

4-3-96

SANDIA REPORT

SAND96-0820 • UC-705

Unlimited Release

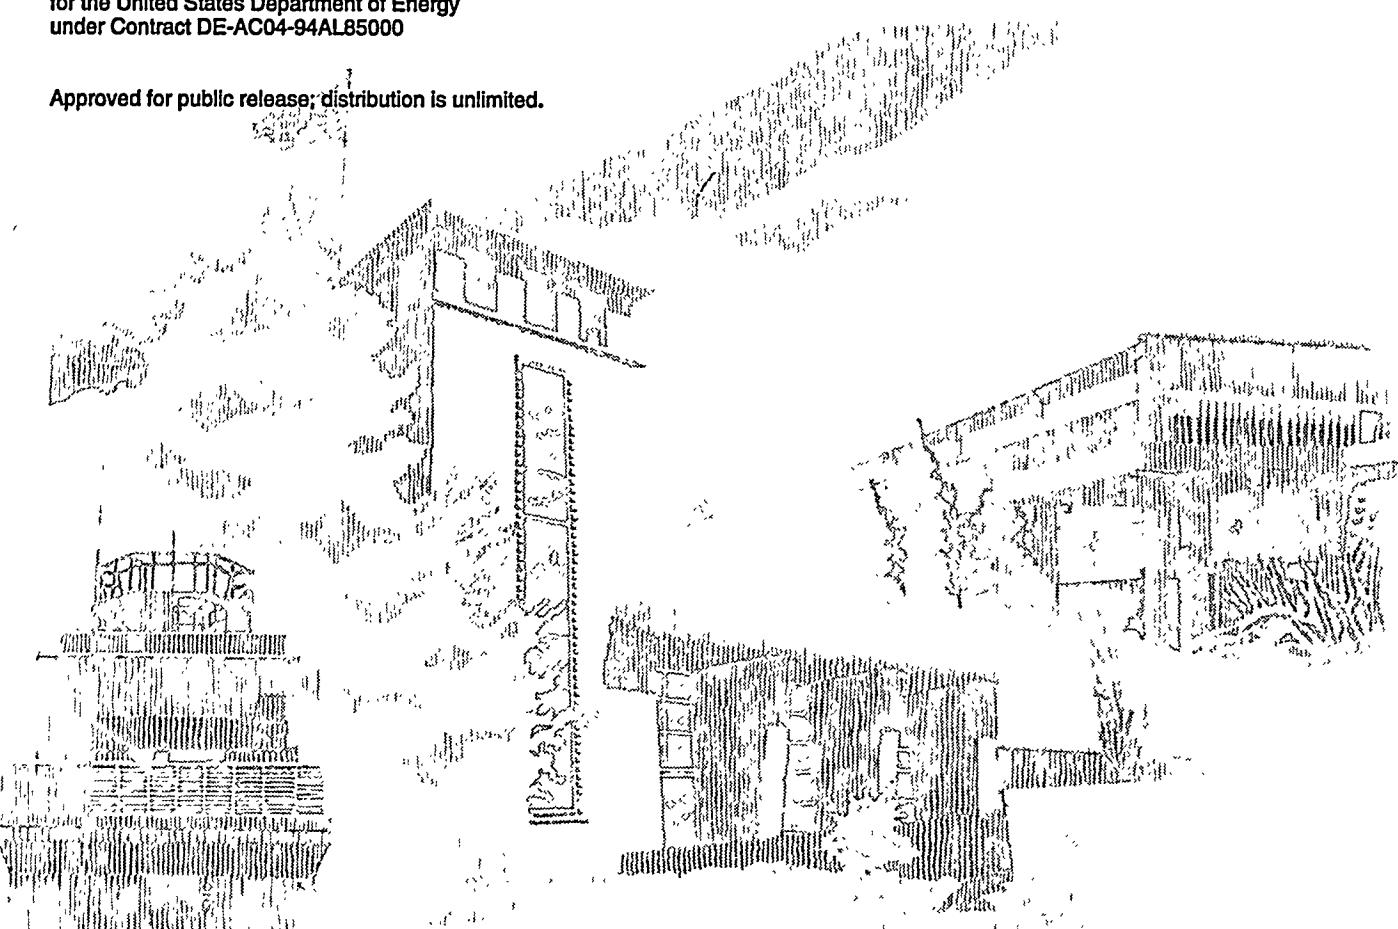
Printed April 1996

Sandia's Network for Supercomputing '95: Validating the Progress of Asynchronous Transfer Mode (ATM) Switching

Thomas J. Pratt, Michael O. Vahle, Steven A. Gossage

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550
for the United States Department of Energy
under Contract DE-AC04-94AL85000

Approved for public release; distribution is unlimited.



SF2900Q(8-81)

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

MASTER

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
Office of Scientific and Technical Information
PO Box 62
Oak Ridge, TN 37831

Prices available from (615) 576-8401, FTS 626-8401

Available to the public from
National Technical Information Service
US Department of Commerce
5285 Port Royal Rd
Springfield, VA 22161

NTIS price codes
Printed copy: A03
Microfiche copy: A01

**SAND96-0820
Unlimited Release**

**Distribution
Category UC—705**

Printed April 1996

Sandia's Network for Supercomputing '95: Validating the Progress of Asynchronous Transfer Mode (ATM) Switching

Thomas J. Pratt, Michael O. Vahle,

and Steven A. Gossage

Advanced Networking Integration Department

Sandia National Laboratories

Albuquerque, NM 87185

Abstract

The Advanced Networking Integration Department at Sandia National Laboratories has used the annual Supercomputing conference sponsored by the IEEE and ACM for the past three years as a forum to demonstrate and focus communication and networking developments. For *Supercomputing '95*, Sandia elected: to demonstrate the functionality and capability of an AT&T Globeview 20Gbps Asynchronous Transfer Mode (ATM) switch, which represents the core of Sandia's corporate network, to build and utilize a three node 622 megabit per second Paragon network, and to extend the DOD's ACTS ATM Internet from Sandia, New Mexico to the conference's show floor in San Diego, California, for video demonstrations. This paper documents those accomplishments, discusses the details of their implementation, and describes how these demonstrations supports Sandia's overall strategies in ATM networking.

Contents

1	Introduction	1
2	History of Sandia's Participation	3
3	Sandia's Supercomputing 95 ATM Networks	4
4	ATM/SONET OC-12 Interfaces	8
5	Sandia's Corporate Video Demonstration	10
6	Lessons Learned	12
7	Conclusion	12
8	Acknowledgements	13

List of Figures

1	Sandia's <i>Supercomputing '95</i> ATM Network	5
2	An OC-12 Based Supercomputing Cluster	9
3	Sandia's <i>Supercomputing '95</i> Video Network	11

List of Tables

1	Circuit Entry Form for the Globeview 2000	6
2	VP/VC Map for the Globeview 2000	6



1 Introduction

The Advanced Networking Integration Department at Sandia National Laboratories has used the annual Supercomputing conference sponsored by the IEEE and ACM for the past three years as a forum to demonstrate and focus communication and networking developments. The department's participation began in Minneapolis at the 1992 conference with a demonstration of the prototype Switched Multimegabit Data and Synchronous Optical Network (SONET) technology that Sandia was intending to use in its consolidation of their supercomputers [?, ?, ?]. As a direct result of this participation, the National Information Infrastructure Testbed (NIIT) was born. The following year's conference in Portland focused on the interoperability of emerging ATM technology and its efficacy in providing high quality video and multimedia capability [?]. This conference resulted in an early pilot of an interconnection of the three DOE DP National Laboratories over a capable wide area network. In Washington the following year, the three Laboratories once again were connected over Sandia's extension of its production networks to the conference's trade show floor [?]. In all cases, noteworthy contributions from Sandia technical partners, especially AT&T and US West, made the results possible and added significance to the accomplishments.

Some common themes and benefits at each of the conferences have been:

- demonstrating the power of DS-3 and SONET communications for enterprise network extension,
- demonstrating the application of routing and local network switching technologies,
- partnering with industry to gain early access to new technology,
- focusing current projects and activities by preparing challenging demonstrations,
- engendering new and evolving old partnerships with industry, academia, and the other government labs and agencies to work on common problems,
- highlighting the synergy that results from the tight coupling of networking and communication technologies and organizations,
- providing a stage to professionally interact with colleagues and associates from other organizations in order to challenge and validate our current thinking.

For the 1995 conference in San Diego [?], Sandia once again elected to use this forum to highlight its current networking and communications activities and to share and validate its current thinking and strategies. Specifically, this year Sandia elected:

- to demonstrate the functionality and capability of an AT&T Globeview 20Gbps ATM switch¹, which mimics the capability recently installed in the core of Sandia's internal network [?, ?],
- to build and utilize, in collaboration with Oak Ridge National Laboratory (ORNL), Intel, and Giganet, a three node 622 megabits per second (mbps) Paragon network,
- and to extend, with DOD partners, the ACTS ATM Internet ² from Sandia to the show floor to allow real-time experimentation with resources (High Performance Storage System, corporate video network, and virtual reality laboratory) back at Sandia's main campus.

Each year, as part of the conference, a team of volunteers constructs a local area network, called *SCI**net*, to interconnect the computing resources in the vendor and research exhibits on the showroom floor. Sandia always participates in this endeavor and expands the network to include wide area connections. Sandia's network extensions demonstrate the frontiers of what is currently possible in high-performance networking and communications for the wide area. Volunteers from Sandia have played major roles in designing and building *SCI**net*³ since its inception at the Albuquerque show in 1991 .

The remainder of this paper describes and documents Sandia's participation at *Supercomputing '95*. The history of Sandia's participation in the conference are discussed in Section ???. The ATM network Sandia built for the conference is described in Section ???, the demostration of an OC-12 interface is described in Section ???, and the extension of Sandia's video network to the show floor is detailed in Section ???. Finally, lessons learned are discussed in Section ??.

¹The Globeview 2000 switch was provided by AT&T through the efforts of O. Suarez and supported at the show by R. Zahorak and A. Ghosh.

²The network extension was supported by Sprint, Army Research Laboratory, and the Air Force Phillips Laboratory. Throughout this paper, this network will be referred to as the Internet Defense Research Engineering Network (IDREN)

³Luis Martinez, Sandia National Laboratories, played instrumental roles in building 1995's *SCI**net*.

2 History of Sandia's Participation

Sandia's networking and communications departments have participated in the conference since 1991 when it was held in Albuquerque, New Mexico. Sandia's participation in that show included providing three DS-3 (44.736 mbps) circuits over an Optical Carrier circuit running at 155.52 mbps (OC-3). Sandia also used a High Performance Parallel Interface (HIPPI) gigabit extension over U S WEST provided dark fiber⁴ between Sandia's Massively Parallel Computing Research Laboratory (MPCRL) and its booth at the Albuquerque Convention Center [?].

At the following year's conference in Minneapolis, Sandia demonstrated the prototype technology it was planning to use in its Supercomputing consolidation that was planned for the following spring [?][?]. The network consisted of SMDS based on prototype AT&T BNS-2000 switches with DS-3 (44.736 mbps) access [?]. This network provided the longest, fastest, private implementation of the technology at the time of the conference [?][?].

The following year at *Supercomputing '93* in Portland, Oregon, Sandia built three high-performance networks [?][?]. The networks encompassed an SMDS network running at DS-3, an Asynchronous Transfer Mode (ATM) network running on a SONET circuit at the OC-3 rate, and a High Performance Parallel Interface (HIPPI) network running over a 622.08 mbps SONET circuit (OC-12). The SMDS and ATM networks extended from Albuquerque, New Mexico to the showroom floor in Portland, Oregon, while the HIPPI/SONET network extended from Beaverton, Oregon to the showroom floor.

In 1994, utilizing private SMDS network facilities at Sandia combined with an AT&T FTS2000 DS-3 circuit terminated in an SMDS Cisco router in Washington, D.C., a transcontinental SMDS network was constructed [?]. This network linked the advanced supercomputing environments of Sandia National Laboratories (SNL), Los Alamos National Laboratories (LANL), and Lawrence Livermore National Laboratories (LLNL) to the show floor at *Supercomputing '94*. A series of demonstrations allowed the Laboratories to showcase the capabilities being jointly developed to support a variety of Department of Energy initiatives. The unique combination public and private network was fully connected at DS-3 (44.736 Mbps) rates utilizing AT&T BNS-2000 switches located at Sandia's New Mexico and California sites and an AT&T⁵ extension to Washington, D.C. AT&T BNS-2000 switches provided connectionless SMDS functionality via ATM cell switching. Sandia's private SMDS network, extending between California and New Mexico over a private DS-3 link, was extended

⁴Fiber without any transmission equipment attached so that no light is present and no service is being offered.

⁵Via the FTS2000 public network.

temporarily to LANL (approximately 90 miles) over a DS-3 circuit provided by U S WEST. Similarly, Sandia's California network was temporarily extended to LLNL via FDDI. AT&T provided the DS-3 interexchange carrier (IEC) circuit to Washington, D.C. from Albuquerque, New Mexico allowing Sandia to place an SMDS router on the show floor in Washington, D.C. The DS-3 link was extended to the *SCI**net* '94 booth by Metropolitan Fiber Service (MFS). An AT&T DDM2000 multiplexer was used to extend the DS-3 circuit from the *SCI**net* '94 booth to Sandia's booth. *SCI**net* '94 provided the dark fibers that allowed an FDDI ring to connect the research booths of LANL, LLNL and SNL. Conference attendees could use LAN-attached workstations at the individual booths to access the Massively Parallel Computing Laboratory (MPCRL) and other computing resources at the National Laboratories including Sandia's California Heterogeneous Environment and Testbed (HEAT).

3 Sandia's Supercomputing 95 ATM Networks

Sandia designed a large ATM network for *Supercomputing* '95 (see Figure ??). The network built for Sandia's booth consisted of five switches, three FORE Systems ASX200, an AT&T GlobeView 2000, and a DEC AN2. With the exception of the DEC AN2 switch, the network was completely interconnected, allowing all users access to the entire ATM network. The AN2 switch operated as a stand-alone switch connecting eight Dec Alpha computers through OC-3 ATM ports and was used to demonstrate low latency distributed computing using Parallel Virtual Machine (PVM). The primary reason that the switch wasn't interconnected to the rest of the ATM network was that all of its ports were required to perform the PVM demonstration.

The hub of the booth network was an AT&T GlobeView 2000 which also was the interconnection point for the external clouds. The GlobeView had 12 OC-3 single mode ports of which 7 were actively used. Three AT&T Enhanced MultiMedia Interfaces (EMMIs) were directly attached to the GlobeView. Using Permanent Virtual Circuits (PVCs) every monitor was able to show video from every camera and video player that was in the booth. The FORE Systems switches were placed at the edge of the ATM network. The first FORE switch contained three OC-12 ATM modules and one OC-3 module. The OC-12 modules contained one port per module. These OC-12 ports were used to interconnect three Intel Paragon on the *Supercomputing* '95 exhibit floor together into a cluster of supercomputers (see Figure ?? and Section ??). The OC-3 module on this switch was used to provide network access to OC-3 ATM devices.

An EMMI, connected via the OC-3 module, was the source for the real-

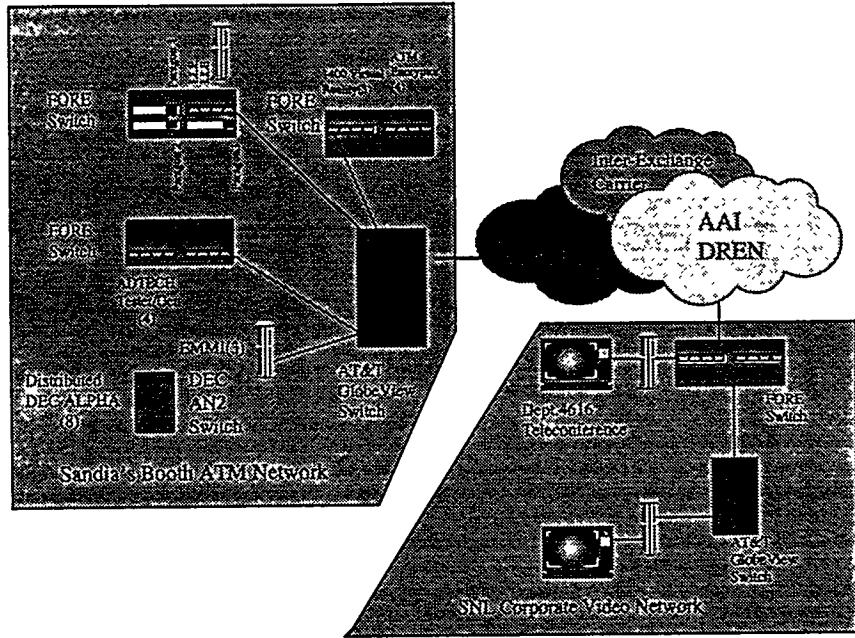


Figure 1: Sandia's *Supercomputing '95* ATM Network

time video application. Also connected by OC-3 ports was the BBN LongLine Emulator (LLE). The LLE was used to simulate the latency of a network going between Sandia in New Mexico and ORNL in Tennessee. The last OC-3 port on the Fore Switch was connected to the Globeview 2000 switch. The second FORE switch contained two OC-3 multi-mode modules giving eight OC-3 ports in total. These ports were set aside for the use of the other researchers in the booth. Four ports were assigned to the End-to-End Encryption project. Two ports were set aside for the Virtual Reality demonstrations. Another port was used to interconnect this switch to the Globeview 2000. The third Fore Switch was connected to the ADTECH 4000 Cell Tester/Generator. This connection made it possible to route any ATM traffic created in the network to the tester. The tester was primarily used by the End-to-End Encryption demonstration. Four of the tester ports were connected to the switch providing 4 separate testing channels for the booth. During the show, the booth's ATM logical network consisted of over 36 permanent virtual channels, PVC. A data form (see Table ??) and spreadsheet (see Table ??) were created to track the PVC and to facilitate communication between the applications and the customer engineer entering PVC on the switches.

The *Supercomputing '95* booth network was connected to Sandia's External Open Network ATM cloud in New Mexico through Sandia's connection to the IDREN. The IDREN network is an existing ATM network that is con-

Table 1: Circuit Entry Form for the Globeview 2000

Circuit Entry Form for the Globeview			
Source End Equip.		Destination End Equip.	
	Name		
	LTP		
	VPI		
	VCI		

Table 2: VP/VC Map for the Globeview 2000

Source				Destination				Circuit Name/Description	
LTP	VPI	VCI	Src. Name	LTP	VPI	VCI	Dst. Name		
1013	0	51	Paragon EMMI	11	0	51	Fore OC-12	Sandia Paragon Demo Video	1
11	0	52	Fore OC-12	1013	0	52	Paragon EMMI	Sandia Paragon Demo Video Rtn	2
1013	0	53	Paragon EMMI	11	0	53	Fore OC-12	ORNL Paragon Demo Video	3
1011	0		Main EMMI	1013	0		Paragon EMMI	Main Paragon Video	4
1011	0		Main EMMI	1013	0		Paragon EMMI	Main Audio Paragon	5
1013	0		Paragon EMMI	1013	0		Paragon EMMI	Sandia Paragon Demo Audio	6
1011	0		Main EMMI	1011	0		Main EMMI2	Main Loop Video	7
1011	0		Main EMMI	1011	0		Main EMMI2	Main Loop Audio	8
10	0		Scinet	1012	0		Steve's PC	Sandia Telecon Video Rtn	9
1012	0		Steve's PC	10	0		Scinet	Sandia Telecon Video	10
10	9		Scinet	1012	0		Steve's PC	Sandia Telecon Audio Rtn	11
1012	0		Steve's PC	10	9		Scinet	Sandia Telecon Audio	12
1010	0		EMMIS3	1011	0		Main	SC Demo Audio Main	13
1010	0		EMMIS3	1011	0		Main	SC Demo Video Main	14
1010	0		EMMIS3	11	9		Scinet	SC Demo Audio NM	15
1010	0		EMMIS3	10	0		Scinet	SC Demo Video NM	16
10	0		Scinet	1011	0		Main	Sandia Telecon Main Video	17
10	0		Scinet	1011	0		Main	Sandia Video Main Video	18
10	9		Scinet	1011	0		Main	Sandia Telecon Main Video	19
10	9		Scinet	1011	0		Main	Sandia Video Main Audio	20

ected to Sandia through the Air Force Phillips Laboratory on Kirkland Air force Base. The IDREN network provider is Sprint. When Sandia became aware that the Sprint was extending the IDREN to *Supercomputing '95* to support researchers at the Army Research Laboratory (ARL), Sandia requested, the Air Force Phillips Laboratory's concurrence, through the IDREN that Sprint provision two Permanent Virtual Paths (PVPs) from San Diego Convention Center to the Air Force Phillips Laboratory. One virtual path was specified to be a 30 mbps path for ATM Adaptation Layer (AAL) 5 type traffic. The second virtual path was to be a 2 mbps Constant Bit Rate (CBR) type circuit to carry AAL1 traffic. Although Sprint seemed able to fulfill this service request through their network, they weren't able to get the PVP circuits transported through the *SCInet* cloud. However, *SCInet* agreed to provide Sandia's booth with three PVCs. In addition, the Sandia booth wasn't able to get any CBR services from

*SCI**net*. Although the lack of CBR caused Sandia's team many problems, the team was able to successfully pass CBR traffic through the *Supercomputing '95* cloud using one of the three PVCs.

The lack of a PVP impacted some of the planned demonstrations. It forced Sandia to limit its wide area activities to just one video source and destination operating at a time. It also forced Sandia to reconfigure the booth's logical ATM network as demonstrations were changed. Since there wasn't any Quality of Service (QOS) on the *SCI**net* PVCs, the CBR cell stream wasn't within specification for the audio traffic that utilized the circuit. This caused audio synchronization loss that manifested itself as everything from an occasional buzzing or popping sound to a complete loss of audio that required the audio circuit to be reset. A complete loss synchronization occurred approximately once an hour. Since *SCI**net* didn't police the PVCs, the network exceeded the 30 mbps data rate for periods at *Supercomputing '95*. But it also meant that there were no guarantees that traffic would not be dropped. There was evidence that cells were occasionally being dropped in the network. But overall the network was able to pass cells reliably enough to provide a quality video picture and a usable audio signal.

The Sprint cloud and IDREN circuits took traffic from the *SCI**net* cloud to the IDREN point of presence at the Air Force Phillips Laboratory. The lowest speed link was a DS-3 circuit that was the local loop connection at Phillips. From Phillips, the circuit was an OC-3 connection into Sandia EON cloud. Just before the conference, two permanent virtual paths were constructed between the Phillips' switch and Sandia's switch. At Sandia, the path was terminated on the Fore System switch that serves as Sandia's IDREN entry point. Virtual circuits were constructed from the video head in Building 880's Technical Control Center to the Sandia IDREN edge switch. A loopback VC on each path was provisioned. This VC provided a way of testing the transmission path from San Diego without having anyone at Sandia working on the circuit. Circuits to Department 4616's office area in Building 880, where there was a live camera. This camera traffic was the only traffic running across the network until Tuesday morning. Although there was only this test circuit running, the traffic from this circuit set a record for traffic sourced from a single site on the IDREN. The traffic averaged 79% of a DS-3. This record was beaten on the following Tuesday, the 5th of December and Wednesday the 6th of December. On Tuesday, the traffic peak reached 89% of the DS-3. This traffic was being generated by the video demonstrations in the booth. On Wednesday, Sprint requested a burst test. For the test period, Sandia was able to generate 99.63% of the DS-3 capacity. Because the DS-3 was being overloaded at the OC-3 to DS-3 transition, the network was dropping many cells. This made the output unusable. The remaining .37% of the DS-3 was being used by Sprint for ATM signaling. It appeared that Sandia's traffic was the bulk of the sustained traffic

that was running on the Sprint *Supercomputing '95* network. Sprint entered the San Diego Convention Center on an OC-3 circuit. On this OC-3, the 15 minute peaks rate on Tuesday was 21%. An approximation is that our traffic was making up over 90% of this traffic. During the burst test with all of the parties participating, the OC-3 peak rate was recorded as 40% with Sandia's traffic making up over 50% of that traffic.

Although the booth network was up and running throughout the conference, many switch problems were seen. In the booth, a Fore Switch operating system would lock up requiring a team member to reset the entire unit. Also, a Fore System switch's operating system added two PVC when only one was entered. The bogus VP could only be deleted by deleting the real VP that had been added. Another problem that caused service outage was overdriving the *SCI*net edge switch. This caused the path to Sandia to be down for several hours. The addition of 15db of optical attenuation between the GlobeView and the edge switch was required to fix the problem. On both the 5th and 6th of December, the GlobeView's operating system disconnected from the switch. With AT&T Hot-line support, the unit was back in service in less than an hour. Although no PVCs were dropped during this time, the outage resulted in a loss of the booth traffic statistics. On the morning of the 7th of December, the external network stopped working because *SCI*net had spent the previous night reconfiguring the ATM network. At Sandia, the operating system of the Fore switch disconnected the console several times. To recover the console required that the switch be rebooted. Also, the switch control network was unreachable whenever the *Supercomputing '95* IP network loss connection to DOE's Energy Science network (ESNET). The Phillips ATM switch also experienced the operating system lockup problem. Human errors were also made. A logical loop was entered into one of the Fore switches causing congestion and affecting user applications.

4 ATM/SONET OC-12 Interfaces

The Linking Massively Parallel Remote Hosts with ATM/SONET OC-12 Interfaces demonstration (see Figure ??) was designed to show what could be done with a high performance I/O interface in the INTEL Paragon. It highlighted the OC-12C ATM board that is being developed by Giganet Inc. for Sandia. The demonstration was physically setup to show three types of performance.

The use of multiple OC-12 interfaces connecting between the machines demonstrated using the interfaces as a high capacity data interconnect between co-located machines. The OC-12 switched connection demonstrated the ability to create a high performance data network. This network allowed us to built a closely coupled supercomputing cluster. The addition of an OC-3 module

Figure 2: An OC-12 Based Supercomputing Cluster

into the switch showed how this high performance network could have connectivity to Wide Area Networks (WAN) and other existing infrastructure. Three applications were run across this network. They were:

1. PCTH,
2. an ORNL material science application and
3. a real-time video application.

For the real-time video application, an EMMI was connected to the Paragon via the switch OC-3 module. This EMMI was the source for the real-time video application. Sandia also connected the LLE to the network. The LLE was used to simulate the latency of a network going between Sandia in New Mexico and ORNL. This allowed us to see how the applications were effected by the delay of a large WAN network. The Supercomputing environment provided a unique environment that brought together researchers, application developers, and communication engineers from ORNL and Sandia together with product designers, application developers, and communication engineers from Giganet and INTEL. The exhibit floor provided a closely coupled laboratory for all the interested parties to come together to do experiments and characterize systems.

PCTH is an impact physic code developed and maintained by Sandia. PCTH is expected to be one of the applications that could be used to demonstrate unique capabilities of a closely couple computing environment that will be set up between ORNL and Sandia. A small problem was run on Paragon in the Sandia booth and the Paragon in the ORNL booth. Communications between the two machine were processed by PVM, a message passing code developed at ORNL. After base-lining the performance, both running on a single machine and running on two machines with minimum delay between them, the traffic was routed from the OC-12 network into an OC-3 network that contained a delay simulator. With zero delay, the performance didn't degrade on the distributed problem. Performance was effected by delays greater than 1 millisecond. With a 25 millisecond delay in the line, the code performance had degraded by a factor of four. These experiments demonstrated that it is possible to hide much of the WAN latency by continuing to do calculations while the communication was taking place. Current versions of the code assumes that communication latency is minimal. The PVM code also could be optimized for the WAN environment. Currently PVM is capable of sourcing data at approximately 2 Mbytes/second. By rewriting and optimizing PVM the performance could be improved.

The real-time video demonstration also utilized two Paragons. An ATM cell stream from the EMMI was routed via an OC-3 module to the first Paragon. That Paragon inverted the cell data bits, recalculated the header and trailer, and passed the data back to the network. The data was then routed to the second Paragon. The second Paragon repeated the operation of the first Paragon. The data was then routed back to an EMMI where it was displayed. The data rate for the video averaged 70 Mbits/sec. The Paragon was required to traffic shape the video data so as not to overrun the EMMI buffers. A live video camera allowed viewer to see their image making the loop trip between the Paragons. The latency of the loop is not discernible to the human eye.

Although data is scare on the ORNL material science applications, it ran consistent with expectations.

5 Sandia's Corporate Video Demonstration

The Corporate Video Extension demonstration utilized the IDREN to provide a video feed both to and from the Supercomputing Conference (see Figure ??) and Sandia in New Mexico. During the conference, special programming from *Supercomputing '95* was presented live on the SNL corporate video network in Albuquerque. Sandia's Video Services organization also placed a greeting to the remote viewers in the blue informational banner that they add to the corporate feed. The feed from the conference presented live demonstrations that were

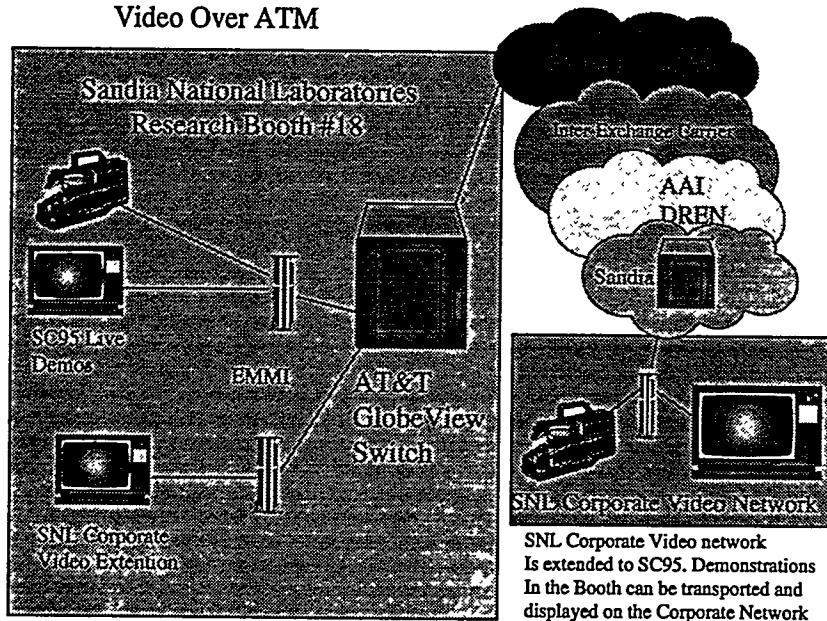


Figure 3: Sandia's *Supercomputing '95* Video Network

going on in our booth. The demos that were selected involved Sandia virtual reality projects. Although we were not able to secure a quality audio circuit back from the convention center, the Video Services' personnel put a musical filler into their transmission. The demonstrations from the live virtual reality was received back at our research booth on the corporate video extension display. The time delay between the transmission to Sandia in New Mexico and was less than noticeable to the viewers. The transmission delay was calculated as being near 100 milliseconds.

For the demonstration, the video feed from Sandia's Corporate Video Pilot Project video network was converted to a motion JPEG AAL5 ATM data stream by the AT&T EMMI. The audio feed from Sandia's Corporate Video Pilot project video head was converted to an AAL1 ATM data stream by the AT&T EMMI. The audio and video data streams were transported across two VCs to San Diego Convention Center. A similar video feed back from the Supercomputing Conference was also provisioned. The slowest link that we used on the demonstration was a DS-3 44.763 mbps telecommunication circuit. This circuit determined the quality factor that we could use on the demonstration. The quality factor that was selected for the corporate video extension transmis-

sion was IJG75 which specifies that a typical images will result in data rates of 18 mbps and a complex images results in a 33 mbps data rate. This quality factor results in a picture that is slightly less than VHS video quality, but the viewers felt that picture quality was very good. The bandwidth of the audio network was 2 mbps. A video conference was run in conjunction with the video extension. The video conference quality factor for the video conference was IJG50 or typical image bandwidth of 11 mbps, complex image bandwidth was 21 mbps. Audio bandwidth was 2 mbps. These two transmissions resulted in the average utilization of 33 mbps. The peak utilization for the demo was 89% of the DS-3 ATM bandwidth or approximately 36 mbps.

6 Lessons Learned

A number of lessons were learned along the way to implementing these networks. The most important included:

- forming mutually advantageous partnerships remains the key to success and must be done well in advance of the conference,
- a well developed plan should be presented at an early *SCI*net meeting to allow the committee time to react,
- the temptation to modify and expand the plan should be avoided as the conference draws near,
- applications should be planned that will utilize the network so that the effort and expense of building the network is justified for everyone concerned,
- Building a complex PVC based ATM network is a labor intensive activity and the PVCs should be mapped out prior to the show.

7 Conclusion

By all measures the conference proved successful for Sandia. The conference provided a forum for Sandia to feature a wide variety of state-of-the-art networking and communications technologies and associated applications. The demonstrations benchmarked the current state of the SONET and ATM technologies, both of importance to many Sandia initiatives, and the evolving partnership with ORNL. The success stories were the culmination of work accomplished by

many people both within and outside of Sandia. In order to meet the challenging goals of the state-of-the-art networks, many teams were formed that crossed corporate and organizational boundaries. The conference also provided an opportunity to identify future goals and plan joint activities. The teamwork amplified the accomplishments and achievements of all the participants. Similarly, the conference provided many Sandians an individual opportunity for professional growth, friendly competition, and professional association. Still, on another level, the conference challenged its participants to take stock of their individual projects and to focus them for the demonstration. In all these ways, Sandia benefited from its participation in *Supercomputing '95*.

8 Acknowledgements

The activities at *Supercomputing '95* was a team effort. Many people took the time to make the conference demonstrations successful. We would like to thank the following people and organizations for their contribution to this effort:

AT&T – Oscar Saurez, Robert Zahorak, and Arun Ghosh;
Giganet Inc. – David Follet, Maria Gutierrez, and Rich Prohaska;
GTE Customer Networks – Ron Wells;
Intel Scalable Systems Division – Peter Molnar, George Carr, Ken Lord, Robert Fugatt, Bob Moore, David Fair, and Johanna Lees;
Oakridge National Laboratory – Tom Dunigan, Buddy Bland, Philip Papadopoulos, and Lawrence MacIntyre;
Phillips Labs – Carlos Fernandez;
Sandia National Laboratories – R. M. Cahoon, Christine Yang, Don Morrow, Alden Jackson, John Naegele, Allen Robinson, Larry Tolendino, Len Stans, Luis Martinez, Mark Schaefer, Frank Castelluccio, George Yonek, and John Eldridge;
Sprint – Laura Klein.

References

- [1] Arthurine Breckenridge and Michael O. Vahle. An account of Sandia's research booth at *Supercomputing '92*: A collaborative effort in high-

performance computing and networking. Technical Report SAND 93-0224, Sandia National Laboratories, Albuquerque, New Mexico, 1993.

- [2] Joseph P Brenkosh. SMDS at Supercomputing '92. *MAN Newsletter*, 4(4), April 1993.
- [3] Joseph P Brenkosh. The BNS-2000 at Supercomputing '93. *AT&T Status Packet Newsletter*, 21:10, Spring 1994.
- [4] Carl F. Diegert. Gigabit-per-second networking a Supercomputing '91. Technical Report 1, Sandia National Laboratories, Albuquerque, New Mexico, March 1992.
- [5] Steven A. Gossage. Delivery of very high bandwidth with ATM switches and SONET. Technical Report SAND92-1295, Sandia National Laboratories, Albuquerque, New Mexico, October 1992.
- [6] Steven A. Gossage and Douglas J. Green. Sandia Labs selects BNS-2000 switch for their super computer network. *AT&T Status Packet Newsletter*, 17, Spring 1993.
- [7] Steven A. Gossage, John H. Naegle, and Nicholas Testi. Early experience with ATM in the scientific research environment. In *Broadband Networking Conference Proceedings*, page sec. 7, Washington, D. C., November 1993. 7th Annual Next Generation Networks Conference.
- [8] Steven A. Gossage and Michael O. Vahle. Sandia's research network for Supercomputing '93: A demonstration of advanced technologies for building high-performance networks. Technical Report SAND 93-3807, Sandia National Laboratories, Albuquerque, New Mexico, December 1993.
- [9] John H. Naegle, Nicholas Testi, Steven A. Gossage, Michael O. Vahle, and Joseph H. Maestas. Building networks for the wide and local areas using asynchronous transfer mode switches and synchronous optical network technology. *IEEE Journal on Selected Areas In Communications*, 13(4):662-672, May 1995.
- [10] Staff. Sandia links sites with high-capacity computer networks. *Sandia Science News*, 28(4), April 1993.
- [11] Staff. Supercomputing '95 final program. Sponsored by ACM SIGARCH and the IEEE Computer Society, December 1995.
- [12] Michael O. Vahle, Steven A. Gossage, and Joseph P. Brenkosh. Sandia's network for Supercomputing '94: Linking the Los Alamos, Lawrence Livermore, and Sandia National Laboratories using Switched Multimegabit Data Service. Technical Report SAND94-3096, Sandia National Laboratories, Albuquerque, New Mexico, December 1994.

Distribution:

1-0103	R. J. Detry, 12100	1-0970	S. G. Varnado, 9400
1-0318	A. Breckearidge, 9215	1-1090	W. D. Swartz, 4411
1-0318	G. S. Davidson, 9215	1-1109	A. L. Hale, 9224
1-0321	W. J. Camp, 9204	1-1111	K. S. McCurley, 9224
1-0322	P. J. Eicker, 9600	1-1375	S. S. Dosanjh, 9221
1-0451	M. R. Sjulin, 9417	1-1375	L. D. Bertholf, 4500
1-0451	T. D. Tarman, 9417	1-1427	W. F. Mason, 4021
1-0458	J. S. Rottler, 9003	1-9003	P. L. Mattern, 1100
1-0619	Print Media, 12615	1-9011	D. L. Crawford, 8900
1-0630	M. J. Eaton, 4010	1-9011	A. W. Jackson, 8910
1-0661	M. H. Pendley, 4612	1-9011	C. L. Yang, 89101
1-0801	M. J. Murphy, 4900	1-9011	F. T. Bielecki, 8930
1-0803	J. F. Jones, Jr., 4600	1-9011	H. Y. Chen, 8910
1-0806	D. C. Jones, 4603	1-9011	J. A. Friesen, 8901
1-0806	J. H. Naegle, 4616	1-9011	J. A. Hutchins, 8930
1-0806	J. P. Brenkosh, 4616	1-9011	J. C. Berry, 8910
1-0806	L. F. Tolendino, 4616	1-9011	J. Costa, 8920
1-0806	L. G. Martinez, 4616	1-9011	J. M. Brandt, 8910
1-0806	L. G. Pierson, 4616	1-9011	P. W. Dean, 8910
1-0806	L. Stans, 4616	1-9011	R. D. Gay, 8920
10-0806	M. O. Vahle, 4616	5-0899	R. E. Palmer, 8901
50-0806	S. A. Gossage, 4616	2-9018	Technical Library, 4414
1-0806	T. J. Pratt, 4616	2-0100	Central Technical Files, 8523
1-0807	B. C. Whittet, 4617		Document Processing for
1-0807	F. J. Castelluccio, 4617		DOE/OSTI, 7613-2
1-0807	G. A. Yonek, 4617		
1-0807	I. C. Alexander, 4617		
1-0807	J. M. Eidridge, 4617		
1-0807	M. A. Schaefer, 4617		
1-0807	R. L. Adams, 4617		
1-0807	R. M. Cahoon, 4918		
1-0807	S. D. Nelson, 4617		
1-0807	T. C. Hu, 4617		
1-0811	C. D. Brown, 13221		
1-0819	A. C. Robinson, 9231		
1-0819	J. M. McGlaun, 9231		
1-0820	P. Yarrington, 9232		
1-0826	J. D. Zepper, 9111		
1-0833	J. H. Biffle, 9103		
1-0841	P. J. Hommert, 9100		

