



Sun Microsystems, Inc.  
2550 Garcia Avenue  
Mountain View, CA 94043  
415 960-1300  
FAX 415 969-9131

For U.S. Sales Office locations, call:  
800 821-4643  
In California:  
800 821-4642

Australia: (02) 413 2666  
Belgium: 32-2-759 5925  
Canada: 416 477-6745  
Finland: 358-0-502 27 00  
France: (1) 30 67 50 00  
Germany: (0) 89-46 00 8-0  
Hong Kong: 852 802 4188  
Italy: 039 60551  
Japan: (03) 221-7021  
Korea: 822-563-8700  
Latin America: 415 688-9464  
The Netherlands: 033 501234  
New Zealand: (04) 499 2344  
Nordic Countries: +46 (0) 8 623 90 00  
PRC: 861-831-5568  
Singapore: 224 3388  
Spain: (91) 5551648  
Switzerland: (1) 825 71 11  
Taiwan: 2-514-0567  
UK: 0276 20444

Elsewhere in the world,  
call Corporate Headquarters:  
415 960-1300  
Intercontinental Sales: 415 688-9000



---

## *References*

1. “Asynchronous Transfer Mode: bandwidth for the Future” , Telco Systems Inc, Norwood, Massachusetts. This document contains an excellent description of what ATM is, and how it will impact network decision makers in the future.
2. “World Multimedia Hardware and Software Markets” Market Intelligence Research Corp.
3. “A brief tutorial on ATM (Draft)” , Zahir Ebrahim, Sun Microsystems February 13, 1992. This document gives an excellent overview of ATM and its place on the TCP/IP stack.
4. “SBus ATM Interface Specification”, Draft Document 1.0, Prakash Kashyap, Sun Microsystems, March 1993. This specification addresses the functionality of Sun’s SBus ATM Interface board, including issues regarding buffering, signalling, bit-rates etc.
5. “A General ATM-PHY Layer Interface Specification”, Version 0.95, Soren Christensen (Sun Microsystems) and Gaurav Garg (Synoptics Communications Inc)
6. “Unix Network Programming”, W.Richard Stevens, Prentice Hall, 1990..

## *Acknowledgements*

The author would like to thank Fred Sammartino for his valuable suggestions and use of his slides, Diane Daniels for her valuable suggestions and editorial contributions, Kevin Kalajan, Doug Ehrenreich, and others who have reviewed this article and offered suggestions.

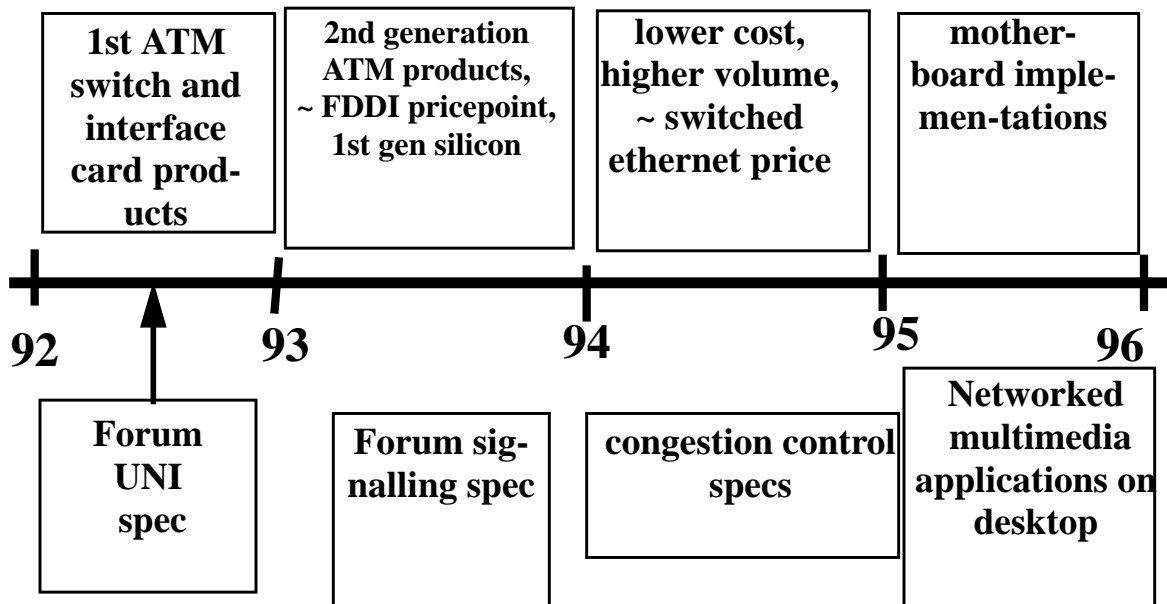
## *Author’s Profile*

Krish Narayan is a Member of Technical Staff at Sun Microsystems Inc., and has over six years of networking experience in his previous position as Manager of Network Support at Tymnet, where he was responsible for successfully installing the first Unix network in Network Operations, and for successfully developing Tymnet’s first Unix-based Network Design System. His current interests are in Networking and Multiprocessing.



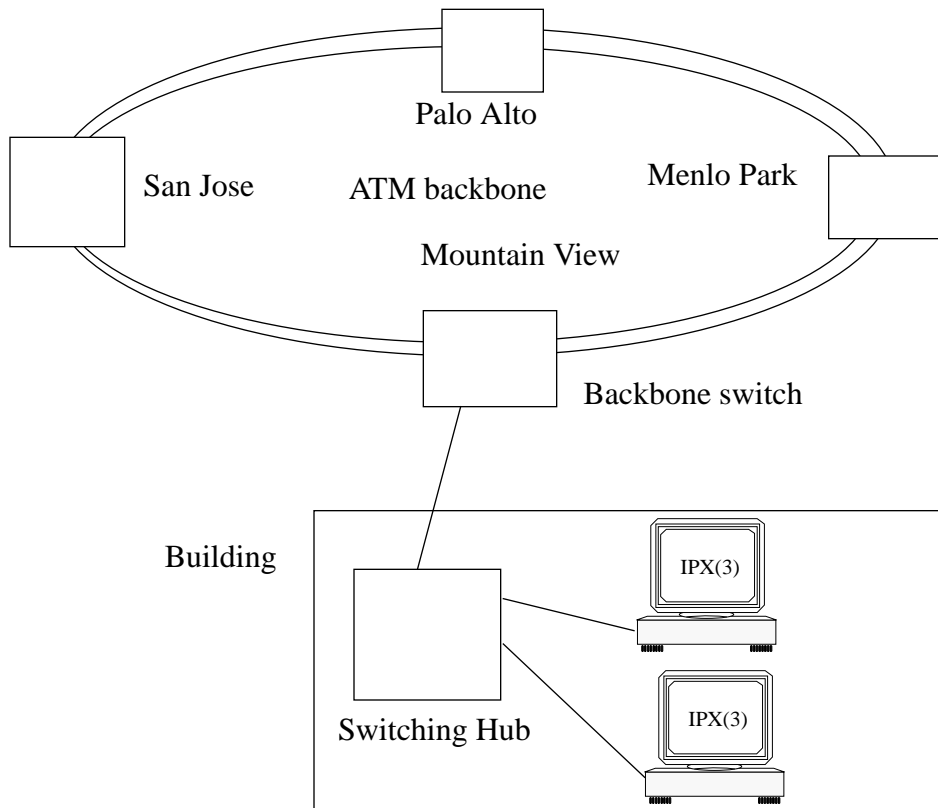
## ATM Futures at Sun

The figure below outlines the possible deployment schedule of Sun's ATM products:



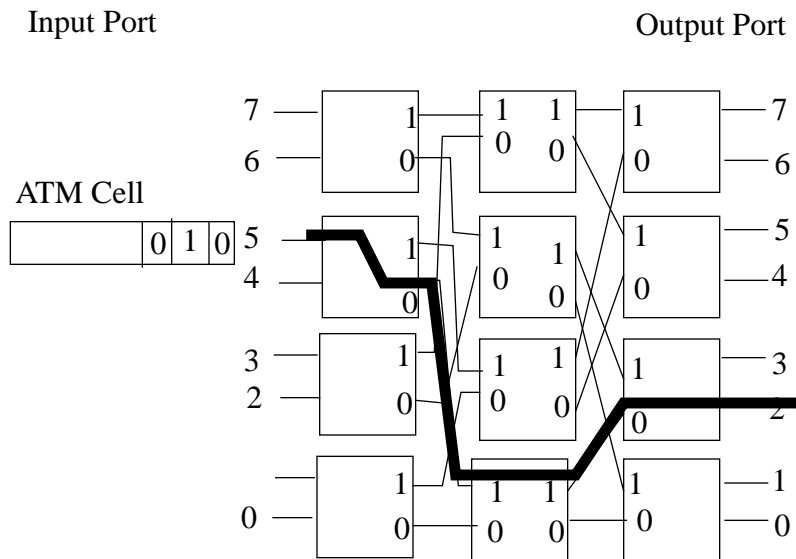
As we can see from the figure, networked multimedia applications are likely to become available in the 1995-1996 time-frame, which is only a couple of years away.

The ATM Forum, of which Sun is a member, is a non-profit organization chartered to accelerate the use of ATM products and services, and has a current membership of almost 200 organizations (this list is growing at a phenomenal rate). The current president of the ATM Forum is Fred Sammartino of Sun Microsystems Inc. The ATM Forum is in the process of defining standards in the areas of Signalling, Network-to-Network Interface, Congestion Control, Traffic Management, and the use of twisted-pair cable for ATM. Since twisted-pair cable installations are almost the rule in most Companies in the USA today, ATM over twisted-pair (when it becomes available) will accelerate it's deployment in many MIS shops.



The backbone network is a “redundant ring” configuration which may consist of at least 2 pairs of single-mode fiber links. One pair will be used for two-way traffic, and the other will provide redundancy. Backbone switch manufacturers include AT&T, Fore Systems, Ericsson, Siemens, Alcatel and several others. ATM Switching hub manufacturers include Synoptics, Newbridge and several others. The switching hubs would be connected to the SparcStations via the ATM Sbus interface card; this would take the place of a standard ethernet connection.

This configuration is only one of several possibilities, which makes ATM one of the more flexible networking solutions available today.



simple decision based on the tag value: 0 goes to the lower port and 1 goes to the upper one. Many other types of switches are available (see reference 1 for more info).

## *Unlimited Area Network and Flexible Network Topology*

Let us now take a look at what a typical ATM network may look like. One possible configuration is illustrated below:



obvious. Existing SNA, X.25 etc products could be ported/streamlined to fit into the new networking paradigm. Multimedia programmers can write applications as before, but now they do not have to worry about differences in LAN/WAN interfaces and physical layers, and the network latencies (delays) that result from them.

In fact, Sun's Multimedia software platform is currently being tailored to enable the creation of the following applications across the broadband network:

Video Conferencing

-Room to Room

- Desk to Desk

-Documents

-Images

-CAD/CAM Files

Remote attendance to talks

Multimedia email

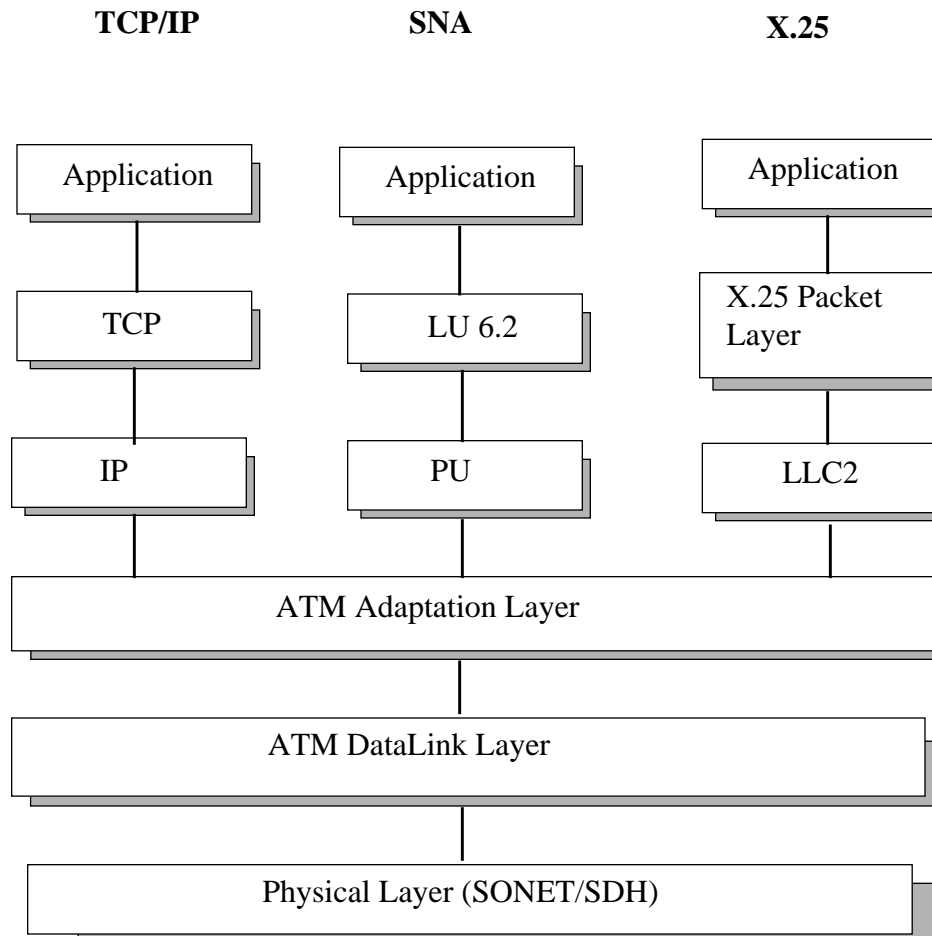
Collaborative work (groupware)

## *Simple Massive Hardware Switching*

Since the network switch's main function is to switch fixed length cells (with no protocol conversions taking place) we can expect to see a considerably different switching methodology than is currently available from VAN providers. ATM switches are built using self-routing procedures where the cell actually finds its own way through the switch fabric using the cell address, rather than having an external process establish and tear down a path.

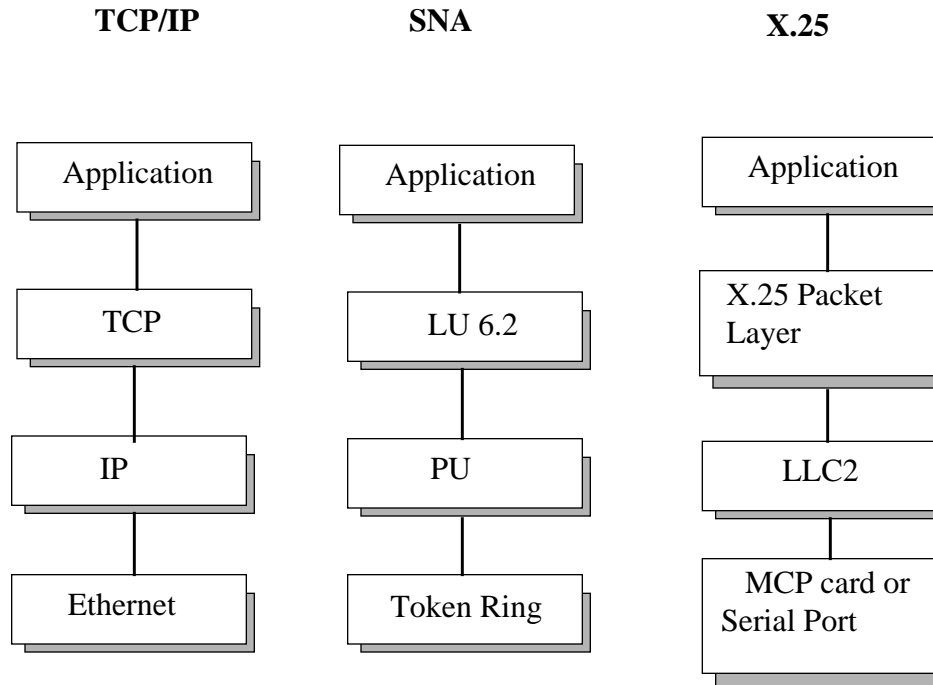
The figure below illustrates how such a switching process might work for a single ATM cell, shown entering port 5<sup>1</sup>. The cell is prepended with a tag 010., indicating its destination is port 2. The reader can demonstrate that no matter which port the cell arrives in, the tag "010" will always send it to port 2. This type of a switch is called a Banyan Switch (not to be confused with the Banyan Vines network), and is a simple two-input, two-output switch that makes a

1.From Reference 1, page 32



**Figure 3. Protocol Delivery Mechanisms**

Thus, an IP packet would get fragmented into ATM cells via the ATM adaptation layer, and the reassembly would be done by the ATM layers on the receiving sides. This would, of course, require that existing Sunlink products be re-written to talk to the ATM SBus card, but the advantages are fairly



**Figure 3. Protocol Delivery Mechanisms**

Experts will argue that each of these protocols can run over any of the available physical interfaces, but that really misses the point. The point really is that the proliferation of physical interface standards impose a considerable burden on the programmer, who must now ensure that their application is “interoperable” across several paradigms. Several implementations of the same code must be developed to ensure this interoperability, and maintenance across several platforms/protocol interfaces can be a nightmare (ask any MIS Manager, they know). And, this doesn’t even begin to address the separate requirements of video and voice, which have their own sets of protocols/standards/physical interfaces.

Let us now look at the situation with ATM, and one can see why it make sense to have a standard physical interfaces, as shown in Figure 4.



Table 4 Quality of Service Parameters

Parameter	Responsibility of
Cell Buffering	Host System
Protocol Integrity	Host System
Error Correction	Host System

The VAN provider is responsible for ensuring that sufficient throughput capacity is available in the backbone network to handle the application demands at an agreed-upon availability level, and that circuits are available for the duration of the session. It is the user's responsibility to ensure that ATM cells are delivered and received error-free, with adequate buffering/flow control as part of the overall congestion management philosophy.

Sun's ATM SBus interface card <sup>1</sup> has been designed to handle most of the user requirements listed in Table 4. Network interfaces (multi-mode fiber, Shielded Twisted Pair, etc) are provided by attaching "daughter cards" to the main SBus board. The SAHI card provides much of the needed functionality that will meet the market demands for high-bandwidth applications in the near future.

## Architectural Consistency

Achieving architectural consistency across LAN/WAN domains is impeded by the proliferation of delivery mechanisms in use by different protocols. For example, TCP/IP uses ethernet and FDDI as its main transport medium. X.25 uses HDLC, and the SNA protocol stack uses token ring (see Figure 3). The OSI and TCP/IP models <sup>2</sup> have been developed to allow for "interoperability" between vendor-specific hosts.

Sun has attempted to meet the market needs for interoperability by developing a suite of Sunlink products for a number of communication protocols (X.25, SNA, Frame Relay etc). As mentioned previously, each of these products has their own physical interface, as is illustrated below:

---

1. SAHI card, which is an abbreviation for SBus ATM Host Interface. See References 4 and 5

2. See Reference 6 for a detailed description of OSI and TCP/IP



DS3 range. Frame Relay attempts to fix this problem in the data networking domain by defining a parameter called the CIR (Committed Information Rate) which guarantees a user a specified rate for an agreed-upon duration.

However, the use of Frame Relay (which is further limited to data comms only) requires that users monitor their network usage very closely. The standards clearly do not impose mandatory requirements for users to implement congestion avoidance and recovery mechanisms. And the main weakness in implementing Frame Relay is in the looseness of these standards.

ATM is expected to provide true bandwidth on demand by implementing two functions: Call management and connection (or session) management. Call management is concerned with issues such as signalling (line provisioning and call setup), flow control, congestion avoidance and recovery. Connection management includes consideration of link utilizations, traffic, etc. and is a part of the "network monitoring" function. Connection management is expected to be handled through close integration of the "backbone" network management system and the "local" network management system, such as SunNet Manager, Sun's Network Management product. Efforts are currently underway to achieve this integration (for example with Synoptics, an ATM switch vendor).

Table 4 below illustrates the ATM quality-of-service (QOS) parameters.

*Table 4* Quality of Service Parameters

<b>Parameter</b>	<b>Responsibility of</b>
Throughput	VAN Provider
Transit Delay	VAN Provider
Switch MTBF	VAN Provider
Circuit Provision Time	VAN Provider
Information Integrity	VAN Provider
Congestion Management	Host System
ATM Cell SAR	Host System



little data would get lost, and in some cases could easily be recovered. Also, the fixed cell size would favor voice and video traffic, which need fixed-size cells for efficient, loss-free transmission and improved switch efficiency.

The 48 bytes of payload may optionally contain a 4 byte ATM adaptation layer and 44 bytes of actual data, or all 48 bytes may be data, based on a bit in the control field of the header. This enables fragmentation and reassembly of cells into larger packets at the source and destination respectively. The control field may also contain a bit to specify whether this is a flow control cell or an ordinary cell, an advisory bit to indicate whether this cell is droppable in the face of congestion in the network, etc.

In the words of Tom Lyon, Distinguished Engineer at Sun: “Sun is very interested in ATM as a general-purpose interconnect. It is the nature of general purpose things that, especially early in their life, they don’t look terribly appealing for any one role. But ultimately the volume attainable by general purpose things make their price/performance compelling for more and more roles. Also, doing more than one interface where one would suffice is just too damn expensive in today’s cutthroat computer industry. Consider the duck. It can walk (slowly), swim(OK), dive(sometimes), and even fly. Elegant, it is not. ATM is a duck with a several billion \$ world-wide R&D budget. Better birds are clearly possible, but incredibly unlikely to achieve the same momentum.”

## *Bandwidth on Demand and Quality of Service*

This section will examine how ATM is expected to meet the market requirements for bandwidth-on-demand, and highlight some quality of service issues.

True bandwidth-on-demand implies that users pay only for the bandwidth they use, and that sufficient bandwidth be made available (on a pre-arranged agreed-upon basis) whenever and wherever they need it.

Current networking paradigms (X.25, SNA etc) do not provide this mechanism. Rather, users pay for a fixed bandwidth connection from the WAN provider, and they are billed on a formula based on characters transmitted and purchased bandwidth. As users’ demand increases, they are obliged to purchase a large chunk (see Table 1) above their current rate, regardless of usage patterns. For example, table 1 shows poor granularity in the DS2 thru



Table 3 SONET Rates<sup>1</sup>

Nomenclature	Rate	Comments
STS-1/OC-1	51.84 Mbps	28 DS1s or 1 DS3
STS-3/OC-3	155.52 Mbps	3 STS-1s
STS-12/OC-12	622.08 Mbps	12 STS-1s
STS-48/OC-48	2488.32 Mbps	48 STS-1s

1. From Reference 1

Contrast these rates with the standard Ethernet rate of 10 Mbps, and we can quickly see the “several orders of magnitude” improvements in optical rates over their digital counterparts. These high rates are also associated with high costs, and ways must be found to reduce the cost/user to acceptable levels in order to be competitive with existing digital rates. One way to do this is to provide “bandwidth on demand” which will be discussed in more detail later.

The ATM cell, shown below, consists of a 5-bit header followed by a 48-bit information field.




Figure 2. ATM Cell

The two end points in an ATM network are associated with each other via an identifier called the "Virtual Circuit Identifier" (VCI label). The VCI is carried in the 5 Byte header portion of the cell.

Why a 53-bit cell? ATM was independently proposed by Bellcore, the research arm of the Regional Bell Operating Companies (RBOCs) in the US, and several giant telecommunications companies in Europe. The idea was to keep the size of the cell small so that if any one cell got dropped enroute due to congestion,



ATM provides increased networking flexibility, since the uniform cell format used in both LAN and WAN environments will greatly simplify interconnecting them.

Most importantly for Multimedia, ATM provides low latency. Real-time applications like desktop conferencing, and distance learning, which use voice and video simultaneously, are now feasible because ATM generates small cells that can be switched at high-rates, like circuit switching.

## What is ATM?

Many tutorials, documents and newsletters contain excellent descriptions of ATM technology<sup>1</sup>. Since not every reader may have access to these documents, we will present a summary of these papers here.

ATM is the complement of STM which stands for "Synchronous Transfer Mode". STM (also called Time Division Multiplexing in the literature) is used by telecommunication backbone networks to transfer packetized voice and data across long distances. Most TDM signals are the familiar electrical signals such as DS1, DS3 etc as shown below:

Table 2 TDM Rates<sup>1</sup>

Nomenclature	Rate	Comments
DS0	64 kbps	1 Voice Channel
DS1	1.544 Mbps	24 DS0s
DS2	6.312 Mbps	4 DS1s
DS3	44.736 Mbps	28 DS1s

1. From Reference 1

Asynchronous Transfer Mode traffic, on the other hand, can be carried on a synchronous optical network (SONET). Sonet rates are shown below:

---

1. See, for example, References 1, 3, and 4



Table 1 How Technology Must Address the Market

The Market Wants	Technology must Provide
Multimedia (integrated voice, data, video)	High bandwidth on demand
Easier LAN-MAN-WAN integration	Architectural Consistency
Economy of Networking	Simple, massive, hardware switching
Simplicity of Networking	Unlimited Area Network
Flexible network topology	Switched mesh rather than star

What benefits does ATM provide for Multimedia, over conventional packet switching protocols like X.25, SNA, Frame Relay etc? Briefly, they are in the areas of scalability, transparency, granularity, networking flexibility and low latency.

ATM is a scalable technology because it's fixed cell format (Section 2) is silent on rates, framing protocols, or physical bearers. This means that a cell generated by a 100 Mbps LAN can be carried over a 45 Mbps DS3 to a central office and switched into a 2.4 Gbps SONET transport system. This illustrates three very different systems (LAN, network transport, switching) with the message in the same format, rate scaled to suit the application. It is very difficult to scale this way with conventional packet-switched networks because the underlying delivery mechanism (ethernet, token ring, etc.) is different for each protocol on each system.

ATM is application transparent, because the cell size is a compromise between the long length packets needed by data communications applications and the short-length repetitive needs of voice and video. ATM will allow free mixture of data, voice and video within the same application without worrying about compatibility problems with LAN or WAN.

ATM provides increased network granularity by allowing for the network to be tailored to the application, rather than have the application be forced to fit into the underlying network limitations.



## *Overview*

Enterprise networking today stands at the threshold of a new era. Driven by a “technology push”, on the one hand, due to the availability of ever-increasing computing power on the desktop, and a “demand-pull” on the other, due to the increasing demand for network services from corporate entities spread across the globe, massive changes are taking place in MIS shops and corporate network planning centers worldwide. These changes are further exacerbated by the diverging networking architectures presented by existing WANs and LANs.

Multimedia traffic requirements (integrated voice, video and data) are primed to explode by almost 500 percent over the next decade. Existing WAN and LAN architectures are optimized separately for data, voice and video. For example, two of the largest data network providers in the USA today operate different networks for data and videoconferencing applications. Frame relay has been designed to efficiently handle bursty, large data packets on a global scale. Transmitting voice and video on a frame relay backbone would be very inefficient.

In this paper, we will describe a new networking technology called ATM (Asynchronous Transfer Mode). We will then proceed to examine how ATM promises to satisfy the market requirements (shown in Table 1 below). Finally, we will describe the current state of ATM product development at Sun Microsystems Inc.

## *The Benefits for Multimedia*

Table 1 shows a brief overview of what the market wants, and what the new networking technology must provide. The following terminology is used:

- LAN - Local Area Network
- MAN - Metropolitan Area Network
- WAN - Wide Area Network

© 1993 Sun Microsystems, Inc.  
2550 Garcia Avenue, Mountain View, California 94043-1100 U.S.A.

All rights reserved. This product and related documentation are protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or related documentation may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Portions of this product may be derived from the UNIX® and Berkeley 4.3 BSD systems, licensed from UNIX System Laboratories, Inc. and the University of California, respectively. Third-party font software in this product is protected by copyright and licensed from Sun's Font Suppliers.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the United States Government is subject to the restrictions set forth in DFARS 252.227-7013 (c)(1)(ii) and FAR 52.227-19.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

#### TRADEMARKS

Sun, Sun Microsystems, Sun Microsystems Computer Corporation, the Sun logo, the SMCC logo, are trademarks or registered trademarks of Sun Microsystems, Inc. UNIX and OPEN LOOK are registered trademarks of UNIX System Laboratories, Inc. All other product names mentioned herein are the trademarks of their respective owners.

All SPARC trademarks, including the SCD Compliant Logo, are trademarks or registered trademarks of SPARC International, Inc. SPARCstation, SPARCserver, SPARCengine, SPARCworks, and SPARCcompiler are licensed exclusively to Sun Microsystems, Inc. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK® and Sun™ Graphical User Interfaces were developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

X Window System is a trademark and product of the Massachusetts Institute of Technology.

THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS PUBLICATION COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THE PUBLICATION. SUN MICROSYSTEMS, INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS PUBLICATION AT ANY TIME.



Please  
Recycle

# *ATM Networking and Multimedia*

*A White Paper*

*Krishnan Lakshminarayan  
Market Development Engineering  
August 1993*



**Sun Microsystems Computer  
Corporation**

2550 Garcia Avenue

Part No.: 800-999-0001  
Revision X, August 3, 1993