CD-quality album on the Internet and just sell that. That would be the ultimate music format."

To see Virtual Radio for yourself, point your Web browser to http://www.microserve.net/vradio.

More radio on the Internet

Check the following sources for more information about radio on the Net:

- A list of radio stations that use the World Wide Web for broadcasting or informational purposes is accessible at http://american.recordings. com/wwwofmusic/radio.html.
- Yahoo's Index of Radio Stations on the Net is at http://akebono.stanford. edu/yahoo/Entertainment/Radio/Stations/.
- The MIT List of Radio Stations on the Internet is at http://web.mit. edu/afs/athena.mit.edu/user/w/m/wmbr/www/otherstations.html.
- WXYC. http://sunsite.unc.edu/wxyc/.

Electronic Mail

Even tools as pedestrian as electronic mail benefit from the addition of multimedia. Although most e-mail that is sent and received today is composed of boring ASCII text, without the interesting fonts, color characters, or formatting tools that are available in the simplest word processors, electronic mail is beginning to see a multimedia shift of its own.

Note

Note

MIME

Today, *MIME* (the *Multi-purpose Internet Mail Extensions* specification) enables Internet users to send e-mail that contains graphics and text and has simple layout features.

MIME is a specification that offers a way to interchange text in languages that have different character sets and exchange multimedia e-mail among various computer systems. In addition to plain old ASCII text, MIME messages can include other character sets, images and sounds, PostScript files, pointers to FTPable files, and more.

Note

For more information about MIME, read the comp.mail.mime FAQ, which is available on the Internet at ftp://rtfm.mit.edu/pub/usenet-by-group/news. answers/mail/mime-faq/ and at http://www.cis.ohio-state. edu/text/faq/usenet/mail/mime-faq/top.html.

Not all e-mail clients can handle MIME, although an increasing number of them do. We use Eudora (for the Mac and Windows) and TCP/Connect II (for the Macintosh) for our MIME e-mail needs.

VideoMail

VideoMail is another example of how multimedia is enhancing electronic mail. This shareware program for the Macintosh makes it easy to e-mail a QuickTime movie of yourself hamming it up.

Using VideoMail, you can digitize yourself (using audio and video, or just audio, depending on the hardware attached to your Mac), enter an e-mail address, and watch VideoMail encode the multimedia message and fire it off to an associate in just a few seconds. The recipient receives e-mail with an encoded attachment — a QuickTime movie — of your face and voice. (See Figure 2-14.)

VideoMail works great for short messages, but you're likely to annoy your friends if you send them a 10-minute video of yourself with a cat on your head. As with all other multimedia, sending audio and video messages with VideoMail takes a good chunk of bandwidth — about 7 Kbps for audio only and 30 Kbps for audio and video. The program's author recommends keeping AV messages to 15 seconds or less, and audioonly messages to under 45 seconds. A 15-second VideoMail message with audio and video can take about 4 minutes to download at 14.4 Kbps. Many users would agree that although that time could be much faster, it approaches an acceptable download time for a message of that length.

VideoMail is not an end-all e-mail program. It was designed to be a simple, send-only multimedia mailer and is a fine complement to a full-featured e-mail program such as Eudora. Perhaps soon, full-featured mail programs will include the ability to capture and send audio and video without the help of a separate program such as VideoMail.



Figure 2-14: VideoMail screenshot.

Note

For more information about VideoMail, check out http://www.spyglass.com/~dtrinka/videomail.html.

Multimedia mail for the masses?

Multimedia electronic mail will likely become a staple of Internet communications in the coming months — certainly sooner than real-time applications such as CU-SeeMe. Live multimedia applications such as CU-SeeMe and Internet Phone need to be able to send and receive information in real time. Should your network connection slow down below a certain limit, packets of information will arrive late or get lost. If you're participating in a videoconference, a constant stream of packets (arriving in more or less the same order in which they were created) is important. Receiving a snippet of video that should have arrived ten seconds ago makes for a conversation that is, at best, hard to follow.

But electronic mail doesn't occur in real time; it doesn't matter terribly whether the message you send is received in ten seconds or two minutes. So, with multimedia e-mail, the information can travel at a more leisurely pace than it can with real-time applications. If packets arrive out of order, they can be reassembled in sequence with no loss of continuity. When your e-mail program barks that mail has arrived, you'll receive a clear message, complete with video and audio.

Some people say that in the Internet's middle future, your service provider's charges may be based on how much information you send and how quickly you want that data to arrive, rather than on how long you're online. Under this scheme, real-time audio and video, which use a large amount of bandwidth and require fast distribution to the recipient, would be expensive. On the other hand, a multimedia e-mail message would cost less to send because although it might use the same amount of bandwidth, data could be sent at a lower priority because precise timing is not critical.

41

Where the MBONE Came From

t's never necessary to know a technology's origins. Technologies exist to be used, and as long as you can do so, it really doesn't matter where they came from. At least, that's the practical view of things. For anyone hoping to *understand* a technology rather than just *use* it, however, origins matter a great deal. And for any technology that promises (or even threatens) to change the nature of entertainment, interaction, and communication, using it without understanding it just doesn't seem sufficient.

Does the MBONE really promise to change the nature of entertainment, interaction, and communication? Actually, it doesn't promise anything at all; nor would its inventors and developers be comfortable suggesting any such significance. But one thing we've all come to realize over the course of this century is that the most lasting technological designs often begin with very little sense of significance, whereas those designed with significance in mind often come up short - or not at all. Who knew, for instance, that the microwave oven or the videocassette recorder would find their way into practically all households in, at least, the U.S. and Canada? Who knew that the home computer would become practically ubiquitous? Who knew, for that matter, that the Internet would take off as it has? In all cases, certainly not the designers. On the other side, of course, we have the technologies that were supposed to change our lives but didn't. Quadraphonic sound. Ted Nelson's Xanadu. The Wankel rotary engine. And who knows how many more?

56

To say that the MBONE promises to enact change means, of course, that it holds the promise to do so. Nothing more, and nothing less. We'll look at some of that promise in Chapter 8, but for now let's examine the reasons for the MBONE's existence. By this, we're not referring to the *technological* history of the MBONE itself, at least not expressly. In this chapter, we'll look at the underpinnings of the MBONE from the standpoint of the major developments over the past decades in communications, broadcasting, and multimedia. It will be all too brief, naturally, but since its goal is to provide a *context* for the MBONE, rather than a detailed history, that brevity will serve us well.

Communications

The history of twentieth century communications begins, quite probably, in the latter part of the nineteenth century. After a series of experiments conducted by a number of people, Alexander Graham Bell put together a device that would become the telephone. Of all the technological inventions of the past 150 years, surely none has affected the course of human communications as much as this one, at least not in the nations where telephone service has become a *sine qua non* of day-to-day existence. Through this one device, businesses thrive, relationships evolve, people's lives are both made and destroyed. It is as difficult for most North Americans to imagine a world without telephones as it is for them to imagine a world without electricity. The two, it seems, go hand in hand.

Since the middle of this century, the telephone has taken on a new role. Whereas previously phone lines were relied on to transmit voice messages, they have since become the primary means for transporting digital data as well. When you log on to your Internet account and download a file from an FTP site, you are using the same telephone system you use when you call your loved ones from the other side of the continent. The difference is that the latter scenario transmits *analog* signals, and the former transmits *digital* signals. Because the phone system is an analog transport system (fiber-optics technologies are changing all this), the digital signals from your computer have to be changed into analog signals for transmission to occur. The device that does this is the *modem* (modulator-demodulator), a device that has become almost as ubiquitous as personal computers themselves.

There's a problem with the combination of modems and phone lines. They aren't very fast. That is, they allow data to be transferred at a limited rate, and the 28.8 Kbps (kilobits per second) modem is approaching the analog phone system's technological limit. (Probably. But that's what they said about 2400 bps modems, too.) By compressing the data it is sending and receiving, your modem can reach an effective transmission speed of 115,000 Kbps, but in practice, this rate is rarely achieved. Even if it were, it would be too slow to handle the huge quantities of data necessary for the transmission of audio and (especially) video files, so from the standpoint of the MBONE, the modem isn't an issue. Far greater speeds are needed, along with increasingly effective data compression technologies.

But the telephone-modem combination has provided the basis for theorizing about the MBONE or any other kind of digital communications technology. With the telephone, long-distance real-time human communication became possible (*telephone* means *to speak far*), and with the modem, direct transfer of data from one machine to another, by means of the widespread global telephone system, also became possible. The main point here is that people from all over the world could communicate with each other or exchange information with each other, and they began to expect to do so more or less instantaneously. People could pick up the phone and say hello from one side of the world to the other, or they could log into their accounts and transfer files to machines every bit as far away.

Or, by using these same computers and these same phone lines, they could replace telephonic communication with computer-based communication. To do this, they used the messaging capabilities of their networked computers — electronic mail, newsgroups, or various forms of network "chat." Electronic mail and newsgroups (the latter primarily through a subnetwork called Usenet) weren't real-time communications; you composed and sent your message, then waited for the recipients to read it, compose a reply, and send it back to you.

As for network chat, which allowed real-time communications between two Internet users, or the more popular Internet Relay Chat (IRC), in which several people linked themselves together simultaneously and communicated that way, the major problem was the fact that communicating was done through typing. To see how significant this problem is, arrange with a friend to use a chat program such as UNIX's *talk* while you're logged on together, and try to carry on a "conversation" through typing. You'll quickly discover that the typing itself becomes a topic of conversation, and that for most people typing impedes the actual communication. Then, for a really fascinating time, hook up to an IRC channel and carry on a conversation with not just one, but potentially a dozen or so other users, all trying to make themselves heard, with several different conversation, except that our eyes are much less effective than our ears at filtering through the din.

Over the years, electronic mail and newsgroups haven't changed much. Nor, for that matter, has postal mail. With their capabilities essentially at their limit, and with the telephone limited by the fact that conversations are timed for the sake of charging money (we've all experienced the nearly debilitating effect of that imaginary clock ticking while we're talking on the phone), new communications mechanisms became desirable. Then, too, there was the advent of videoconferencing, in which people from distant locations set up equipment and conducted live conversations with a full video complement. That meant that you could see each other while you talked, and therefore you could draw diagrams, demonstrate things, and otherwise make use of the sense of sight, something that the telephone and its associated audioco. ferencing (conference calls) didn't allow. The problem with videoconferencing was that it was (and is) extremely expensive, prohibitively so for anyone outside a well-funded organization.

During the past few years, two communications methods have appeared on the Internet, neither of them new but both promising to be revolutionary. The first is real-time audio communication; the second is real-time video communication. Real-time audio communication effectively means the ability of two Internet users to converse with each other by voice and in real-time — in other words, the ability to use the Internet like a telephone. The demands of real-time video over the Internet are much higher, and until a few years ago, this technology was considered impossible. The MBONE is the primary means of effecting real-time video, and its research was begun in earnest, in fact, nearly four years ago.

The first basis for the coming of the MBONE, then, is the constantly growing need for more and better forms of human communication.

Broadcasting

Today, it's extremely difficult to imagine life without broadcasting. Broadcasting is the means by which we receive, by far, the greatest portion of the information we take in each day, and to be cut off from it is considered both unusual and, if done for any length of time, dangerous. At the same time, the sheer volume of advertising has probably led to a growing inability to separate important from unimportant information, and indeed, to think about what we receive at all. That's why many of us vacation in places that actually deny our access to information. We seem to realize, somewhere underneath it all, that information overload is unhuman.

Given a choice, though (not to mention absolute political and societal power), few of us would eliminate the means by which broadcasting is accomplished. Radio is a continuous noisemaker, but it's a noise most of us have come both to accept and to expect. Television is criticized by many as notoriously ineffective as a worthwhile entertainment medium, but the TV set has become an almost inevitable fixture in today's home (in the "have" countries, at least). Another, older form of broadcasting is the newspaper, which does not broadcast "signals" per se but which unquestionably operates as a daily disseminator of information in much the same way as radio and TV. In fact, as has been well noted, a newspaper like *USA Today* is little more than a print version of televised information, complete with graphical design and brevity of detail. Still, newspapers and similar publications are not typically considered broadcasting, at least not electronic broadcasting, and we'll stick to that norm here, restricting our discussion to radio, television, and computer networks.

Radio

Radio came first. It broadcast sound, and when it did so, it completely revolutionized information dissemination. How? Consider the days before radio. To hear a political speech or debate, you and hundreds of others went to a meeting hall. To hear a new musical composition, you and hundreds of others went to a concert hall. To hear the

58

50

news, you and hundreds (well, maybe not hundreds) of others went to the town square. If you wanted a slightly dated version of the news, you bought a newspaper, took it home, and read it.

All these events still go on, of course. But radio changed them. With radio, it was suddenly possible for you and thousands of others to hear a political speech or debate while sitting, isolated from one another, in the comfort of your living room. You could sit in the same room, again all by yourself, and hear musical concerts and even full operas. And the news could be offered to you every hour of the day, constantly updated and less time-consuming to digest. You and everyone in your town, your entire state for that matter, could hear the same material without ever appearing in the same physical location.

What this did was to change forever our expectations about receiving information. For information to be considered useful, it had to be timely, and it also had to be conveniently available. Newspapers quickly became a means of fleshing out pieces of information and of providing information of more personal interest but less regional or national interest. Real news was available immediately after it happened, and it was suddenly your responsibility to keep up with it all. It's possible, in fact, that the advent of broadcast radio was the beginning of today's obsession with absolutely up-to-date current events on a global scale. Before radio, you could be up-to-date about local events through the "grapevine," but a knowledge of the outside world had to wait for the newspapers. And even here things were changing, as the grapevine quickly became a function of another growing technology, the telephone. The radio and the telephone, exclusively aural media both, cast a huge shadow over the future history of the dissemination of information.

Radio changed one other thing — the experience of theater. On the one hand, you were able to receive broadcasts of theatrical events, such as operas. On the other, a new type of theater emerged, short dramas and comedies designed and written specifically for the new medium. Knowing that you were willing to sit in your living room with the radio on, producers decided to supply entertainment (both frivolous and serious) of a kind never before available. The weekly broadcast of ongoing series began with radio, taking advantage of the the medium's immediacy, convenience, and ephemerality. These were programs for the moment, and they became fully accepted as such.

Television

Television expanded on radio's capabilities by adding video. That's so obvious it seems pointless to write down, but it's no less significant for being obvious. Radio offered a limited view of information and performance, after all; you could hear, but you could not see, and as such it was precisely the opposite of the newspaper and the magazine. Cinema fed both senses — seeing and hearing — but it was much less immediate than radio. A medium that would combine radio's immediacy and comfort with the visual capabilities of cinema, newspapers, and live theater was yet to be developed. That medium proved to be television.

60

It is impossible to overstate the significance of television as a medium of information in this century. Liking or disliking this medium has nothing to do with it; the fact remains that it has utterly dominated all other methods of acquiring information. Think of the great moments in television history — assassinations, space journeys, natural disasters, political events, van chases, Sally Field's Oscar acceptance speech — and they all point to television. Think of the enormous effect of the coverage of the Vietnam War on the U.S. populace, another television experience, and then, years later, of the brilliant (if scary) manipulation of television by the U.S. military during the Gulf War. For information to be considered completely useful today, it must not only be immediate and convenient, it must also engage our visual and auditory senses together. Especially the visual.

In many ways, television is a perfect medium for disseminating information. And, to be sure, it has single-handedly conveyed more information to more people than any other medium. As it turns out, however, it has become severely limited in this capability, largely because of a commercial model that preaches audience size above all other concerns. In other words, and ironically, television has become limited by the very fact that it is a *broad*cast medium, that the broadness is expected (by its inevitable financiers, the advertisers) to grow larger at all times. The result has been a limitation in programming types and topics, a point that remains true even with the variety of specialty channels coming to the fore.

Perhaps the most important effect of television on the MBONE has been our acceptance of television's audio and visual ability to convey information. As soon as we realized that we could experience live video on television, it was only a matter of time until we wanted to use the medium for an enormously important genre in business communication, the meeting. Meetings are nothing more than a limited number of people getting together to discuss issues of common interest, and their main limitation is the fact that all the participants have to be in the same room at the same time. Why not have a meeting with people at a remote location, with their faces displayed on a television screen and their voices coming in over the airwaves? And, in return, they'll see and hear you via TV screens as well.

The idea caught on, and it became known as videoconferencing. This idea is at the heart of the MBONE, but it's important to realize its close relationship with television, particularly television's capability to convey information on a real-time basis using both video and audio capabilities. Videoconferencing via the MBONE is one of the most exciting immediate applications for business, and videoconferencing is a seemingly natural extension of the communications capabilities of television itself.

In discussions of broadcasting, radio and television are two obvious phenomena. Less obvious, perhaps, but equally important for our purposes, is the advent of computer networking, in particular wide-area networking and internetworking. Had it not been for the capability of computers to exchange data with each other across large distances, the MBONE could not have come into existence. What must be said, however, is that the MBONE uses the Internet's broadcasting capabilities, and these aren't frequently talked about.

The Internet is well known as a person-to-person communications system. Electronic mail, in fact, remains its primary application. But electronic mail and newsgroups are broadcast technologies as well. When you send a message to a newsgroup or a mailing list, you are broadcasting on a one-to-many basis. You, an individual, broadcast information to whoever happens to be tuned into the channel — such as the newsgroup or mailing list — to which you transmit. This is, however, broadcasting of a different kind, because unlike with TV or radio, the message is not received in real time. To watch a newscast, all viewers turn on their TVs at an agreed-upon time (11 p.m., for instance); newsgroup subscribers, by contrast, can read your message whenever they feel like it, secure in the knowledge that it will be there. When it comes to time-dependent broadcasting, the Internet is only beginning to demonstrate its usefulness.

But the MBONE is not bound to time-dependant broadcasting — "on demand" technology allows MBONE users to download speeches, lectures, and concerts at any time (stopping and restarting at the user's whim) without waiting for a predetermined "broadcast" time.

Multimedia

The third technology necessary for the MBONE's development was computer multimedia. Much hyped and coming to us with endless promises, multimedia is nothing more than the incorporation of video and audio technologies into computer programs. Nothing more, but also nothing less. Most of us can remember when computers couldn't do these things at all.

A long, long time ago (which in computer terms means about 30 years), computers didn't have graphics. In fact, they didn't even have monitors on which graphics might be displayed. It wasn't until a full ten years after the beginnings of the Internet/ ARPAnet in 1969 that a home computer with even rudimentary graphics capabilities became available, and from the time it debuted, the Apple II became a mecca for multimedia designers. People bought the Apple II, it seemed, for two main reasons: They wanted to use a revolutionary program called VisiCalc, the very first spread-sheet package; and they wanted to play games.

The Apple II had sound in the form of beeps. Soon, a peripheral became available that allowed better sound. With that sound, and combined with graphics that looked notably like line drawings, artists and game designers had a multimedia environment. Back then, they didn't call it a "multimedia environment," but they had it nonetheless.

While this was going on, of course, serious multimedia research was taking place in computer labs around the world. Most of us saw the results in media such as cinema, where computers were beginning to be used for the creation of special effects. But for home computer buyers, two hugely popular products brought multimedia's potential home. First came the Atari 2600, a game machine that was in its day what the Nintendo became about several years later. This was a machine that connected directly to the television, the most popular multimedia device of all time, and delivered interactive entertainment with both graphics and sound into the living room.

And then, in the early 1980s, the Commodore 64 exploded onto the scene, changing the public's perception of computers for good. Right out of the box, the 64 offered sophisticated graphics and sounds capabilities, and in a move atypical for the company, Commodore actually promoted it. It became the favorite design platform for designers of games, educational software, easy-to-use art programs, and early music packages. It's easy to dismiss the 64's influence because of its rapid decline, but doing so would be wrong. The machine mattered.

Shortly after the introduction of the 64, IBM introduced its first PC, and multimedia development took a huge step backward. The IBM-PC was expensive, monochromatic, and text-only. Its idea of sound was a beep whenever something went wrong, and it offered no graphics capabilities whatsoever. It was bland, it was ugly, and corporate America thronged to it. Graphics and sound might be nice for home use, as the Commodore 64 proved, but real computer users didn't need these things at all.

In 1984, Apple brought the Macintosh onto the scene. In many ways, it was a less capable multimedia machine than the Commodore 64, and well behind the soon-to-beintroduced Atari ST and Commodore Amiga. But Apple did two things very right. They introduced the graphical user interface and the mouse to the public, and they marketed their machine extremely well. Artists and others in the nonbusiness and nonscience communities picked up on the Mac immediately, and it wasn't long until it became the favored machine for graphics designers and desktop publishers. It remains so to this day.

When the Amiga hit the market in 1985, it had built-in multimedia capabilities beyond all but the most multimedia-driven computers today. It had separate chips for sound and graphics (taking the load off the main processor), and it spawned games with graphics and animation quality that PC designers are only now beginning to match. The Amiga never became the machine it could have been for a variety of reasons, but its influence was widespread.

From the standpoint of multimedia, the decade since the the arrival of the Macintosh and the Amiga has been little more than an attempt by PC designers to catch up to these two machines and their enhancements. Sound cards and video boards are now available by the score for PC owners, while the Mac continues its multimedia excellence. CD-ROM at least partly took care of the huge expansion of file sizes and now allows computers to play lengthy video sequences with top-quality sound. Computer makers want their machines to look and sound like enhanced televisions (miniature cinemas if you will), and this goal is beginning to be realized.

For the MBONE, computer multimedia was a prerequisite. The MBONE relies on the fact that video and audio can work together, and that computer users can display and hear it. Without multimedia-capable computers, an MBONE broadcast would be pointless, even worse than listening to a TV station through your radio. The MBONE's future relies on the convergence of multimedia with several other related technologies.

62

6.5

Summary

The key to the MBONE's existence is convergence. Communications concepts and technologies are added to broadcast concepts and technologies, and then are mixed with multimedia concepts and technologies, with the final result being nothing less than an interdependent joining of all three. Unfortunately, like most media convergences, the result often receives less credit than it should because it seems like such a natural idea in the first place. The MBONE, a technology that is remarkable for existing at all, will almost certainly become a technology that Internet users — especially those who join during the next few years — will take entirely for granted.

But then, that's what happens to all successful technologies. Ask most 20-year olds if they're in awe of television's capability to bring live video and audio signals into their house from somewhere halfway around the globe, and you're likely to be greeted with stares of incomprehension. For that matter, try to explain to them how truly wonderful it is that your PC plays the video introduction to Wing Commander III without burping even once; the incomprehension continues. Then turn around and wax eloquently about how amazing it feels to be talking on the phone with someone three time zones away, or that you're listening to a baseball game in another city, or that you turn the key in your ignition and the car seems to start by itself.

If the MBONE works well, it will quickly be taken for granted. Indeed, its name may even disappear. You'll walk into a room and ask, "How's the MBONE performing today?" and someone at the table will say, "MBONE? No idea. Let's just get on with the videoconference." It's only a matter of time.

64 Part I + Tomorrow's Technology Today