

SANDIA REPORT

SAND97-0642 • UC-705

Unlimited Release

Printed March 1997

Windows NT 4.0 Asynchronous Transfer Mode Network Interface Card Performance

Lawrence F. Tolentino

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia
Corporation, a Lockheed Martin Company, for the United States
Department of Energy under Contract DE-AC04-94AL85000.

Approved for public release; distribution is unlimited.



Sandia National Laboratories

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
P.O. Box 62
Oak Ridge, TN 37831

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

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

NTIS price codes
Printed copy: A03
Microfiche copy: A01

SAND97-0642
Unlimited Release
Printed March 1997

Distribution
Category UC-705

Windows NT 4.0 Asynchronous Transfer Mode Network Interface Card Performance

Lawrence F. Tolendino

Advanced Networking Integration Department

Sandia National Laboratories

P. O. Box 5800

Albuquerque, NM 87185-0806

February 18, 1997

Abstract

Windows NT desktop and server systems are becoming increasingly important to Sandia. These systems are capable of network performance considerably in excess of the 10 Mbps Ethernet data rate. As alternatives to conventional Ethernet, 155 Mbps Asynchronous Transfer Mode, ATM, and 100 Mbps Ethernet network interface cards were tested and compared to conventional 10 Mbps Ethernet cards in a typical Windows NT system. The results of the tests were analyzed and compared to show the advantages of the alternative technologies. Both 155 Mbps ATM and 100 Mbps Ethernet offer significant performance improvements over conventional 10 Mbps shared media Ethernet.

Acknowledgments

Thanks to John Naegle, and Steven Gossage, both 4616, for reviewing drafts of this document and making many helpful suggestions. Thanks also go to Michael Vahle, 4616, for suggesting the development of a metric based on CPU utilization and network performance (the figure of merit M) for the direct comparison of network interface cards. Special thanks also go to Erik Tolendino for running the Adaptec network interface card performance tests during "Take Your Sons to Work Day".

Contents

INTRODUCTION.....	4
SOME SANDIA STRATEGIC DIRECTIONS	4
ATM TESTBED.....	5
TESTING WINDOWS NT 4.0 WITH ATM NETWORK INTERFACE CARDS	6
NETWORK INTERFACE CARD FIGURE OF MERIT	8
ATM 155 MBPS VS. ETHERNET 100 BASE-T	9
CONCLUSIONS	11
APPENDIX A - ATM THROUGHPUT TEST PROTOCOL	14
TTCF TEST	14
FTP - FILE TRANSFER PROTOCOL TEST	15
APPENDIX B THE ATM TESTBED	17
APPENDIX C - THE IMPACT OF TCP WINDOW SIZE ON ATM PERFORMANCE	18

Figures

Figure 1. Performance of ATM Network Interface Cards.....	8
Figure 2. The Figure of Merit for ATM Network Interface Cards Compared to an Ethernet Network Interface Card	9
Figure 3. Performance of ATM and 100 Base-T Ethernet Network Interface Cards.....	10
Figure 4. Figures of Merit Compared for ATM and 100 Base-T Ethernet Network Interface Cards.....	11
Figure 5. Advanced Networking Integration Department ATM Testbed.....	17
Figure 6. The Current LANE Packet Size of 1500 Bytes Limits ATM Performance.....	18

Introduction

The Advanced Networking Integration Department at Sandia National Laboratories is responsible for evaluating emerging networking and communications technologies with respect to corporate requirements. Technologies that meet Sandia's needs are then integrated into the Sandia infrastructure. As part of this evaluation, responsibility the department has investigated Asynchronous Transfer Mode (ATM) technology for over 5 years and has determined several benefits that ATM could provide our corporate communications infrastructure. In order to evaluate and develop ATM technologies for deployment throughout Sandia's communications infrastructure, an ATM testbed has been operational since the first ATM equipment was available. As the technology matured, ATM effectively demonstrated capabilities beyond that of legacy network technologies in the Local Area Network (LAN) and Wide Area Network (WAN) environments. Several publications report the results of these demonstrations.^{1,2,3} In addition, the recently formed Corporate Information Officer (CIO) organization began the definition of an Integrated Information System, IIS, for the laboratories. The definition of the IIS caused attention to be focused on the Windows NT desktop and server operating systems.

Some Sandia Strategic Directions

Sandia National Laboratories charged the CIO organization with developing and implementing a Corporate Information Plan as well as the IIS. As part of those efforts, the CIO defined certain standards with regard to desktop operating systems and workstations. One of the first definitions for desktop systems included the Microsoft Windows NT operating system and Intel Pentium and Pentium Pro based desktop workstations. In addition, UNIX and Windows based network servers were chosen to provide critical network services such as electronic mail

¹ Gossage, S. A., "Delivery of Very High Bandwidth with ATM Switches and SONET", SAND92-1295, November 1992.

² Naegle, J. H., Testi, N., and Gossage, S. A., "Developing an ATM Network at Sandia National Laboratories", Telecommunications, Vol. 28, No. 2, February 1994, pp. 21-23.

(E-mail), Domain Name Services (DNS), file services and print services. Initial testing of Intel Pentium based systems running the Microsoft Windows NT operating system indicated that conventional shared media based Ethernet Local Area Networks, LANs, could seriously hamper the performance of these systems. Indeed, first looks indicated that these systems would easily exceed the 10Mbps clock rate of standard Ethernet.

Concurrent with the definition of desktop and server platforms, the CIO organization was defining an internal corporate data network, INTRANET, to provide corporate information to any customer, in any location, at any time. World Wide Web technology (including the TCP/IP network protocol, the HTTP protocol, and the HTML document language) forms the basis for the Sandia INTRANET. The Sandia INTRANET includes significant graphical content and demands quick response times. Such requirements place a very high load on the network infrastructure and place a premium on high throughput and minimal latency as well as powerful desktop workstations. These developments demonstrated the need to investigate alternatives to ETHERNET and experience pointed in the direction of ATM.

ATM Testbed

The Advanced Networking Integration Department has evolved an extensive ATM testbed over the past five years. Using this testbed, ATM network components such as switches and routers as well as Network Interface Cards (NIC) were evaluated. In addition, the testbed provided the means to evaluate the performance of servers and workstations in an ATM environment. The work accomplished by the testbed allowed Sandia to commit to installing an ATM based data communications backbone infrastructure in 1995. This backbone built upon the Synchronous Optical NET (SONET) transport network that was installed beginning in 1991. The current testbed consists of ATM switches, ATM/Ethernet switches, ATM connected routers and ATM network test gear. (See Appendix B)

³ Naegle, J. H., *et. al.*, "Building Networks for the Wide and Local Areas Using ATM Switches and SONET," IEEE J. Select. Areas Commun., vol. 13, no. 4, pp. 662-672.

The testbed has evolved into an ATM network that supports 10's of desktop workstations and PCs used to perform day-to-day work for the department. This evolution has allowed the staff to evaluate the performance of ATM in typical office and engineering environments. At this point, the network has shown the stability to perform useful work and provided the opportunity to accurately test client and server systems. Given the Sandia strategic directions and a stable, reliable testbed it was decided to test the performance of ATM based Network Interface Cards. The cards were tested running a commercial operating system, Windows NT, in typical office or research laboratory applications.

Testing Windows NT 4.0 with ATM Network Interface Cards

A Gateway 2000 PC based on the Intel 150 MHz Pentium Pro CPU equipped with a PCI bus and standard SCSI disk subsystem was chosen to test. Windows NT version 4.0 was installed on the system. In order to benchmark the system, a SUN Sparc20 running the Solaris 2.51 operating system was selected to function as a network server. Both systems were connected to the ATM switches in the testbed network and configured to use LAN Emulation networking or LANE, a typical and common use of ATM. While the majority of Sandia's scientific and engineering customers use TCP/IP based networking, LANE provides a way to connect with Sandia's legacy Ethernet networks, via ATM regardless, of the protocol. For the scientific and engineering customers classical IP over ATM would have provided better performance but would have prevented other customers from taking advantage of ATM.

It was recognized that the use of LANE and its attendant 1,500 byte packet size would limit the performance of the Windows NT desktop. However, the goal was to determine performance in a typical installation rather than maximum possible performance. (The LANE standard does support larger packet sizes but no vendors are currently supporting packet sizes greater than 1,500 bytes; therefore, the throughput for certain performance tests is reduced by a factor of two or three because of the relatively small packet size.)

The determination to ascertain performance in typical real world settings also led to choosing two performance tests; TTCP and FTP. TTCP utilizes the complete TCP/IP protocol stacks in the client and server systems to move data between computer memories. Such activity is typical of an application where a compute server generates images to be viewed on a user's desktop. The second test, FTP, utilized the File Transfer Protocol to transfer a large file to or from the desktop client system. File transfer is a common desktop operation that is used in many applications such as WEB browsing. Three network interfaces were tested and are reported in this paper; the Interphase 5515 ATM NIC, the Adaptec ATM ANA5240 NIC and the 3Com 509 10 Mbps Ethernet NIC. The Ethernet card was included in the test as a standard for comparison because it represents the most common interface installed at Sandia. Each test was conducted through switched virtual circuits. No shared media network circuits were involved; therefore, the impact of other network traffic on measured throughput was minimized.

Cautionary Note: The tests reported in this paper should not be construed as indicative of the absolute throughput that could be obtained on other PC compatible systems running Windows NT 4.0. Rather the tests indicate the relative performance and CPU loading that can be expected when running Windows NT 4.0 using the network interfaces tested. As always, your mileage may vary.

Each card and associated device driver was tested in the Gateway PC using the TTCP and FTP tests, transmitting and receiving data with the SUN server. The complete test protocol is documented in Appendix A. Each test was run four times and data throughput and CPU percentage utilization were recorded. The importance of recording data throughput is obvious while CPU utilization is less so. It is desirable to evaluate the CPU utilization associated with data transfer because driver efficiency affects the overall performance of the desktop systems. Efficient well-written drivers are certainly desirable and are preferred over those that are less efficient. Figure 1 records the test results averages. While average test results are reported, it is interesting to note that test results varied by less than 5% for each test.

G6-150 Windows NT 4.0 Workstation NIC Performance								
Test	TTCP Receive		TTCP Transmit		FTP Receive		FTP Transmit	
	Data Throughput Mbps	CPU Utilization %						
NIC								
Interphase ATM 5515, driver of 9/12/96	20.65	100	45.91	79	15.5	100	30.1	92
Adaptec ATM ANA5240, driver 4.0	27.7	89	35.8	66	21.8	98	28.4	75
3Com509 ISA Ethernet	7.9	74	6.5	47	7.8	82	8	57

Figure 1. Performance of ATM Network Interface Cards

As one can see from the results in Figure 1, the ATM connected client system outperforms the Ethernet connected system by a wide margin. However, the performance comes at the expense of increased CPU utilization. Given this situation it is not obvious which NIC and associated driver is "best." In order to provide an easier way to compare these interfaces, a figure of merit was defined.

Network Interface Card Figure of Merit

In order to compare network interface cards a figure of merit, *M*, a dimensionless number, was defined. *M* was defined as the measured throughput in Megabits per second divided by the measured CPU percentage utilization represented as a decimal fraction,

$$M = \text{Mbps} / (\% \text{CPU utilization}).$$

This figure of merit was chosen as a way to demonstrate a benefit-cost ratio for each type of network interface card. The benefit is data throughput delivered while the cost is the amount of the CPU capacity consumed. Therefore, the larger *M* grows the greater the benefit per unit cost. For example, a throughput of 10Mbps that required 100 % CPU utilization would have a

figure of merit of 10; whereas, a throughput of 7.5 Mbps that required only 50% CPU utilization would have a figure of merit of 15. Figure 2 plots the figures of merit for the tests performed.

Windows NT 4.0 Network Interface Card Performance




Figure 2. The Figure of Merit for ATM Network Interface Cards Compared to an Ethernet Network Interface Card

ATM 155 Mbps VS. Ethernet 100 Base-T

With the advent of 100 Base-T Ethernet (100 Mbps), another alternative has been provided to those who wish to provide data rates higher than 10 Mbps to the desktop. The 100 Base-T Network Interface Cards, NICs, represent a very attractive alternative to the ATM NICs from a cost versus performance viewpoint. Therefore, the two alternatives of 100 Mbps Ethernet and 155 Mbps ATM have been compared on a typical PC platform. Unfortunately, the Gateway G5-150 Pentium Pro system that was used for the former tests was no longer available for testing so a Gateway 2000 G6-200 Pentium Pro system was substituted. Using a different NT test platform precluded direct performance comparisons with the previous tests. (When a long term

test platform becomes available, the tests will be rerun with an expanded set of NICs providing direct performance comparisons.

As before the PC platform was running Windows NT 4.0 and was tested with a SUN SPARC 20, running the Solaris operating system, version 2.51. The Sparc20 was connected to the ATM testbed via fiber optic cable and a 155 Mbps ATM NIC. Two NIC cards were tested in the NT system, the 3COM595 PCI bus 10/100 Mbps Ethernet NIC and the Interphase 5515 155 Mbps ATM NIC. The Ethernet card was connected to the testbed via twisted pair cable and the Cisco Catalyst 5000 Ethernet switch while the Interphase ATM card connected to the testbed via fiber optic cable and the Cisco LS1010 ATM switch. The actual throughput from the tests was recorded independently from the CPU utilization. The process of measuring the CPU utilization actually lowers throughput by several percent.

All results are reported from the perspective of the Windows NT system and are recorded as average values in the following illustrations. To provide a direct comparison of the two NIC alternatives the figure of merit, M, for each test was calculated in the same fashion as was done for previous tests.

G6-200 Windows NT 4.0 Workstation NIC Performance								
Test	TTCIP Receive		TTCIP Transmit		FTP Receive		FTP Transmit	
	Data Throughput Mbps	CPU Utilization %						
NIC								
Interphase ATM 5515, driver of 9/12/96	47.9	92	53.7	57	30.6	100	41.3	61
3COM595 100 Base-T Ethernet	44.4	75	46.2	58	26.2	100	38.4	56

Figure 3. Performance of ATM and 100 Base-T Ethernet Network Interface Cards

100 Mbps Ethernet Versus 155 Mbps ATM




Figure 4. Figures of Merit Compared for ATM and 100 Base-T Ethernet Network Interface Cards

Conclusions

One may observe, from the results of the two sets of tests, that receiving data requires more CPU resources than transmitting data. It is hypothesized that this difference in performance is due to at least two factors. First, receiving data is an interrupt driven process, requiring context switching, which increases the load on the CPU. Second, data can be received over several logical “ports”; therefore, a table search is required to route the data to the correct process.

One may also observe that the Ethernet drivers appeared to be more efficient with regard to CPU utilization. This may be indicative of the experience manufacturers have with writing Ethernet drivers.

There are several significant conclusions that can be drawn from these tests. The first is that the performance improvement provided by an ATM LAN connection over a conventional 10 Mbps Ethernet connection is significant. Not only can a significant performance improvement be measured in throughput, but it can also be "felt" by the user. This feeling comes through improved responsiveness when doing such activities as WEB browsing or accessing remote files.

The second conclusion is that, for data communications based applications, the 100 Mbps Ethernet NIC performs as well as the 155 Mbps ATM NIC on a typical desktop system. Given the perceived cost advantages of switched 100 Mbps Ethernet (100 Base-T), 100 Mbps Ethernet is very attractive. However, ATM throughput is artificially constrained when packet size is limited to 1,500 bytes to provide compatibility with existing LANs. One may expect there to be a significant difference between ATM and 100 Base-T Ethernet when the packet size for ATM is increased to approximately 10,000 bytes. It is not apparent that the packet size for 100 Base-T Ethernet can be increased at all. One should be cautious comparing the Ethernet and ATM technologies solely on the basis of throughput performance. Other factors must be considered when one chooses a technology. These factors include: maintainability, reliability, suitability for voice, video and data services and the ability to scale from local area network to wide area network environments.

The third conclusion is that the "standard" desktop chosen by Sandia is seriously constrained by a 10 Mbps Ethernet local area network connection. Ideally, no new networks would be designed and installed at Sandia that were based on the 10 Mbps shared media Ethernet model. In the near term (<3 years), the minimum that Sandia should strive for would be 10 Mbps switched access for all new local area networks. Such a design may grow to provide either switched 100 Mbps Ethernet or 155Mbps ATM as the customers demand.

Finally, one may conclude that Sandia must design all future local area networks based on a switched architecture. The switched architecture must be able to provide throughput of 100 Mbps or more. The fundamental limitation that held the ATM performance to approximately 40 Mbps was the requirement to be compatible with the legacy Ethernet LANs and their 1,500 byte

maximum packet size. Please see Appendix C for an illustration of the impact of packet size on performance. Techniques exist today to increase that packet size to approximately 10,000 bytes. It has been shown, through tests on the Sandia High Performance Storage System, that an increase in packet size would result in test throughputs of over 100 Mbps.

Appendix A - ATM Throughput Test Protocol

Each Network Interface Card and associated driver were tested with the Test TCP (TTCP) test and File Transfer Protocol (FTP) test. These two tests were chosen because they represent typical operations customers are expected to perform. The TTCP test represents memory-to-memory data transfer that might be expected between a computational server and client desktop. The FTP test represents a file transfer between two client systems or a client system and file server.

TTCP TEST

Each interface card and driver combination was evaluated using the TTCP memory-to-memory test. This test evaluates network throughput utilizing the complete operating system including the TCP/IP protocol stack; however, there is no disk I/O or file system support associated with the test. Two particular implementations of TTCP were used for the tests. TTCNT software, TTCP compiled for Windows NT, was used on the Windows NT system while GTTCP software, a Sandia modified version of TTCP, was used on the SUN workstation.

TTCP test parameters used represent the most common settings

n=2000, source buffers written to the network

b=32768, TCP socket buffer size

l=8192, length of buffers read or written to the network

s, discard all data

TTCP Test Procedure

The TTCP tests were controlled from the Windows NT system using Telnet to log on to the SUN Sparc20 running GTTCP and a DOS command window to control the TTCNT software on the NT workstation. Throughput was recorded from the GTTCP performance report as the GTTCP timing is more precise. The following typical test measures TTCP

throughput when the SUN Sparc20 receives data transmitted from the Windows NT workstation.

1. Open a DOS command window, change to the `ttcp` directory
2. Telnet to the SUN, change to the `/opt/metrics` directory
3. type `gttcp -r -s -b32768 -l8192` in the Telnet window to set up the receive
4. type `ttcpnt -t -s -b32768 -l 8192 -n2000 134.253.4.206 (or .71)` in the DOS command window to start the test
5. record the throughput reported by the GTTCP program

The above test was run four times and an average throughput calculated. Then the test was reversed with the Windows NT workstation receiving the data transmitted from the SUN workstation. This receive test was also run four times and the results averaged. Several tests were also run with the Windows NT CPU performance monitor turned on in order to measure CPU utilization. The data throughputs were not recorded when measuring CPU utilization on the Windows NT system because measuring CPU utilization affected data throughput.

FTP - File Transfer Protocol test

The FTP test used a standard file, `flute2.wav`, an 11 Mbyte file, to measure system performance. Using this file size allowed the SUN Sparc20 and the Windows NT systems to cache the file eliminating (to a large extent) any effects of the disk subsystems. All file transfers were controlled from the NT workstation and the notion of transmit and receive are used with respect to the Windows NT system.

FTP Receive Test - The Windows NT workstation receives files from the SUN workstation.

1. Open a DOS command window on the NT workstation
2. type `cd \test` to access the test file area

3. type **ftp 134.253.4.206 (or .71)** to use the Windows NT FTP application to connect to the SUN workstation
4. type **cd /opt/netnscape/data/samples/audio** to access the test files directory
5. type **bin** to use binary mode data transfer
6. type **get flute2.wav** to receive the 11 Mbyte file
7. record the data throughput reported by FTP

(repeat steps 6 to 7 four times and average the results)

FTP Transmit Test - The Windows NT workstation transmits files to the SUN system.

This test is run identically to the receive test except that **put** is typed instead of **get** at step #6.

The FTP tests were also run with the Windows NT CPU monitor active but the throughput results are not included in the report. Only the Windows NT CPU utilization is recorded for those tests because measuring CPU utilization reduces performance.

Appendix B The ATM Testbed




Figure 5. Advanced Networking Integration Department ATM Testbed

Appendix C - The Impact of TCP Window Size on ATM Performance

The constraint that ATM networking “Look like a legacy LAN” exacts a high performance penalty. Experience has led one to believe that as one increases the TCP window size network throughput would increase. However, the current LAN emulation, LANE, packet size of 1,500 bytes lowers anticipate performance. The following graph shows the performance of an early model Pentium system using an ATM network interface card and running Windows NT. The Test TCP program was used to transfer data from the Pentium system to an ATM connected SUN Sparc20 UNIX workstation. The systems used LANE to transfer data. The effect of the 1,500 byte legacy Ethernet packet limit shows up dramatically in Figure 6 as the TCP window size was increased.




Figure 6. The Current LANE Packet Size of 1500 Bytes Limits ATM Performance

The dip in performance at 1,500 bytes is evident and can be explained as the segmentation necessary to send large TCP data buffers via 1,500 byte packets.

Distribution:

1	MS0449	P. E. Sholander, 6512
1	0449	M. R. Sjulín, 6512
1	0630	M. J. Eaton, 4010
1	0661	M. H. Pendley, 4612
1	0803	J. F. Jones, 4600
1	0805	M. M. Miller, 4911
1	0806	C. D. Brown, 4621
1	0806	J. M. Eldridge, 4616
1	0806	D. E. Ellis, 4621
5	0806	S. A. Gossage, 4616
1	0806	P. S. Hamilton, 4616
1	0806	T. C. Hu, 4616
1	0806	L. G. Martínez, 4616
5	0806	J. H. Naegle, 4616
1	0806	T. J. Pratt, 4616
1	0806	D. H. Renaud, 4616
1	0806	L. Stans, 4616
15	0806	L. F. Tolendino, 4616
1	0807	I. A. Alexander, 4417
1	0807	T. V. Archuleta, 4417
1	0807	D. P. Evans, 4417
1	0807	J. H. Maestas, 4417
1	0807	S. D. Nelson, 4417
1	0807	B. C. Whittet
1	0812	C. A. Harris, 4914
1	9003	D. L. Crawford, 8900
1	9011	P. W. Dean, 8910
1	9011	R. E. Palmer, 8901
1	9012	J. E. Costa, 8920
1	9018	Central Technical Files, 8940-2
5	0899	Technical Library, 4414
2	0619	Review & Approval Desk, 12690 for DOE/OSTI,