Scientist at the Hewlett-Packard Laboratories. Mark’s current research includes multimedia networking, high-speed gigabit networking, and the architecture/construction of the Bay Area Gigabit and Sprint Broadband MAN ATM testbeds. He has been with the Hewlett-Packard Company for 14.5 years. Mark is a member of the IETF and is the chair of the IP-over-ATM Working Group. He is the author of RFC1577 *Classical IP and ARP Over ATM*. He is a past member of the CSNET Executive Committee and a past member of the board of trustees for the Corporation for Research and Educational Networking (CREN = CSNET+BITNET).

BERRY KERCHEVAL received the Bachelor of Science degree from the University of California in 1977. He is a Member of the Research Staff at the Xerox Palo Alto Research Center in Palo Alto California. He has worked widely both in research and industry at such places as Lawrence Livermore National Laboratory and Protocol Engines, Inc. His research interests include high speed networks, multimedia and collaborative technologies.

and many willing hands to do it. Some of the challenges that will be addressed in the near-term are:

- *Switched Virtual Circuits (Q.93B)*. Rapid deployment of operational signalling code is critical to any growth in BAGNet as the existing BAGBone switches just can't handle any more PVCs.
- *Multicast implementations*. This is important to efficient use of ATM bandwidth for teleseminar use. The wide deployment of multicast in BAGNet is problematic, though, as the "right way" to do it is still a matter of some debate. Initial point-to-multipoint PVCs that connect each host to all others will provide a sort of primitive broadcast capability, but true multicast is an important feature currently lacking.
- *Security*. Putting a host on BAGNet and leaving it inside the local site's network firewall is tantamount to tearing down the firewall altogether. It leaves the local network at the mercy of the security of all the other BAGNet participants. THis is unacceptable to many site administrators, so while the3 short term solution is to place BAGNet hosts outside the local firewall on an "insecure" subnetwork, ultimately some form of connection authentication must be implemented. Active debate on the best way to do this is ongoing in the ATM over IP Working Group of the IETF.
- *Applications*. All this talk is great, but we have to get some good applications written to take advantage of the BAGNet. SO far some limited experiments with video and audio have been made, but now that the network infrastructure is in place the work of getting applications deployed can start.
- *Multivendor Interoperability*. Our experience with differing interpretations of the LLC/SNAP encapsulation standards has made us aware of the importance of interoperability testing. THis will be an ongoing part of BAGNet, to ensure that all participants can freely exchange packets with each other.

V. Summary

BAGNet is real. We have real ATM switches deployed by a real PTT sending real cells between real computers.

The right standards appear to be in the right place at the right time. The only problem is waiting for implementations to arrive from the vendors and developers.

We've had some "learning experiences", but nothing insurmountable, just the usual growing pains in any large multi-party enterprise.

ATM will grow up (and we're helping!)

And in the mean time, we'll have fun and build something useful to enrich the San Francisco Bay Area and the world in general.

For More Information

Information on the Bay Area Gigabit Testbed can be found on the World Wide Web at:

`http://george.lbl.gov/BAGNet.html`

Information on the IP over ATM Working Group of the Internet Engineering Task Force can be found on the web at:

`http://matmos.hpl.hp.com/`

References

[1] ATM-FORUM, *User Network Interface (UNI) Specification Version 3.0*, ISBN 0-13-225863-3, Prentice Hall, December 1993.

[2] Hoffman, Eric, *VINCE RELEASE 0.6 ALPHA*, release note as sent to the atm@hpl.hp.com mailing list on December 21, 1993.

[3] Heinanen, Juha, *Multiprotocol Encapsulation over ATM*, RFC-1483, March, 1993.

[4] Laubach, Mark, *Classical IP and ARP over ATM*, RFC1577, January, 1994.

[5] Johnston, B., Johnson, M., and Swinehart, D., "Bay Area Gigabit Testbed (BAGNet) - Overview", available via anonymous ftp from `george.lbl.gov:~/ftp/pub/BAGigabit/Bagnet.Release.Overview.ps`, January, 1994.

Biographical Sketches

MARK LAUBACH holds a BEE. and M.Sc. from the University of Delaware. He is an Engineer/

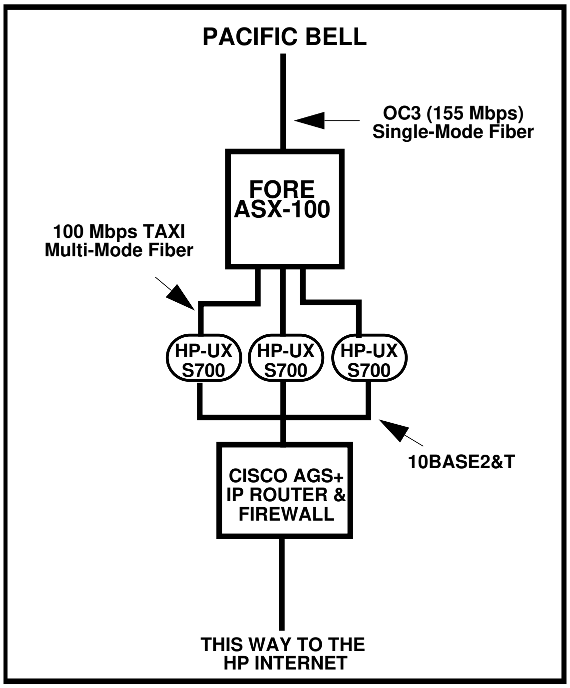

Figure x. HP Labs ATM Testbed

- Firewalls kept in IP (and higher), possible filters in the HP-UX ATM driver
- Fiber to the desktop
- ATM <> CATV interoperability

Our application experiments will include:

- IP applications that run fast: NFS, AFS, email, X, etc.
- Video over ATM in HP-UX using FORE EISA ATM host interface and MediaMagic EISA Video Card with JPEG
- Internet multicast backbone (MBONE) audio, video, and shared white board tools for desktop teleconferencing and collaborative work.
- Video on Demand

III. What's good and bad about it?

III.A. Good Things

After the fiber was installed, it was trivial to bring up the ATM connection between the two machines. We say trivial because we were able to bring up the connection between the two workstations very easily, however there are some caveats and observations:

From the point of view of the network integrator, SONET is a *wonderful* physical interface. Our SONET interfaces have line quality indicator lights on them, so when they say the link is up, it really is.

Permanent virtual channels (PVCs) are just no fun. We were able to cleanly set up the connection between our workstations because the folks at PARC and HPL have had much experience playing around with ATM. We pity larger networks as without signalling, logical IP subnets must be fully meshed. Switched virtual channels (SVCs) are really the way to go.

III.B. Bad Things

A problem that occurred during an early BAGNet test was that the software drivers for some of the ATM Host Interface cards used by PARC used drivers based on the Classical IP and ARP Internet draft which became RFC1577 [4] and implemented the default LLC/SNAP encapsulation standard. FORE Systems, Inc. however, implemented null encapsulation (VC multiplexing method) as specified in RFC1483 [3] as one of the two choices for IP over ATM encapsulation. HPLabs used FORE interfaces at the time of the test. This left PARC unable to interoperate with HPL! Fortunately though, PARC did have a couple of FORE host interface cards and we were able to interoperate

This standards clash was quite upsetting at the time, since it made it likely that some host would have to route packets from the hosts with FORE cards to the others (which would give unacceptable performance); or worse, that BAGNet would be partitioned into two logical networks that could not communicate with each other.

Luckily, Fore has recently released a patched driver (at least for some platforms) that corrects this problem and supports LLC/SNAP encapsulation. We hope all sites will be able to freely exchange packets by June of 1994.

IV. Challenges for the next 2 years.

Clearly, we are just beginning to get BAGNet operational. THere is still a lot of work to be done, PARC has pulled fibers to most of the offices in the Computer Science Laboratory. We have a mixed bag of Sun equipment ranging from Sparc IPCs to SparcServer-2000s. for ATM host adapter cards we use Fore SBA-100,and SBA-200s, and some prototype Sun SAHI cards. for software we use some Fore drivers for the SBA-200s, a modified Sun driver for the SAHI cards and our own, written to allow some signalling experiments, for the SBA-100.

We provided some input to Bellcore during the design of their Q.PORT ATM signalling code, and as a result porting that code to our switch was straightforward. We plan to finish deploying this package internally so we can "use SVCs in anger" and perform interoperability experiments with other signalling code developers.

We've been active on the multimedia front, too, and are building some ATM-capable devices that will allow audiovisual transmissions without monopolizing the users workstation.

II.G. And a HPLabs layout.

At HPLabs, Palo Alto we are actively involved in ATM technology research. At the core of our ATM activities is the construction of an ATM testbed sandbox. It is fundamental to our efforts that we understand the ATM from the ground up. Specifically, we are:

- Deploying ATM as a local area network replacement in our laboratory and workgroup scenarios. This includes the preparation necessary for installing a well thought out fiber plan. Our local network will be implemented in a seamless fashion to allow IP interoperability.
- Participating in several Bay Area ATM testbed opportunities: the Sprint Broadband MAN trial, the Bay Area Gigabit Testbed (BAGNet), the Smart Valley network, and the National Information Infrastructure Testbed.

Our investigations in ATM at HPL will eventually lead us to:

- Investigating early Q.93B signalling implementation with our other partners on the Sprint BMAN and BAGNet testbeds.
- Investigating applications that benefit from having ATM as the physical layer: high bandwidth applications, guaranteed quality of service, obligatory video teleconferencing, etc.
- Working in the standards group to help guide the future of ATM: ATM Forum and the Internet Engineering Task Force (IETF).

HP Labs has actually been waiting for the next generation backbone technology for years now. It has turned out that we will be just deploying ATM as it now appears more viable and flexible than FDDI. We have spend considerable deliberate time preplanning fiber installations throughout the Palo Alto campus. We have deployed both single mode and multimode over the past couple years between all buildings as appropriate. We even have multiple path connectivity planned so that we can avert catastrophic backhoe or earthquake failure for our critical paths. We tend to be progressive in the investigation of new technology. Where we find ourselves successful, we try to leverage our experiences over to our corporate networking group.

Inside our local ATM testbed, we are using FORE Systems Inc. ASX-100 ATM switches configured with four OC3 ports, four 140 Mbps TAXI ports, and four 100 Mbps TAXI ports. Additionally, we will be using an experimental ATM switch called "Sapphire" developed by HP Labs, Bristol. Is has six 100 Mbps TAXI ports and will be used as an extension of the FORE switch. Our hosts will be HP 9000 Series 700 HP-UX workstations. All ATM host interfaces will be EISA bus based. We will dual home all ATM hosts to an Ethernet which will be connected on the outside of our corporate IP firewall. Strict security measures will be maintained on these machines. (See Figure x.)

Our interoperability goals for the HPLabs ATM testbed are:

- IP over ATM is our interoperability standard for seamless internetworking over ATM and between ATM and non-ATM networks.
- SNMP Management of all ATM switches and host interfaces
- PVCs initially for public connections, and SVC experiments using FORE SPANs protocol and Q.93B for local and HP geographic connections
- Experiment with Q.93B when available

ATM switches must be few, due to cost reasons, yet support enough PVCs and bandwidth for the network. We figure on a minimum of two trunks connecting the two switches.

Given these constraints and an conservative estimate on the number of PVCs to allocate on each trunk (448 PVCs per port = 512 - 64 for headroom) we can calculate the number of PVCs the will traverse the Oakland-Palo Alto trunks. This will be the product of the number of hosts attached via the Palo Alto switch times the number of hosts attached via the Oakland switch. Rearranging the prior equation and solving for the number of hosts/site, we get:

$$\sqrt{((448 \times 2) / (9 \times 6))} = 4.1 \text{hosts/site}$$

for the two trunk interconnects.

This means that the interconnects are our choke point for PVC allocation.

- Point-to-multipoint PVCs will be available on a limited basis, however most available host software will not initially be able to use it.

II.E. Add a good dash of Consensus

During the course of our IP down Under planning meetings, we came to the following sets of design goals:

PacBell will initially configure the complete BAGBONE allotment of PVCs, assuming a preset allotment of 2, 3, or 4 hosts at individual sites, selected per site. Tools will be created to aid in the tedious task of configuring PVCs throughout the network.

BAGNet will not be used as an alternative path for production Internet use; i.e., BAGNet will not be used to avoid paying Internet access/use fees.

The BAGBONE will not be used as a transit network for non-BAGNet use.

As operation experience dictates, we will schedule high bandwidth needs. High bandwidth will be defined later.

No routing protocols on the BAGBONE (initially), i.e., we'll use static configurations.

Each site will have the option of making SNMP available to the network. Some sites are using the ATM switch fabrics for other users and will not be making SNMP available.

BAGNet will follow RFC1577 and RFC1483; i.e., LLC/SNAP encapsulation will be required on all BAGBONE virtual channels.

BAGBONE IP address assignments will be coordinated the old fashion way, i.e. with manual host tables, followed shortly by DNS support.

If a site needs more than 3 or 4 hosts, then they will need to implement a local IP routing solution.

Every PVC will be configured for the full link bandwidth (i.e. 155 Mbps). Our traffic management will be best effort with no peak limiting scheduling.

We would like PacBell to urge its switch vendor to implement Q.93B signalling as soon as possible. PVCs are far too painful.

The BAGNet application experiments will consist of:

- Initially, getting the network up and running.
- Trying to break the network by stressing various loads and evaluating throughput.
- Running the Internet Multicast Backbone (MBONE) tools over the BAGBONE, including mapping Class D IP addresses to point-to-multipoint VCs.
- Developing our motivating teleseminar application.

At the time of this writing, the current BAGNet installed base consists of XEROX PARC, NASA Ames, and HP Labs. We expect the other site to come on line by the end of May 1994. We have achieved limited operational testing at this time and are expecting significant results by the time this material is presented.

II.F. Pour out a Xerox Layout

Xerox PARC is connected to BAGNet by a SONET OC-3 link to the Palo Alto office of Pacific Bell.

The local ATM switch at PARC was built at PARC under an ARPA research contract to investigate the use of high-bandwidth networking in collaborative multimedia applications.

I.C. BAGNet Topology

The main switching nodes for BAGNet are distributed between two sites. One is in Palo Alto and servers the southern Bay Area and the San Francisco Peninsula. The other is located in Oakland and servers the eastern Bay Area. (See Figure 1.)

II. Recipe for Engineering the BAGNet

II.A. Start with ATM Technology.

ATM is a good choice for the base of a BAGNet. It provides higher bandwidth than traditional local are network technologies and does not use a shared medium, so that data flowing between one pair of hosts do not interfere with data between any other hosts.

For bandwidth-hungry applications like the distributed video, this helps a lot. With the emergence of the Multicast Backbone on the Internet, some local Ethernets are being heavily affected by this high-volume traffic.

II.B. Add Cooperation from the Local PTT

Pacific Bell initially installed two Newbridge ATM switches to act as the BAGNet backbone (or BAGBone as it is sometimes called). All sites were connected by OC-3 SONET links that have proved to be extremely reliable. A few sites will be upgrading to OC-12 as soon as equipment is available.

Pacific Bell personnel have been extremely responsive to the needs of the BAGNet participants. Their willingness to "go the extra mile" has been important to the early successes of the project.

II.C. Pour in a Good Dose of Standards.

> *The good thing about standards is that there are so many of them.* - Anonymous

ATM standards are currently being worked on in the ITU-TS, ANSI T1S1, and the Internet Engineering Task Force. The ATM Forum, an industrial consortium, is also working on implementation references for private ATM networks and for private-to-public interconnections.

The ATM Forum has recently published its User Network Interface (UNI) Specification, Version 3.0 [1]. The UNI 3.0 implementation reference, specifies the initial aspects of using Q.93B signalling for ATM local area networks.

The Internet Engineering Task Force (IETF) IP-over-ATM Working Group has produced two Request for Comment (RFC) proposed standards for the encapsulation of IP packets over ATM Adaptation Layer 5 [3] and for the implementation of IP and ARP within a logical IP subnet over ATM [4]. Additional Request for Comments are expected to be out soon for specifying path MTU discovery and for an implementation (coding) guide when using switched (dynamic) virtual channels and the ATM Forum's User Network Interface (UNI) Version 3.0 Implementation Reference [1].

II.D. Mix with the Problem's Constraints

Engineering implies the creation of a satisfactory solution to a problem, crafted with available technology, while obeying certain rules or constraints. We have a fine collection of constraints within BAGNet. Some of these are imposed by the available technology, and some are self-imposed.

- PacBell has selected a vendor whose ATM switches:

  - will not support Q.93B signalling (SVCs) until early 1995.

  - provide only 512 configurable VC entries per port.

- Each host in the IP backbone subnet must be able to communicate with all other hosts on the backbone via a fully connected PVC mesh.
- Per port PVC problems: if each site has 6 IP hosts and there are 15 sites, there will be 90 total hosts on the backbone. Each host must be connected to each other hosts at the 14 others site. This requires roughly:

  $$6\text{ hosts} \times (6\text{ hosts} \times 14\text{ sites}) = 504\text{ PVCs}$$

  behind each site port, which is uncomfortably close to the 512 VC / site port limit allowing no headroom.
- Interconnect trunk PVC problem: the interconnect trunks between the Oakland and Palo Alto

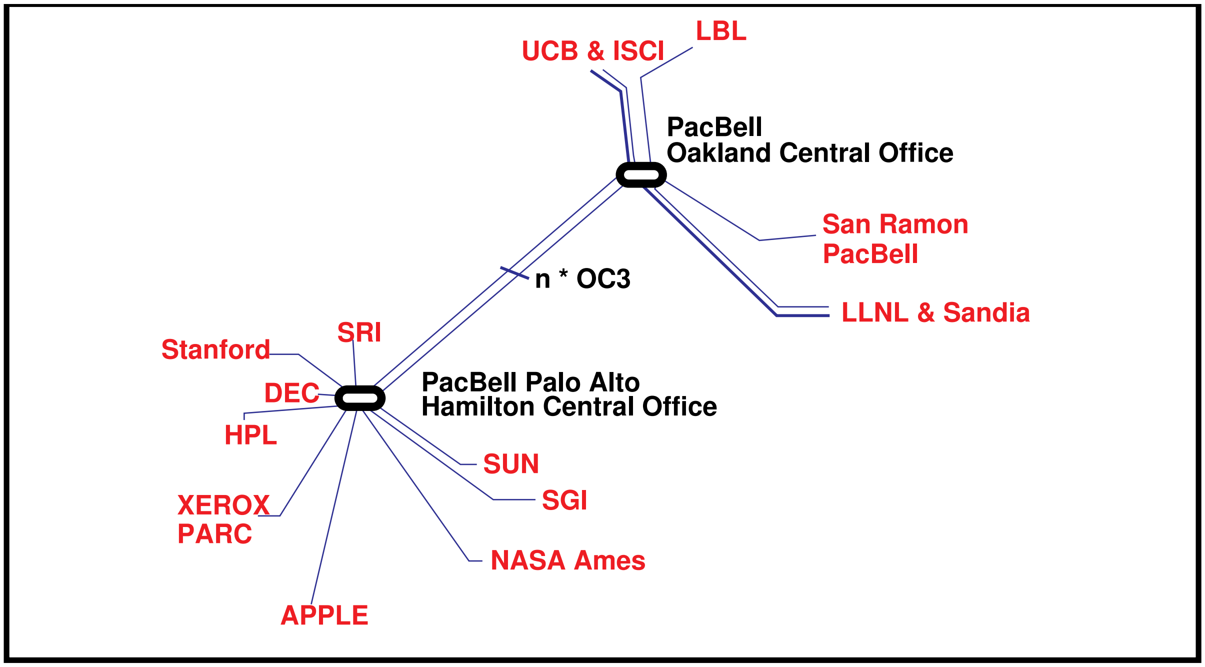

Figure 1. BAGNet Topology

only.

After several years of planning, and many funding and program management discussions, BAGNet was awarded a grant from Pacific Bell's CalREN (California Research and Education Network) program for a two-year period. The two-year program began on December 30th, 1993 as the first two BAGNet sites were connected.

CALREN is Pacific Bell's unique U.S. $27 million trust for stimulating the creation of high-speed applications to run on the communications superhighway. CalREN award winners are selected by an external council made up of technology and community-of-interest experts.

I.B. Application Focus

BAGNet's is directing its efforts toward a "common focus" application. This allows participants to work toward a well defined goal and avoids diffusion of effort.

This "common focus" application is a generalized teleseminar capability. The goal is to provide at the desktop the same richness of group interaction found in a spirited debate among scientists, educators, politicians, and others. [5]

There is a wealth of technical information available in the Bay Area in the form of seminars presented at universities, government laboratories, and technology companies. For many researchers who potentially would attend seminars, or professionals trying to do continuing education; schedule, time, distance, and commuting congestion constraints prohibit attendance at all but a small portion of the seminars that would be beneficial to their work or skills enhancement. The "teleseminar" application will bring this interactive information to the desktop of the participant. In one mode, a primary speaker can be lecturing, arguing with the audience, presenting multimedia material (audio, video, still images and text, impromptu whiteboard drawings), and so on. The audience can interact with the speaker or each other. In all of these modes, the speaker can "see" the audience, the audience can see the speaker and each other. Another mode might provide meetings, either scheduled or arranged casually through e-mail between any interested participants on whatever topic they like. Unlike "traditional" videoconferences with their expensive rooms that must be reserved days in advance, there will be no incremental expense associated with the use of this technology, which the participants hope will lead to widespread synergistic sue of it.

Bay Area Broadband Testbeds and Technology

Mark Laubach <laubach@hpl.hp.com>
Berry Kercheval <kerch@parc.xerox.com>

Abstract

BAGNet is a gigabit network testbed in California's San Francisco Bay Area funded by the California Research and Education Network (CalREN). Fourteen organizations are participating in the project to build and use a multimedia network infrastructure. The "featured application" is a distributed teleseminar system that will enable rich group interactions of the kind found in traditional face-to-face meetings, using both prepared multimedia presentations and live audio and video of speakers and audiences.

The ATM backbone provided by Pacific Bell has been deployed and connections to all of the participants have been complete. Preliminary tests of connectivity proceed.

In addition to the teleseminar work, experiments in ATM network infrastructure multimedia technology are planned.

This work was supported by Xerox. Portions were also paid for by ARPA under contract DABT63-92-C-0034; no official endorsement is implied.

I. Introduction

HP Labs and Xerox PARC are participating with thirteen other organizations to construct a nationally recognized gigabit ATM testbed in the Bay Area (BAGNet). Pacific Bell (PacBell) is the telecommunications provider for this network. Funding should be provided by PacBell's recently established California Research and Education Network (CALREN) program and by other government funding. BAGNet will be a two year project investigating a teleseminar and other applications built over PacBell's production ATM network services offering.

The BAGNet membership is: Apple Computer, Digital Equipment Corp. (DEC), Lawrence Berkeley Labs (LBL), Lawrence Livermore National Labs (LLNL), Hewlett-Packard(HP), the International Computer Science Institute (ICSI), the University of California (Berkeley Campus), the National Aeronautic and Space Agency Ames Research Center (NASA Ames), Sandia Laboratories, Silicon Graphics, Inc. (SGI), Stanford Research Institute (SRI), Stanford University, Sun Microsystems, Xerox PARC, and Pacific Bell.

Pacific Bell formally turned on BAGNet as part of a marketing trial on December 30th, 1993. The first two sites connected were Xerox PARC and NASA Ames. The remainder of the sites will be connected by the end of May, 1994.

BAGNet is providing ATM via OC-3 SONET access for its membership. PacBell will be providing production quality services and support for the network, in much the same way they provide their other tariffed services (voice, ISDN, etc.). There will be some room for experimentation with the ATM networking layer in BAGNet.

I.A. How the BAGNet and CalREN came to be

In 1989, fourteen organizations came together and began planning a gigabit testbed opportunity for the San Francisco Bay Area. The original intentions were to follow in the same "gigabit footsteps" track as the original five U.S. national gigabit testbeds: AURORA, BLANCA, CASA, MAGIC, and VISTA. BAGNet was to be different in its funding approach. The five U.S. testbeds were funded from a combination of government grants via the Corporation for National Research Initiatives (CNRI) and from the Advanced Research Projects Agency (ARPA). The original BAGNet model called for most of the costs being shared amongst the industrial participants with some government funding for the educational and national laboratory participants