

JANUARY 1995


COMPUTER

Innovative technology for computer professionals

ENGP
DISP
FEB 29 1995
Display periodical.
Non-circulating until:

*****5-DIGIT 95
22072280 011995 300
WASHINGTON UNIVERSITY
LIBRARIES
SERIALS DIV FM-25
UNIVERSITY OF WASHINGTON
JAN 20 1995
LIBRARIES

ENGINEERING



The Road to Software Maturity

Contents

Membership

Magazine of the

IEEE Computer Society

Circulation: *Computer* (ISSN 0018-9162) is published monthly by the IEEE Computer Society. IEEE Headquarters, 345 East 47th St., New York, NY 10017-2394; IEEE Computer Society Publications Office, 10662 Los Vaqueros Circle, PO Box 3014, Los Alamitos, CA 90720-1264, phone, (714) 821-8380; IEEE Computer Society Headquarters, 1730 Massachusetts Ave. NW, Washington, DC 20036-1903. Annual subscription included in society member dues. Nonmember subscription rates: available upon request. Single copy prices: members \$10.00; nonmembers \$20.00. This magazine is also available in microfiche form.

Postmaster: Send undelivered copies and address changes to *Computer*, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08855. Second class postage is paid at New York, New York, and at additional mailing offices. Canadian GST #125634188. Printed in USA.

Copyright and reprint permission: Copyright © 1995 by the Institute of Electrical and Electronics Engineers, Inc. All rights reserved. Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limits of US copyright law for private use of patrons: (1) those post-1977 articles that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Dr., Danvers, MA 01923; (2) pre-1978 articles without fee. For other copying, reprint, or republication permission, write to Permissions Editor, *Computer*, 10662 Los Vaqueros Circle, PO Box 3014, Los Alamitos, CA 90720-1264.

Editorial: Unless otherwise stated, bylined articles, as well as product and service descriptions, reflect the author's or firm's opinion. Inclusion in *Computer* does not necessarily constitute endorsement by the IEEE or the Computer Society. All submissions are subject to editing for style, clarity, and space.



Feature Articles

16 SEI Capability Maturity Model's Impact on Contractors

Hossein Saiedian and Richard Kuzara

With strong DoD sponsorship, more companies will probably base their software process improvement efforts on SEI's Capability Maturity Model, but the opposition is loud and clear.

28 Maintaining Data-Driven Rules in Databases

Avigdor Gal and Opher Etzion

A new model with an invariant-based language effectively handles data-driven rules in databases and uses the semantic properties of these rules to meet high-level language requirements.

40 Adaptive Parallelism and Piranha

Nicholas Carriero, Eric Freeman, David Gelernter, and David Kaminsky

Most desktop computers are idle much of the time. Piranha allows waste cycles to be recaptured by putting them to work running parallel applications.

Computing Practices

50 To Whom am I Speaking? Remote Booting in a Hostile World

Mark Lomas and Bruce Christianson

Today's networked computer systems are very vulnerable to attack. The collision-rich hash function described here permits a secure boot across a public network with no security features.

Cybersquare

55 In Search of Self-Belief: The "BOP" Phenomenon

Robert L. Glass

In this fanciful tale, software practitioners lose their self-belief to the forces working against them. There may be a lesson or two as we learn how they regain it.

58 The Challenges of Real-Time AI

David J. Musliner, James A. Hendler, Ashok K. Agrawala, Edmund H. Durfee, Jay K. Strosnider, and C.J. Paul

This article examines the emerging research area of real-time artificial intelligence, surveying new approaches that combine the best of both fields and describing some important applications.

Special Report

68 1994 Gordon Bell Prize Winners

Alan H. Karp, Michael Heath, Don Heller, and Horst Simon

For the first time in the Gordon Bell prize competition, performance numbers in fractions of a teraflop were reported. Bell was so pleased that he doubled the performance awards.

Hot Topics

Editor: Ronald D. Williams, University of Virginia, Thornton Hall/Electrical Engineering, Charlottesville, VA 22903; phone (804) 924-7960; fax (804) 924-8818; Internet, rdw@virginia.edu

Videoconferencing on the Internet

Ronald J. Vetter
North Dakota State University

Videoconferences are becoming increasingly frequent on the Internet and are generating much research interest.¹ Readily available software tools enable real-time audio and video channels as well as shared whiteboards that allow groups to collaborate on distributed group work more quickly and easily than ever (see sidebar on available tools).

The Internet infrastructure is beginning to support videoconferencing applications in several ways. First, the emerging multicast backbone (or Mbone) can efficiently send traffic from a single source over the network to multiple recipients.¹ At the same time, many workstations attached to the Internet are being equipped with video capture and sound cards to send and receive video and audio data streams. The price/performance of these hardware devices has finally reached a level that makes wide-scale deployment possible, which is perhaps the most important factor in the recent growth of videoconferencing applications.

The CU-SeeMe videoconferencing tool is also becoming very popular.² Because CU-SeeMe was designed to run on a Macintosh AV, which has built-in audio and video support, it is an inexpensive way to undertake videoconferencing on the Internet. A CU-SeeMe reflector (that is, a Unix host running appropriate control software) is the multicasting point for CU-SeeMe participants in a single videoconferencing group, whereas Mbone users use "mrounters" to support their multicast packets.

One important problem that must be addressed is the optimal placement (on the Internet) of mrounters and/or reflectors. That is, certain configurations of mrounters/reflectors will result in better utilization of network bandwidth than other placements. For example, a reflector for a local CU-SeeMe conference session should reside "electronically close" to the site with the majority of participants. When there are hundreds of potential users and many mrounters/reflectors, the problem becomes even more acute, and it is easy to set up configurations that waste bandwidth because traffic flows over inappropriate links.

Problems experienced in the virtual classroom

My work with a project to demonstrate an electronic or virtual classroom over the Internet using available videoconferencing tools has revealed several additional prob-

Large software projects require many specialists, such as programmers, database administrators, quality assurance personnel, testers, and technical writers. Implicitly, at-home workers need at least dedicated telephone lines and two-way fax and e-mail communication with coworkers and clients.

Additional issues involve dealing with performance appraisals, awards, and especially the delicate subject of employment termination. Right now, there are more questions than answers about home-office work.

Normal programming-office arrangements

As a management consultant, every year I visit 20 to 30 companies in the US and abroad. The most frequently encountered programming office arrangements have 5-foot partitions separating

- small one-person cubicles of about 48 square feet,
- two-person cubicles of about 90 square feet, or
- three-person cubicles of about 140 square feet.

I also see open office arrangements of perhaps 50 software personnel in large rooms, with file cabinets, fax machines, and supplies located along the periphery.

A common problem with these arrangements is a very short "mean time to interrupt." My informal observations in the US indicate that personnel in shared work spaces experience some kind of interruption every 10 minutes (an incoming/outgoing phone call, a social visit, or a noise or comment from the surrounding cubicles that interrupts concentration).

Assume every worker receives a phone call or a visitor about once an hour. With five or six workers in close proximity, the cumulative number of interruptions can reach significant levels very quickly. It does not have to be that way, of course, but it often is.

Several years ago, I visited a Japanese software factory that had an open office arrangement. According to the managers and personnel queried, this situation caused no apparent reduction in productivity. After a few hours at this factory, I noticed an interesting phenomenon: In a room with about 80 software personnel, barely a sound could be heard.

Several design and code reviews were taking place in separate areas set aside for this purpose, but from a distance of about 30 feet the voices were not even audible. As I was noting this, another visiting American came into the complex. We shouted greetings to each other from across the room and started talking while still more than 50 feet apart.

It suddenly occurred to me that we were making more noise than anything else in the entire building. In the US, it's not unusual to carry on rather loud conversations at a distance. Although there's no reason why we can't, we simply have not developed a group-cooperation culture that lets us work in close physical proximity without interrupting one another. Typical US programming offices are rather noisy places.

Programming is an intense mental activity that requires some periods of quiet concentration without interruptions. That is why, in the US, software office space significantly

videoconferencing tools have several limitations. (I realize that these tools may not have been designed for such an environment, but my goal is to point out important issues in distance-learning video/audio applications.)

The most troublesome problems, unrelated to conferencing tools, were hardware and operating system software problems. When teaching a class remotely and electronically, it's important that computer hardware is up and running whenever the class meets. Although this may seem easy to do, my experience proved that such reliability was hard to achieve. There were several occasions during the semester when a particular piece of hardware was down or an operating system had been upgraded, causing the virtual classroom to become inoperative. (Because the virtual classroom used workstations in a general-purpose laboratory, it was used by other faculty, staff, and students and was reserved only for a given time slot.)

Audio tools. One problem with audio transmission in

a virtual classroom environment is the number and availability of microphones. Since all audio tools allow only a single microphone per workstation, persons wishing to speak at a site with only one machine must move to the microphone and ask the coordinator to allow them to speak.

Although the audio quality was generally acceptable, there was often a large amount of disturbing feedback when the microphones at multiple sites were left "open" during a discussion. Using a push-to-talk mechanism helped reduce the feedback, but it limited interaction among on-site participants because it proved cumbersome for the coordinator to continually invoke the push-to-talk button for others. Another difficulty is coordinating when someone should start talking: Participants often begin talking too soon and thus truncate their messages.

Whiteboard tools. The performance of some whiteboard tools proved to be inadequate. It sometimes took several minutes to broadcast a simple graphic image to multiple

AVAILABLE CONFERENCING TOOLS

Collage is a shared whiteboard tool developed by the National Center for Supercomputing Applications for X Window environments. It allows multiple participants to share a common whiteboard on their desktops.

CU-SeeMe from Cornell University is a software platform that supports video and audio conferencing over the Internet. Originally designed for use with a Macintosh AV computer, CU-SeeMe is intended to provide useful conferencing at minimal cost. Receiving requires only a Mac with a screen capable of displaying 4-bit gray scale and a connection to the Internet. Sending requires a camera and either a Mac AV or a SuperMac VideoSpigot board, Quicktime, and SpigotVDIG extensions added to the system folder. Two configurations are possible: a one-to-one or, by use of a reflector, a many-to-many configuration. PCs running Windows are also now supported.

A CU-SeeMe Reflector is a Unix platform running the reflector program that allows multiparty conferencing with CU-SeeMe. The CU-SeeMe Reflector was constructed out of necessity, since there was no support in the Macintosh TCP/IP facilities for multicast. You need to use a CU-SeeMe reflector to have a multiparty conference using CU-SeeMe software on the Internet. Without reflectors, only point-to-point connections are possible. A reflector program can be used to inject CU-SeeMe audio and video into the MBone.

IVS (INRIA Videoconferencing System) is a software tool that supports audio and video conferences over the Internet. It includes a software codec with an integrated dynamic-admission-control mechanism and a protocol to manage the participants in a videoconference. IVS is based on the CCITT H.261 video compression standard.

MBone (Multicast Backbone) is a virtual network on "top" of the Internet that provides a multicasting facility to the Internet. MBone is composed of network routers or "mrounters" that support multicast. The principal MBone application tools are SD, NV, VAT, and WB (see below).

Nevot (Network Voice Terminal) is an audio-conferenc-

ing tool providing multiple party conferences with a choice of transport protocols over the Internet. It was developed by Henning Schulzrinne of AT&T Bell Laboratories.

NV (Net Video) is a videoconferencing tool that lets users transmit and receive slow-frame-rate video across the Internet. Video streams can be sent point-to-point or to several destinations simultaneously using IP multicast. Receivers need no special hardware, just an X display. The frame rate varies with the amount of motion and the bandwidth available. Frame rates of 3-5 frames per second are common for the default bandwidth of 128 Kbits per second. NV was developed by Ron Frederick of Xerox Palo Alto Research Center.

SD (Session Director) is the session director tool that announces and launches conferences on the MBone. SD provides a dynamically updated list of available sessions (for example, VAT audio conferences, NV or IVS videoconferences, and whiteboard conferences) and an easy way to join any available session or to create and advertise new sessions. The tool was written by Van Jacobson of Lawrence Livermore Laboratory (LLL).

VAT (Visual Audio Tool) is a software tool that supports multiple audio channels between conference participants over the Internet. It was developed by Van Jacobson and Steve McCanne at LLL.

WB (Whiteboard) is a collaborative software tool that supports a shared desktop whiteboard among a group of distributed users on the Internet. It was developed by Van Jacobson and Steve McCanne of LLL.

(See URL <http://ugwww.ucs.ed.ac.uk/~jaw/vconf.html> for the availability of these and other Internet videoconferencing tools. Another URL that provides a good review of desktop video conferencing tools is: http://www2.ncsu.edu/eos/service/ece/project/succeed_info/dtvc_survey/survey.html. For CU-SeeMe information, anonymous ftp to gated.cornell.edu and go to the /pub/video directory.)

participants. This is unacceptable in a real-time environment. It appeared that the problem was mainly one of scalability, since adding additional participants worsened the problem. For example, when using the Collage whiteboard tool with more than two or three remote participants, the responsiveness of the Collage tool diminished to the point where it became difficult to use. The drawing tools and simple text displays worked satisfactorily.

Video tools. At times during a classroom presentation, the video streams overwhelmed the network and caused all lab workstations to "lock up." Video sources should be able to reduce their output streams (this is called scalable rate control) during periods of network congestion, because otherwise it would be easy to flood the network with packets and adversely affect other users in a shared environment such as a campus. Such an approach doesn't make one very popular with the local network administrators!

Also, it is often impossible to tell if remote sites are "alive" and receiving video during a classroom presentation. None of the video tools I used provided feedback indicating that the images were actually being displayed on the appropriate workstations. We found ourselves often asking, "Are you still there? Are you receiving the video OK?" and so on. We recommend that tools designed for videoconferencing applications be equipped with mechanisms to inform the sending site that a receiving site is active and able to display the video (and audio for that matter).

Physically Close, Electronically Distant

We found that even though two sites are physically close, they can be electronically distant. That is, two communicating sites may be many hops away on the Internet, so that the bandwidth between them is limited by the slowest link connecting them. For example, during our virtual classroom experiments, two sites—one at North Dakota State University in Fargo, North Dakota, and the other at Moorhead State University in Moorhead, Minnesota—were only five miles apart physically, but thousands of miles apart on the Internet (see Table 1). Thus, it is very important to know how many hops separate distributed sites when using conferencing tools over electronic networks, as this will adversely affect the quality collaboration.

Table 1. Results of a "trace route" from a North Dakota State University machine to a corresponding Moorhead State University machine.

Hop number	Internet host/gateway and associated IP number	Round-trip count in milliseconds for three probes.		
1	ndgate.NoDak.edu (134.129.107.1)	0	10	0
2	spokane1-gw.nwnet.net (192.147.162.26)	110	140	120
3	seattle2-gw.nwnet.net (192.80.16.130)	130	130	140
4	seabr2-gw.nwnet.net (198.104.193.194)	240	110	260
5	enss143-fddi.nwnet.net (192.147.179.2)	230	100	80
6	t3-3.cnss88.Seattle.t3.ans.net (140.222.88.4)	80	100	80
7	t3-0.cnss8.San-Francisco.t3.ans.net (140.222.8.1)	210	230	240
8	t3-3.cnss25.Chicago.t3.ans.net (140.222.25.4)	280	560	450
9	t3-0.enss130.t3.ans.net (140.222.130.1)	490	340	200
10	cicnet-fddifw.ctd.anl.gov (192.5.180.23)	280	320	350
11	dgb-anl-2.cic.net (131.103.25.113)	350	280	370
12	dgf-fddi0.cic.net (131.103.1.19)	370	180	240
13	umn2-dgf.cic.net (131.103.9.2)	290	300	290
14	MIXNet-gw.mr.net (192.207.245.2)	240	210	320
15	MSUS-gw.MR.Net (137.192.4.4)	280	280	290
16	StCloud-GW1.MSUS.EDU (134.29.252.2)	400	200	420
17	Moorhead-GW1.MSUS.EDU (134.29.250.2)	370	230	180
18	134.29.99.253 (134.29.99.253)	160	220	160
19	coms4.moorhead.msus.edu (134.29.97.4)	190	160	220
20	dragon.moorhead.msus.edu (134.29.98.1)	350	160	410

VIDEO AND AUDIO CONFERENCING are an increasingly important way of carrying out collaborative group work. One of the first uses of the national information superhighway will likely be in the area of remote-distance education through virtual classrooms. Therefore, it is important that conferencing tools be designed to handle this application's unique requirements. As conferencing systems like Mbone and CU-SeeMe develop, we must be ready to explore their possible uses so that their limitations can be promptly recognized and corrected.

Ronald J. Vetter is an assistant professor of computer science. He can be contacted at the Computer Science Department, IACC Building, Room 258, North Dakota State University, Fargo, ND 58105-5164; e-mail, rvetter@plains.nodak.edu.

References

1. M. Macedonia and D. Brutzman, "Mbone Provides Audio and Video Across the Internet," *Computer*, Vol. 27, No. 4, Apr. 1994, pp. 30-36.
2. "Frequently Asked Questions (FAQ) on CU-SeeMe Conferencing," anonymous ftp, [gated.cornell.edu, file /pub/video/CU-SeeMe.FAQ.7-6.txt](http://gated.cornell.edu/file/pub/video/CU-SeeMe.FAQ.7-6.txt), July 6, 1994.
3. R. Vetter, "Virtual Classroom Easier Than You Think!" *Computer*, Vol. 27, No. 6, June 1994, p. 104.

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.