

# Benchmarking a Network of PCs Running Parallel Applications<sup>1</sup>

Jeffrey K. Hollingsworth

Erol Guven

Cuneyt Akinlar

Dept. of Computer Science  
University of Maryland  
{hollings, eguven, akinlar}@cs.umd.edu

## Abstract

*We present a benchmarking study that compares the performance of a network of four PCs connected by 100 Mb/s Fast Ethernet running three different system software configurations: TCP/IP on Windows NT, TCP/IP on Linux and a light weight message passing protocol (U-Net Active messages) on Linux. For each configuration, we report results for communication micro-benchmarks and the NAS parallel benchmarks. For the NAS benchmarks, the overall running time using Linux TCP/IP was 12 to 500 percent less than the Windows NT TCP/IP configuration. Likewise, the Linux U-Net based message passing protocol outperformed the Linux TCP/IP version by 5 to over 200 percent. We also show that using Linux U-Net we are able to achieve 125 micro-second latency between two processes using PVM. Finally, we report that the default math libraries supplied with NT (for both gcc and Visual C++) are substantially slower than the one supplied with Linux.*

## 1. Introduction

The low cost of commodity hardware and software has encouraged their use for parallel computation. However, the high overhead associated with message passing and other system services in such environments often limits their performance. In this paper, we present performance measurements of a network of PCs connected with 100Mb/s Fast Ethernet running the Numerical Aerodynamics Simulation (NAS) parallel benchmarks[1] developed by NASA. Our goal in this paper is to assess the impact of different low-level communication systems and operating system services on the performance of parallel applications. A secondary goal is to identify specific bottlenecks in system software that limit parallel application performance.

All of the parallel applications used the Parallel Virtual Machine (PVM) message passing library[5] developed at Oak Ridge National Laboratories. The benchmarks were run on Windows NT using TCP/IP, and on Linux using TCP/IP and U-Net Active messages[13], a low latency

message passing protocol. In order to better understand and interpret the application results, we also tested the performance of the TCP/IP stack using a simple ping-pong benchmark. We used the Pentium processor's hardware counters[8] to isolate performance differences between these two platforms.

To pinpoint the differences between the two operating systems, we also investigated the performance of a pure computation kernel extracted from one of the NAS benchmarks. We used this kernel to compare the performance of the memory hierarchy, compiler, and math libraries. To obtain a detailed understanding of this benchmark, we used the Pentium hardware counters to measure floating point operations, cache activity, and instructions executed.

The organization of this paper is as follows: Section 2 describes our experimental configurations; Section 3 presents our micro-benchmarks of communication performance; Section 4 reports the results for the NAS benchmarks; Section 5 describes our evaluation of the compiler and memory hierarchy; Section 6 reports our experience with math libraries; and Section 7 compares our results to previous work. Finally, Section 8 summarizes the conclusions drawn from our experiments.

## 2. Experimental Configuration

Our experimental configuration consisted of four nodes. Each node contained a 120 MHz Pentium processor, 64 megabytes of memory, a 10 Mb/sec Ethernet card, and a SMC EtherPower 100 Mb/sec Fast Ethernet card (DEC 21140 chipset). The 10 Mb/sec Ethernet card is needed to run PVM with U-Net because U-Net cannot share the network interface with the kernel TCP/IP stack and we use the TCP/IP stack for PVM's process management functions such as *pvm\_spawn*. Thus, in experiments involving U-Net, PVM uses 10Mbps/s Ethernet card through kernel TCP/IP and U-Net uses Fast Ethernet card with its own device driver.

The benchmarks were run under three different software configurations:

- Windows NT 4.0 with PVM using the standard TCP/IP stack and the Fast Ethernet card.

<sup>1</sup> Supported in part by NSF Grant ASC-9703212, and DOE Grant DE-FG02-93ER25176.

- Linux 2.0.1 with PVM using the standard TCP/IP stack and the Fast Ethernet card.
- Linux 2.0.1 with PVM using U-Net Active messages over Fast Ethernet card (described below).

U-Net is a protected, user-level network architecture for low-latency and high-bandwidth communication. The U-Net driver provides best effort delivery of data between two processes. To run applications, we implemented the basic PVM message passing operations (point-to-point and multicast) as a library on top of the Generic Active Message[12] library supplied with the U-Net driver. This allowed us to run unmodified PVM applications in all of our test configurations. To provide source compatibility we replicated PVM's pack and un-pack interface at the cost of an extra data copy per message send.

### 3. Communication Micro Benchmark


In this section, we present the results of a simple ping-pong benchmark to measure the round-trip time of a message. Three versions of the same benchmark were implemented: One using core Active Messages; one using PVM message passing library over TCP/IP; and one using our implementation of PVM message passing library over U-Net Active Messages.

We define the round-trip time as the time required for a message to be sent by the sender, echoed by the receiver, and received again by the sender. In order to reduce sampling errors, the message is sent back and forth 3,000 times and the average of the last 2,950 of these round-trip times is taken.<sup>1</sup> Since each message is sent twice, the bandwidth is computed as half the round-trip time. Figure 1(a) shows bandwidth (MB/s) versus message size achieved by PVM-NT, PVM-Linux, PVM over U-Net Active messages and core U-Net Active messages.


The graph in Figure 1(a) shows the micro benchmark results when direct routing is used with PVM. Linux/TCP incurs a cost of 527 microseconds for a four byte message; whereas the same messages take 151 microseconds for PVM/U-Net. As the message size is increased, the high per message cost of the TCP version is amortized over a larger message and the performance gap between the two versions narrows, but is still over 40%. The four byte message NT/TCP performance was 1.3 milliseconds which is a factor of 2.4 slower than Linux/TCP and over 10 times slower than U-Net.

To quantify how much of the performance difference between Windows NT and Linux is due to their respective TCP/IP implementations, we repeated our micro-benchmark study using raw TCP/IP rather than PVM. We used a simple ping-pong benchmark to measure round-trip latency and bandwidth between two nodes on an otherwise idle Ethernet segment. The graph Figure 1(b) shows the

bandwidth achieved by PVM and TCP/IP on two operating systems. The results clearly show that latency in the Linux implementation of TCP/IP protocol is much less than Windows NT implementation.



(a)



(b)

Figure 1: Bandwidth vs. Message Size.

As shown Figure 1, bandwidth in Windows NT TCP/IP is half of that in Linux for small messages. This cost is primarily due to start-up cost in Windows NT TCP/IP stack. As the message size is increased, bandwidth becomes the bottleneck of communication. Although the difference in round-trip latency of two systems becomes smaller as the message size is increased, for 10KB messages Windows NT it is still 60% slower than Linux.

We performed a black-box analysis of the NT kernel versus the Linux kernel to see if we could infer the source of the differences in TCP performance. To do this, we used the Pentium hardware counters to measure the behavior of the memory hierarchy and instructions executed on each platform. A summary of our findings appears in Figure 2. The values in this figure were gathered by sending 3,000 messages, each 10,000 bytes long between two nodes. Since the Pentium counters can only measure two values at

<sup>1</sup> The first 50 iterations were dropped to allow for any dirty virtual memory pages from other applications to be flushed.

once, the test program was run multiple times under each operating system configuration. The measured elapsed wall-time for different executions on the same OS configuration varied by less than 3% for both platforms. The first row of the figure shows that Windows NT requires 60% longer to execute the test program based on the elapsed cycles. Two factors contribute to NT's slow performance relative to Linux. First, the total number of instructions executed under NT is 32% more than under Linux. Second, the average number of cycles per instruction (Avg. CPI) is 7.58 on NT, but only 6.00 on Linux. For both operating systems, the average cycles per instructions is substantially higher than would typically be expected for the Pentium processor. However, since this test program consists of only calls to message passing routines, the high CPI values are unsurprising because of poor cache locality due to frequent kernel-user space copy operations.

To obtain a better understanding of the source of the increased cycles per instruction on Windows NT, we used additional hardware counters to measure the performance of various architectural features of the Pentium processor. Specifically, we measured the caches, translation look-aside buffers (TLB), utilization of the second integer execution unit (V-pipe), and branch prediction logic. The results of these measurements also are shown in Figure 2. For each hardware feature, we report the number of operations where that feature slowed program execution, the percent of instructions that experienced this slowdown, the approximate number of cycles required to handle this operation, and the amount this type of operation increased the average cycles per instruction.

To convert from operations to cycles for each of the hardware features, we used published information from the Pentium Programmer Manual[7]. Specifically, 5 cycles for unpredicted branches, 3 cycles each for misaligned data read references, 6 cycles per instruction for v-pipe utilization, and at 41 cycles each for TLB misses were used. In addition, instruction cache misses are assumed to average 14 cycles each. For the data read misses, we used the count of pipeline data memory read stalls reported by the

Pentium monitor.

The last two columns of Figure 2 show the increase in CPI for NT versus Linux. Overall data read stalls are responsible for 37% of the increased CPI under NT, code cache misses, 15%; misaligned data references, 11%; and differences in V-pipe utilization, 11%. The data read stalls include both the time required to service data cache misses and the stalls caused by memory operands on integer instructions other than *mov*. For utilization of the second integer unit (called the V-pipe), the Linux version was able to execute 23.6% of the instructions in this unit and NT was able to execute 20.7% of its instructions there. For this metric, a higher percentage of operations is desirable since it implies that fewer instructions will execute in the primary execution unit. For V-pipe utilization, the approximate cycles and CPI values are an indication of the average number of cycles per instructions *saved* by using the second unit. In addition, branches that were taken, but not predicted by the branch prediction logic are responsible for approximately 7% of the additional cycles per instruction. Finally, code TLB misses and data TLB misses contribute approximately 5% and 4%, respectively.

In summary, the longer code paths (indicated by increased instruction counts), and poor memory hierarchy utilization explain the differences in time required between Linux and NT for the micro-communication benchmark.

To verify that the CPU time for protocol processing time by the NT operating system was responsible for the overall difference in TCP performance, we configured our Ethernet cards to use classic 10 Megabit-per-second Ethernet (by replacing the Fast Ethernet hub with a standard one). When run at lower speeds and large packet sizes, the NT operating system is able to achieve data transfers of 1.1 MB/s, which is nearly wire speed. This confirmed our suspicion, that protocol processing in NT is being masked by the low bandwidth of 10 Mbps Ethernet.

Another significant component of overhead required to send messages is the PVM system itself. By comparing the curves showing the raw TCP results versus the PVM results for the same operating system we see that a consider-

	Linux				Windows NT				Extra CPI	Percent Extra CPI
	Operations	Pct. Instr.	Approx. Cycle Count	Avg. CPI	Operations	Pct. Instr.	Approx. Cycle Count	Avg. CPI		
<b>Instructions executed</b>	164,773,816		989,074,214	6.00	217,488,387		1,648,579,760	7.58	1.58	
<b>Data read Misses</b>	6,073,500	3.69	283,205,815	1.72	12,342,575	5.68	502,519,616	2.31	0.59	<b>37.5</b>
<b>Code cache miss</b>	12,469,533	7.57	162,103,929	0.98	20,542,717	9.45	267,055,321	1.23	0.24	<b>15.5</b>
<b>Misaligned data refs</b>	8,148,426	4.95	24,445,278	0.15	23,361,052	10.74	70,083,156	0.32	0.17	<b>11.0</b>
<b>V-Pipe Utilization</b>	38,829,901	23.57	(233,081,049)	(1.41)	44,963,826	20.67	(269,900,655)	(1.24)	0.17	<b>11.0</b>
<b>Un-predicted Branches</b>	5,986,357	3.63	29,931,785	0.18	12,982,186	5.97	64,910,930	0.30	0.12	<b>7.4</b>
<b>Code TLB miss</b>	367,307	0.22	15,059,587	0.09	872,205	0.40	35,760,405	0.16	0.07	<b>4.6</b>
<b>Data TLB miss</b>	42,122.00	0.03	1,727,002	0.01	1,727,002	0.79	14,783,083	0.07	0.06	<b>3.6</b>
<b>Other</b>									0.15	<b>9.3</b>

**Figure 2: Comparison of TCP Performance.**

able overhead is introduced by PVM over raw TCP/IP. For small messages, overhead is more than twice the pure communication cost of TCP/IP in both Linux and NT. The overhead is still over 60% for large-sized messages. The primary reason for this additional overhead is the extra copy required to implement the pack and unpack semantics of PVM. Other message passing interfaces such as MPI[6] and Nexus[4] do not require this step. In Nexus, for example, the overhead added by the message passing layer has been reported as less than 10  $\mu$ s.

#### 4. Application Benchmarks

The Numerical Aerodynamics Simulation (NAS) parallel benchmarks are a set of eight programs designed to measure the performance of parallel supercomputers. We used a version of the benchmarks implemented using PVM message passing primitives. The NAS suite consists of three C and five Fortran 77 programs. Five of the programs are general-purpose kernels that occur frequently in scientific applications, and three are kernels taken from computational fluid dynamics. This suite of test programs is useful for our study since it includes programs with varying computational granularity and differing message sizes.

In this section, we describe the performance results for each benchmark. The execution times denote the wall clock time and are the average of five runs. Figure 3 shows the performance results of all of the benchmarks for the three configurations. Figure 4 shows the communication statistics for the applications. The same results are presented graphically in Figure 5.

Performance results for Embarrassingly Parallel Kernel (EP) are shown in Figure 3. Since there is no communication, U-Net and Linux TCP results are almost the same. However, the same benchmark takes longer to execute in NT due to the increased execution time of functions in the NT math libraries as shown in Section 6.

The Multigrid (MG) benchmark was run for a data size of  $128^3$ . For this application each processor sends messages ranging in size from 100 bytes to 32K bytes. PVM over U-Net clearly out performs the TCP based configurations. The computation time is almost equal in all platforms but communication time is less in U-Net. This behavior can be explained by the results in Figure 1 that show the overhead of TCP is around 1.5 times that of U-Net for 10K messages (an average message for MG).

For the Integer Sort benchmark (IS) used  $2^{20}$  elements and keys in the range of  $2^{18}$ . Over 75% of the messages sent by this program were one megabyte or larger. As a result of these large messages, bandwidth is the dominate factor; thus, all three platforms perform almost identically.

Block tridiagonal solver (BT) and Pentadiagonal solver (SP) benchmarks are both communication bound and send many medium-sized messages ranging from a few hundred to a few thousand bytes each. Due to these numerous messages, the U-Net versions of the applications run

25 to 130% faster than the Linux TCP version. The NT version takes two to three times longer than the Linux version due to the gap in communication speed between the platforms for messages of 1KB (shown in Figure 1).

App	OS-Comm	2 Processors		4 Processors	
		Total Time	Comm Time	Total Time	Comm Time
EP	Linux-TCP	143.7	0	81.0	0
	Linux-UNet	139.7	0	82.1	0
	NT-TCP	262.2	0	145.3	0
MG	Linux-TCP	33.3	18.9	24.6	17.2
	Linux-UNet	27.3	12.6	18.3	10.9
	NT-TCP	37.3	22.1	27.4	20.0
IS	Linux-TCP	18.0	11.6	16.3	14.0
	Linux-UNet	17.0	10.1	15.4	13.1
	NT-TCP	20.3	13.8	16.4	14.7
BT	Linux-TCP	107.7	81.5	58.3	45.2
	Linux-UNet	70.0	44.4	46.4	33.5
	NT-TCP	262.0	184.6	158.0	144.5
SP	Linux-TCP	107.6	91.2	59.5	52.2
	Linux-UNet	45.5	29.8	33.4	26.7
	NT-TCP	292.6	277.6	199.6	193.2
FT <sup>2</sup>	Linux-TCP	-	-	22.5	17.8
	Linux-UNet	-	-	15.3	10.8
	NT-TCP	-	-	38.1	32.1
LU	Linux-TCP	15.6	10.9	15.7	13.1
	Linux-UNet	7.6	3.0	5.4	3.2
	NT-TCP	103.0	98.0	100.2	96.6

Figure 3: Comparison of Execution Times.

Program	Bytes	Msgs	Msg Size	Bytes/sec <sup>3</sup>	Msgs/sec
MG	1,532,360	181	8,466	140,583	16.6
IS	33,558,120	41	818,491	2,561,689	3.1
BT	88,708,888	52,923	1,676	2,648,027	1,579.8
SP	30,948,112	22,703	1,363	2,865,566	2,102.1
FT	11,010,384	42	262,152	3,440,745	13.1
LU	244,944	5,105	48	9,174	191.2

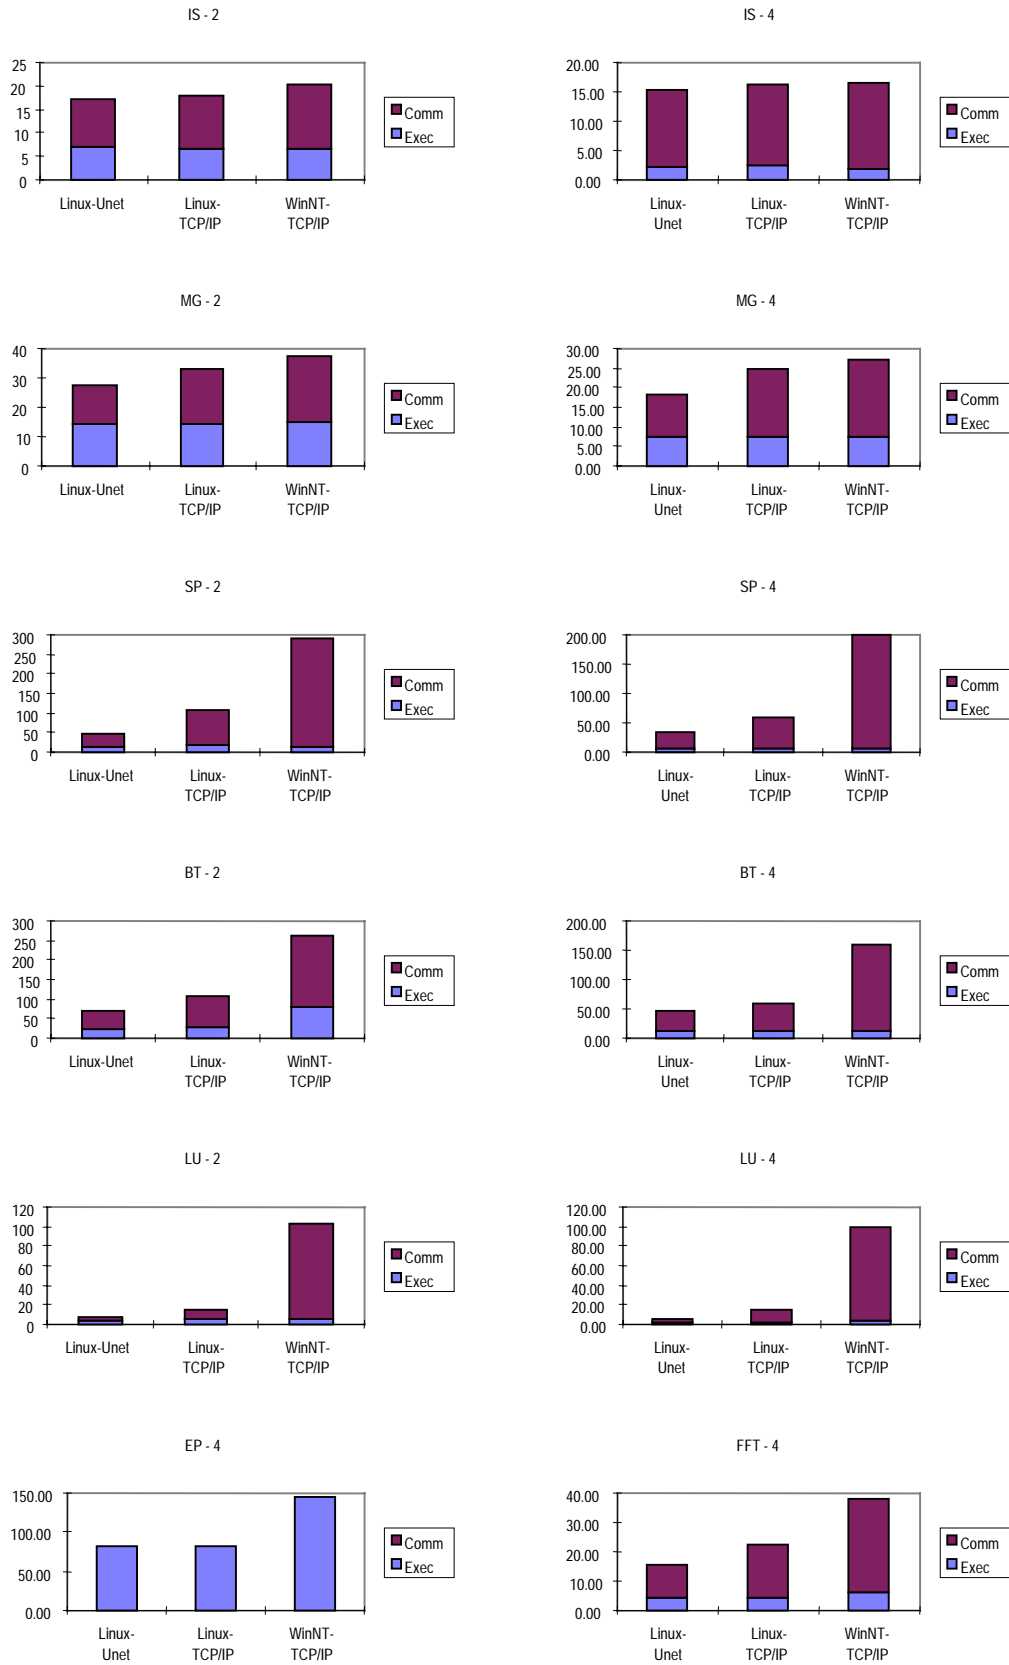
Figure 4: Message Statistics for NAS Applications.

Although Fast Fourier Transform (FFT) sends a small number of large messages (about 256KB each); even this does not save the benchmark from being communication bound as shown in Figure 3. The U-Net communication time is still about 45% faster than Linux TCP. However, the communication cost reduction is not as large as it was for the LU benchmark because the message size is larger.

The Lower block-Upper block matrix decomposition (LU) benchmark sends very small messages, less than 50 bytes each. Because of these small messages U-Net out performs the others by the largest margin on this application - a factor of 2.9 faster than Linux and 18.6 times faster than NT.

<sup>2</sup> FT could not be run on two nodes due to a lack of memory.

<sup>3</sup> The aggregate rate for all processes in the entire application.




**Figure 5: Performance Results for All Benchmarks (All Y-axes in seconds).**

## 5. Instruction and Memory Behavior


Many scientific programs are constrained not only by computation, but also by the amount of memory available and the performance of the memory hierarchy. Due to the importance of the memory hierarchy to application performance, we compared the performance of the memory hierarchy under each OS configuration for the NAS benchmarks.

Figure 6 shows the floating point operations performed in user mode. In both systems, the results are similar. In EP, BT, and SP, NT performs more operations than Linux; these extra operations are due to the floating point operations performed as part of the NT TCP/IP implementation. (During our micro-benchmarking of TCP/IP, we verified that NT performs user space floating point operations during communication.)



**Figure 6: FPU Instructions (normalized to NT).**

Figure 7 shows the user mode instruction cache misses for the TCP versions of the NAS applications. NT clearly has bad cache locality. For the three applications with significant instruction cache misses (LU, BT, and SP), NT's miss rate is 8 to 400% higher than Linux. Chen et al.[2] have shown a similar result for sequential workloads. We speculate that these cache misses result from frequent context switches among system services and applications caused by message passing. All three applications with poor instruction cache miss rates frequently send messages. These three applications also had the greatest slowdown relative to Linux/TCP. Also, the poor cache locality is consistent with the micro-benchmarks presented in Section 2.




**Figure 7: Code Cache Misses.**

## 6. Math Libraries

Based on the observation that the EP application showed such a large variation in performance when run on NT and Linux, we decided to investigate why two identical

compute-bound jobs would have different behavior. We discovered that the problem was due to differences in the way the math libraries are implemented on the two platforms.

We extracted the main computational loop of the EP benchmark and compiled it with gcc and Visual C++ on NT, and gcc on Linux. Figure 8 shows the results of running this program. This benchmark makes extensive use of math library function calls. Gcc and Visual C++ on NT produce similar results. However, both NT versions execute more than twice the number of instructions executed by Linux gcc, yet the number of floating point operations performed is approximately the same. By disassembling the math library functions for each program, we verified that the default math libraries on NT for both gcc and the Microsoft C compiler generate software implementations of rounding down and up (floor, ceil), absolute value (fabs), and logarithm (log10). The Linux version of the math library used the Pentium hardware instructions for these functions. We then replaced the version of the math library with one that used the native instructions for these functions. The performance under Windows NT using the modified math libraries is within 10% of the performance under Linux. The results of that experiment are shown in the bars labeled "fixed math".



**Figure 8: Comparison of Math Libraries.**

## 7. Related Work

Chen et al.[2] have studied the interactions of application programs on different personal computer operating systems (Windows, Windows NT, and Net BSD). However, the applications were traditional sequential UNIX applications. Our work differs in that we measure the performance of communication intensive parallel applications.

Nahum et. al[11] studied the cache behavior of network protocols on actual and simulated versions of the MIPS R4400 processor. Their study showed that with direct mapped caches and TCP checksums disabled, instruction cache performance dominates the memory hierarchy's performance. Our results show that for Pentium processors, that the performance of the data cache is more important than the instruction cache. However, for a configuration comparable to ours (8KB two-way set associate caches and

TCP checksums enabled), their data indicates that the data cache has a higher miss rate than the instruction cache.

An alternative to provide higher communication throughput is third party transfers where a user process supplies a file handle and a network connection to the kernel, and the kernel completes the transfer. Such systems have been proposed by Fall and Pasquale[3], Miller and Tripathi[10], and commercially by Microsoft[9]. However, third party transfers are only appropriate for cases where data is merely being shipped from one device to another (e.g., file servers or static web page servers). For other applications such as parallel computing, database processing, or dynamic web page creation, efficient data movement from user space process to the network is required.

## 8. Conclusions

In this paper, we presented the performance of a network of four PCs connected with Fast Ethernet running the NAS parallel benchmarks under two different operating systems and two message passing systems: PVM and U-Net Active Messages. The Embarrassingly Parallel benchmark, which involves just computation and no communication between processes, shows how well PCs perform on computationally bound, floating point intensive workloads. However, EP performed worse on NT than on Linux due to the poor performance of its math library.

The more common barrier is the high latencies incurred by message passing libraries. We presented the execution times of benchmarks when a low cost message passing library, U-Net Active Messages, is used instead of TCP. Message passing costs are much lower for small messages in U-Net than TCP. The LU benchmark is a typical example of this kind of a program; communication cost is reduced by a factor of three compared to TCP/IP-based PVM. But even in the FFT benchmark, where cooperating processes send large messages (on the order of 250KB), U-Net performs better than TCP.

We also presented results for PVM under two operating systems: Linux and NT. For most of benchmarks, execution times are longer with NT due the poor performance of its TCP/IP stack. As we showed in Section 3, round trip time for TCP is much higher for NT than it is for Linux. In terms of processor performance, NT also shows bad cache locality. For two of the applications NT had instruction cache miss rates several times higher than Linux. We also showed that the math library implementations on both systems are considerably different. The gcc math library in Linux is much more efficient than either the Visual C or gcc math libraries in NT. This is very important drawback for NT, especially for scientific applications.

In conclusion, the standard distributions of NT and Linux are not able to achieve consistently high performance for parallel applications. However, with the addition of low latency message passing libraries, Linux achieved a significant fraction of the raw hardware performance.

## Acknowledgments

The authors thank Matt Welsh and Anindya Basu from Cornell University for their great help and patience for answering our many questions and addressing the problems we encountered during the installation and use of the U-Net software and Active Messages. We thank Ismail Hari-taoglu from the University of Maryland for his help in running the Windows NT benchmarks. We also thank Rich Gerber for his helpful suggestions about the structure and presentation of this paper.

## References

1. D. H. Bailey, E. Barszcz, J. T. Barton, and D. S. Browning, "The NAS Parallel Benchmarks," *International Journal of Supercomputer Applications*, **5**(3), 1991, pp. 63-73.
2. J. B. Chen, Y. Endo, K. Chan, D. Mazieres, A. Dias, M. Seltzer, and M. D. Smith, "The Measured Performance of Personal Computer Operating Systems," *15th SOSP*, Dec. 1995, pp. 299-313.
3. K. Fall and J. Pasquale, "Improving In-Kernel Data Paths," *IEEE International Conference on Multimedia Computing and Systems (ICMCS)*, 1994, pp. 100-109.
4. I. Foster, C. Kesselman, and S. Tuecke, "The Nexus Approach to Integrating Multithreading and Communication," *J. Parallel and Distributed Computing*, **37**(1), 1996, pp. 70-82.
5. A. Geist, A. Beguelin, J. Dongarra, W. Jiang, R. Manchek, and V. Sunderam, *PVM: Parallel Virtual Machine*, 1994, Cambridge, Mass: The MIT Press.
6. W. Gropp, E. Lusk, and A. Skjellum, *Using MPI: Portable Parallel Programming with the Message Passing Interface*, 1995: MIT Press.
7. Intel, *Architecture and Programming Manual*. Pentium Processor Family Developer's Manual. Vol. 3. 1997.
8. T. Mathisen, "Pentium Secrets," *Byte*, **19**(7), 1994, pp. 191-192.
9. Microsoft Corporation, *WDM Connection and Streaming Architecture*, 1997.  
<http://www.microsoft.com/hwdev/pcfuture/csa1.htm>.
10. F. W. Miller and S. K. Tripathi, "An Integrated Input/Output System for Kernel Data Streaming," *Multimedia Computing and Networking*, Jan. 1998.
11. E. Nahum, D. Yates, J. Kurose, and D. Towsley, "Cache Behavior of Network Protocols," *SIGMETRICS*, May 1997, Seattle, WA, pp. 1-12.
12. T. von Eicken, D. E. Culler, S. C. Goldstein, and K. E. Schause, "Active Messages: a Mechanism for Integrated Communication and Computation," *19th International Symposium on Computer Architecture*, June 1992, Gold Cost, Australia, pp. 256-266.
13. M. Welsh, A. Basu, and T. von Eicken, "Low-Latency Communication over Fast Ethernet," *Euro-Par '96*, Aug. 1996, Lyon, France, vol. I, pp. 187-194.