

# Improving End System Performance for Multimedia Applications over High Bandwidth Networks

S. ZEADALLY, G. GHEORGHIU AND A.F.J. LEVI

*Department of Electrical Engineering, University of Southern California, University Park, DRB 116, Los Angeles, CA 90089-1111*

**Abstract.** Digital video services, scientific visualization and other multimedia applications require delivery of high network throughput to end user applications. In this paper we identify bottlenecks in the data path between high-speed networks and applications. Using performance of multimedia applications as a metric, the effectiveness of solutions to reduce network, operating system, and user bottlenecks is explored experimentally.

**Keywords:** multimedia, networking, operating systems, performance, TCP/IP

## 1. Introduction

The past few years have seen development of applications with high bandwidth requirements such as medical image transfer, video conferencing, scientific process simulation, and visualization. Popular local area networks such as Ethernet and Token Ring (4 Mbit/s and 16 Mbit/s) are incapable of providing the bandwidth needed by these multimedia applications. New network technologies such as Fiber Distributed Data Interface (FDDI), Fast Ethernet (100BASE-T), and Asynchronous Transfer Mode (ATM) have emerged and are capable of providing high bandwidth to users' desktops. However, the challenge remains for operating system designers and application developers to deliver the bandwidth of these networks to end user applications. In order for applications to reap the benefits of high-speed networks, the entire path from network to application must be optimized. This involves removing bottlenecks introduced both at the operating system and application levels as depicted in figure 1.

This paper describes our experiences delivering multimedia services over high-speed networks and our attempts to optimize application performance. We deal with *real-life multimedia applications* as opposed to using *raw data*. The use of raw data in performance tests provides an upper bound on achievable performance that can be delivered by the underlying software and hardware. However, raw data does not identify other bottlenecks that are normally associated with actual applications such as the presentation of information (e.g., video display in a window) or the effect of running multiple applications concurrently.

Related work undertaken by other researchers in similar areas include: Keller et al. [14] showed the effectiveness of shared memory in the delivery of digital video to user applications. However, no optimization was attempted at either operating system or network layers, which led to poor overall application performance due to various protocol and system overheads (e.g., data copying). In [23], Yau and Lam proposed an architecture which provides

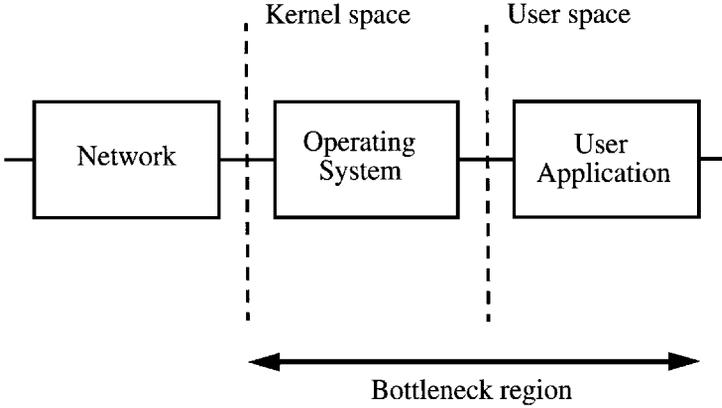


Figure 1. Illustration of bottlenecks in typical multimedia network-application data path for conventional end systems considered in this study.

operating system support for multimedia networking. Their work focused on optimizing efficient data transfers between user and kernel, and exploit independently scheduled kernel threads for control transfer between user and kernel in order to minimize system calls. Although simulation experiments were used for demonstrating the effectiveness of the architecture, issues related to multimedia data presentation and real-life multimedia application performance were not discussed. Saha et al. [20] improved network-application data path at the operating system level using kernel data paths between network adapters and CODEC adapter. Similarly, they did not discuss network-application performance.

This paper differs from these efforts in that we apply optimizations wherever possible to the *entire* data path between network and applications as opposed to focusing on just one stage in the transfer whereby later stages introduce other overheads which degrade the overall application throughput. For instance, no matter how fast the network delivers data to the host, if the operating system itself cannot transfer the data to the application at a high enough sustained rate, then there will be a bottleneck introduced in the network-application data path.

The structure of this paper is organized as follows. In Section 2, we discuss techniques that we applied to increase application-application throughput. Section 3 describes the experimental setup we have used for our multimedia applications. Section 4 presents an analysis of in-host data movement for our scientific graphics visualization application and digital video playback. Section 5 discusses the results obtained by implementing the various optimization techniques given in Section 2. In Section 6, we analyze the overall performance when multiple applications are running. Finally, Section 7 makes some concluding remarks and presents future work.

## 2. Improving application-application throughput

To achieve high application-application throughput in a network environment, it is essential that high sustained throughput be delivered by the network, operating system, and application.

Recent advances in the performance of network physical layers have essentially solved the network bandwidth problem. As a result, it is now feasible to deliver large volumes of data at high rates with minimal loss. In this work, we have used a conventional Ethernet network, a higher speed FDDI network, and an experimental network called Jetstream [21] (Section 3) in order to study their impact on application performance. The basic physical layer characteristics of these networks include: Manchester code Ethernet signalling at 20 Mbit/s giving a maximum data rate of 10 Mbit/s, 4b/5b coding on FDDI signalling at 125 Mbit/s delivering a maximum data rate of 100 Mbit/s, and 16b/20b coding for Jetstream signalling at 1 Gbit/s giving a maximum data rate of 800 Mbit/s. These maximum data rates are reduced by media access control, network subsystem, application, and of course, limitations imposed by the host architecture.

The last few years have witnessed significant hardware improvements that have led to the development of powerful computers. Some of these improvements include: increased CPU performance, high bus bandwidth, large memories, and fast disk systems. However, there has been little change in the structure of conventional operating systems such as UNIX, consequently the availability of new hardware technologies has not been exploited to the fullest. This has made existing operating systems become a bottleneck in end systems. Well-known overheads include data copying, network protocol processing, context switches, and interrupts [13, 16, 18]. Several techniques have been proposed and implemented to avoid these overheads [6, 7, 10, 19]. In this work, we minimized data copying by using a UNIX kernel which supports “single-copy” TCP/IP, a modified version of TCP/IP. Throughout this paper, we refer to “two-copy” TCP/IP as the standard version that normally comes with UNIX operating systems. In this case, data transfer between network and application normally involves two copies: the first copy is between a network buffer and kernel buffer followed by a copy from the kernel buffer to a user application (for an incoming packet). The reverse takes place for an outgoing packet. However, in the case of a single-copy TCP/IP implementation we use for our experiments in this paper, there is only one data copy between network and application thereby eliminating the copy to kernel buffer. Moreover, the single-copy implementation of TCP also supports RFC 1323 window scaling [11], and is capable of calculating checksum during data movement.

It is not easy to come up with general techniques to increase throughput at the application level. The main reason is that different applications have different requirements and each is implemented in its own way. However, it is true that multimedia applications have a common element: they all present information (e.g., video display) to the end user. Most desktop applications running on UNIX platforms are built on standard X window systems to increase their ease of use, and offer a common look and feel to users. The X window system has become the de facto standard graphical user interface for UNIX systems. We argue that there is scope for improving application performance in the X environment in the area of data presentation. In this context, we note that without careful tuning, data display by the X server can degrade performance in the final delivery of information to the user. Our choice was to use the X shared memory extensions [3] in order to speed up image display.

The usual way to display an image is to use the X11 library call **XPutImage()** on an application's data. The call to **XPutImage()** moves data from the application's buffer via Inter-Process Communication (IPC) to a private buffer of the X server (using UNIX domain

sockets when the X client and the X server are on the same machine). The data is then moved by the X server to the frame buffer. With X shared memory extensions support, there is no data movement involved between the application and the X server. Instead the image data is placed into a memory segment that is shared between the application and the X server. In this case, a call to **XShmPutImage()** allows the X server to move data directly from the shared memory segment containing the application image data to the frame buffer.

### 3. Experimental arrangement for multimedia applications

The experiments described in this section have been carried out between two HP 9000 Series 700 workstations (99 MHz PA-RISC) which reside on the Jetstream network. Jetstream is a Gbit/s token-ring network which uses copper coaxial or fiber optic cable for the physical link. The network adapter for the HP 9000 Series 700 workstations is made up of two cards—one is called Afterburner which is equipped with 1 MByte of video random access memory used in a dual ported configuration. Afterburner is the host interface; the other card, Jetstream, is the link adapter. The shared memory present on the Afterburner board enables the support of *single-copy* implementations of network protocols such as TCP/IP and UDP/IP. Further details of Afterburner and Jetstream are given in [21] and [4].

The hardware architecture of the host is illustrated in figure 2. The Standard Graphics Connector (SGC) [5] is the system bus used on the HP 9000 Series 700 workstation and has a maximum data transfer rate of a Gbit/s. However, it is only possible to achieve a maximum of 400 Mbit/s transfer rate by the CPU between memory and input/output (I/O) space (e.g.,

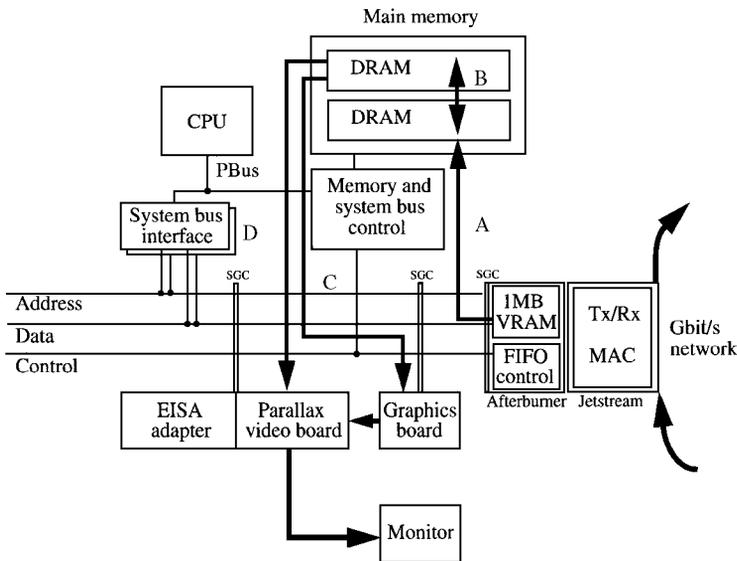


Figure 2. HP 9000 Series 700 architecture data paths for application displaying network data and video playback. A network adapter connects the workstation to the Jetstream network. A Parallax video board performs video decompression and drives the monitor. Graphics data is passed from main memory to the graphics board and overlaid on the monitor by the Parallax video board. A: Network—Application (main memory), B: Application—X server, C: X server—Framebuffer, D: Main memory—Parallax board.

graphics devices or network interface) [8]. A principal feature of the hardware architecture is the memory and system bus controller chip which connects the CPU to memory and the I/O system components. The system bus controller chip communicates to the SGC bus via two system bus interface chips. The workstations used in our experiments have 128 MByte of main memory and two GByte of hard disk. The operating system used was HP-UX 9.01 with single-copy TCP/IP support.

We have used two different applications in our experiments: one is a scientific visualization application and the other is digital video playback.

The visualization application uses a series of gray scale images (8 bit/pixel) of a volume rendered CAT-scan medical image of a child's head. The head can be rotated and viewed at different angles. In our experiment, the images are stored on a remote server and sent over the network to the client machine which displays the images. The image set consists of 20 image frames each of size  $512 \times 512$  pixels (almost 2.1 Mbit per frame) and stored in pixmap format. This makes direct display by the X server possible without requiring any further manipulation. The user interface to the visualization application supports simple operations such as play, stop, and rewind.

The video application uses the PowerVideo700 hardware video codec (compression/decompression) from Parallax [17] which supports Motion-JPEG. Motion-JPEG applies JPEG (Joint Photographic Experts Group, a standardized image compression technique for still images) to individual frames of a video sequence. The video board is capable of handling high quality video at 30 frames per second in real-time during either recording or playback sessions. The PowerVideo700 is an overlay card which resides in one of the EISA slots of the EISA interface attached to the SGC bus as shown in figure 2. We have used the MovieTool software from Parallax for recording and playing digital video stored as Motion-JPEG files.

For our experiments, the files used during video playback were stored on the hard disk of a remote machine. This disk was mounted on the host machine using NFS via the Jetstream interface (details are given in Section 4). In a typical video playback session, the use of NFS enables the host machine to receive compressed video clips over the Jetstream network. The video board decompresses the incoming compressed video stream. The analog signals originating from the graphics card are digitized and the result is overlaid with the uncompressed video image. After the overlay is completed, the entire frame is converted back to analog and sent to the monitor (figure 2).

In all measurement tests for both the video and scientific visualization applications, we use average playback frame rate as our quantitative metric to characterize application performance. Moreover, in all experiments, these applications were run as normal user processes, along with the usual system processes and daemons in the background. Furthermore, all our tests were conducted using *unloaded* networks, which carried no other traffic but the one from the applications under investigation.

#### 4. Analysis of in-host data movement

In this section, we identify the data paths used when running the video and visualization applications based on the architecture presented in figure 2.

To understand the impact of the underlying architecture on application performance, we measured the throughput at different stages when moving data from the network to the

X window display for the visualization application. The major aim of performing such analysis is to identify areas where performance can be improved, and at the same time assess the suitability of the HP 9000 Series 700 workstation architecture in supporting high-speed network applications. Our observations are applicable to other similar networked multimedia applications (e.g., medical imaging, video conferencing).

Equation (1) summarizes the inverse of total throughput  $R$  for typical applications that read data from the network and use the X window system for display.

$$\frac{1}{R} = \frac{1}{Z_{\text{Network-Application}}} + \frac{1}{Z_{\text{Application-X server}}} + \frac{1}{Z_{\text{X server-Frame buffer}}} \tag{1}$$

where  $Z$  represents the throughput at different key stages of the data path from network to frame buffer as depicted in figure 2. Furthermore, Eq. (1) applies to conventional bus-based systems of the type used in our experiments.

To understand the impact of running multiple network applications on the performance of the end system, we have chosen to simultaneously run both the video and the visualization applications over Jetstream. For this to be possible in the case of the video application, we used NFS over the Jetstream network interface as shown in figure 3. This enables playback

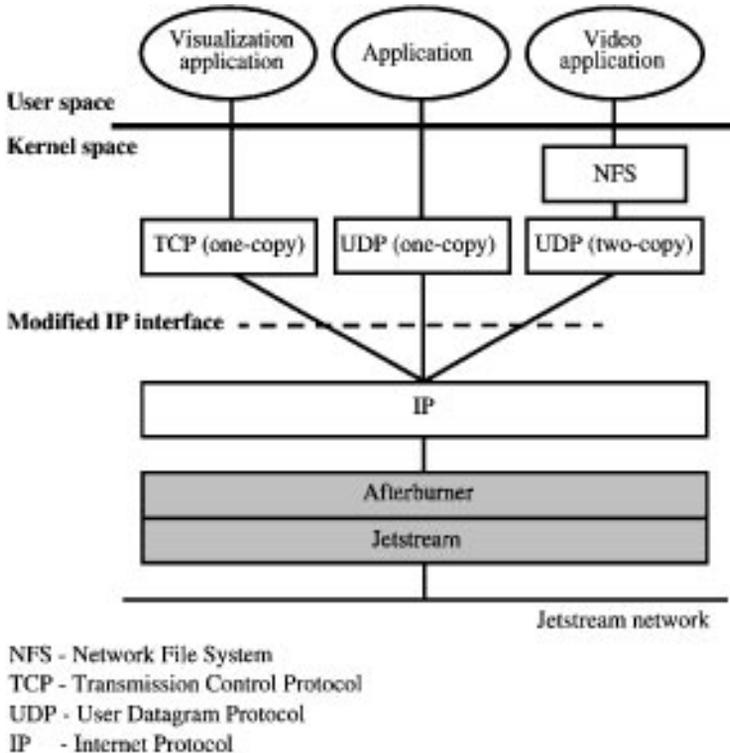


Figure 3. Protocol stacks with one-copy and two-copy support. IP layer has been modified to allow NFS support (which uses two-copy UDP/IP) to co-exist with applications using single-copy UDP/IP.

of digital video clips stored on a remote disk which has been mounted on the host machine (used as an NFS client). Our UNIX kernel supports NFS 2.0 which uses UDP.

Although the UNIX kernel we have used does support both single-copy TCP/IP and single-copy UDP/IP stacks, it was not possible for NFS to use the single-copy UDP/IP stack. This is because NFS does not understand the buffer structures used in the single-copy implementation of UDP/IP. We did not consider it worth modifying NFS to allow it to support our single-copy protocol stacks. The justification was that we do not see significant throughput improvement with a version of NFS that allows single-copy protocols since the disk (at the NFS server) will still be the bottleneck (although latency would be slightly better). For our experiments, it would have been sufficient to use a single-copy TCP/IP stack (for the visualization application) and a two-copy UDP/IP stack (for video application over NFS) by simply disabling single-copy support for UDP/IP in the UNIX kernel. However, the disadvantage of this approach is that it prevents other applications from using single-copy UDP/IP. Our solution was to modify the IP layer in order to distinguish packets destined for NFS which will use two-copy UDP/IP from all other incoming network packets which will use single-copy UDP/IP or single-copy TCP/IP (figure 3).

### 5. Network measurements and results

Initial experiments were conducted using *raw data* to investigate how much of the available network bandwidth can actually be delivered to the application. Figure 4 presents the

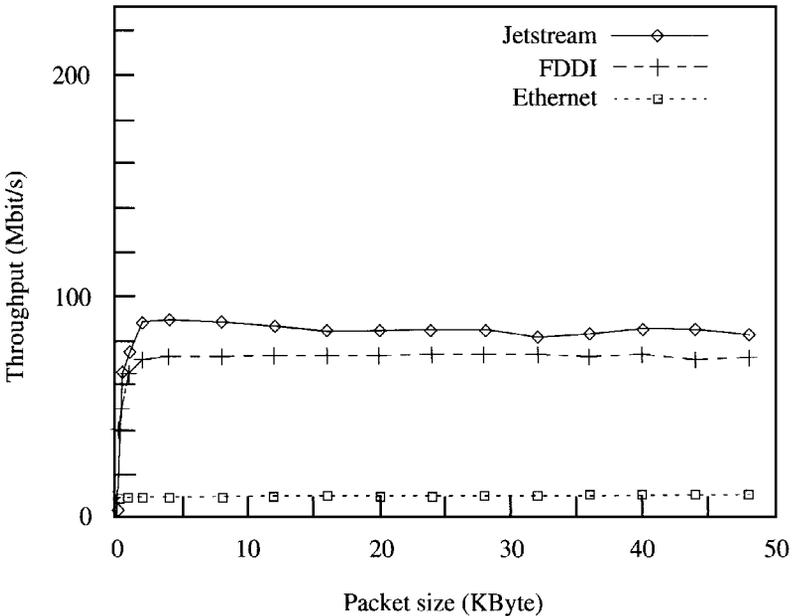


Figure 4. Measured raw data throughput over Ethernet, FDDI and Jetstream using two-copy TCP/IP and 32 KByte socket buffer size.

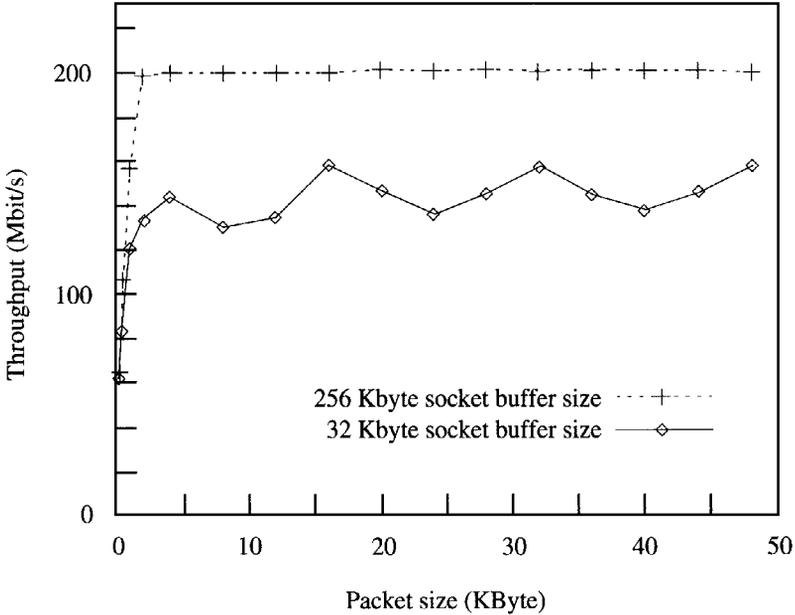


Figure 5. Measured raw data throughput over Jetstream using single-copy TCP/IP and socket buffer sizes of 32 KByte and 256 KByte.

observed throughput using a two-copy TCP/IP stack over Ethernet, FDDI and Jetstream. The laboratory Ethernet has been used for Ethernet tests. In the case of FDDI, an EISA FDDI adapter designed for HP 9000/700 EISA systems was used to attach to an FDDI network. Performance measurements were made with a tool called *netperf* [12] which measures the transfer of data from a producer process (generating the data) to a consumer (receiving data) running on a remote machine. The maximum network data bandwidths for Ethernet, FDDI, and Jetstream are 10 Mbit/s, 100 Mbit/s, and 800 Mbit/s respectively. However, the maximum raw data throughput values obtained using two-copy TCP/IP in our experiments were around 9.6 Mbit/s, 74 Mbit/s and 90 Mbit/s for Ethernet, FDDI, and Jetstream respectively. These results confirm our initial assumption that the end system has become the bottleneck in a high-speed network environment.

Figure 5 shows raw data application throughput using single-copy TCP/IP over the Jetstream network. The use of a single-copy kernel increases throughput by almost 56% using 32 KByte socket buffer size. Increasing the socket buffer size, if the kernel allows it, results in higher throughput values. For instance, the single-copy kernel enabled us to specify a socket buffer size of 256 KByte. In this case the maximum throughput achieved was around 200 Mbit/s.

Having established the achievable throughput possible using raw data, we then ran the visualization application to measure the maximum throughput for a real application running over various networks, using a two-copy TCP/IP kernel and 32 KByte socket buffer size. In contrast to raw data throughput results, we obtained frame rates and corresponding throughput values given in Table 1. Two important observations can be made: first, performance

Table 1. Measured graphics visualization application frame rate and the equivalent throughput using two-copy TCP/IP and 32 Kbyte socket buffer size over Ethernet, FDDI and Jetstream.

Network type	Ethernet	FDDI	Jetstream
Frame rate (frames/s)	4	11	22
Throughput (Mbit/s)	8.4	23.1	46.2

based on raw data is not enough to characterize application throughput; second, in the case of slow networks like Ethernet, software (e.g., operating system) on end systems is able to deliver most of the available network bandwidth to the application. However, as network speed increases, the discrepancy between network bandwidth and actual application throughput is increasing.

Next, we ran the graphics visualization application over Jetstream using single-copy TCP/IP with a large socket buffer size of 256 KByte and obtained a frame rate of 33 frames/second. This corresponds to a throughput of 69.3 Mbit/s (33 frames per second multiplied by 2.1 Mbit per frame). To verify the correctness of Eq. (1), we measured the throughput values corresponding to the three stages of the network-application data path as follows:

- The application reads data from a socket into its buffer. The fact that we are using a single-copy TCP/IP stack allows direct data transfer from a network interface buffer to the application's buffer (avoiding the additional copy to a kernel buffer). The throughput obtained was 200 Mbit/s ( $Z_{\text{Network-Application}}$ ).
- The application acting as an X client sends the data to the X server using inter-process communication. The rate of data movement is 240 Mbit/s ( $Z_{\text{Application-X server}}$ ).
- The X server then moves the data to the frame buffer. The transfer rate obtained along this path is 230 Mbit/s ( $Z_{\text{X server-Frame buffer}}$ ).

Using the measured throughput values, we calculated the total throughput  $R$  from Eq. (1) and obtained 74 Mbit/s. This value is slightly higher than the observed Jetstream throughput of 69.3 Mbit/s. The difference of 4.7 Mbit/s is due to the fact that we did not take into account various overheads such as system calls, context switches, interrupts and memory allocation by the X server during data copying.

We now discuss how the various optimizations mentioned in previous sections can be applied to the data path described by the three terms of Eq. (1).

- *Network:* We use the Jetstream Gbit/s network capable of supporting high bandwidth applications. It is worthwhile noting that an individual application will benefit from an increase in network bandwidth, but this is not necessarily true when multiple traffic streams are competing for the available bandwidth. In this case, bandwidth allocation schemes [1, 22] are required for guaranteed performance. We do not discuss these issues here, however, since this is not the topic of this paper.

- *Operating system:* A UNIX kernel that supports single-copy TCP/IP has been used. This allows direct data copy from the network to the application. As a result, data movement, considered to be the major bottleneck in current operating systems, is minimized. Moreover, network protocol overheads such as checksum calculations have also been significantly reduced. We have therefore optimized the first term of Eq. (1) with  $Z_{\text{Network-Application}}$  being 200 Mbit/s. This value is the maximum throughput obtained when using raw data as illustrated in figure 5.
- *Application:* We have used the X shared memory extensions to eliminate data movement from the application to the X server. This optimizes the overall throughput by eliminating the second term  $Z_{\text{Application-X server}}$  of Eq. (1). It is also worth noting that applications should exploit the capability of using large socket buffer size whenever the kernel allows it. Although this is done at application level, it influences the throughput between network and application at the operating system level.

Figure 6 summarizes the effects of the various optimization approaches on the performance of the visualization application. It is interesting to note from the graph that for socket buffer sizes up to 48 KByte (the maximum allowable by standard UNIX kernels), the application performs better using a two-copy TCP/IP and X shared memory extensions than using a single-copy TCP/IP without X shared memory extensions.

The underlying architecture did not allow further optimization of the last term of Eq. (1). This is because the system and memory bus controller shown in figure 2 becomes the

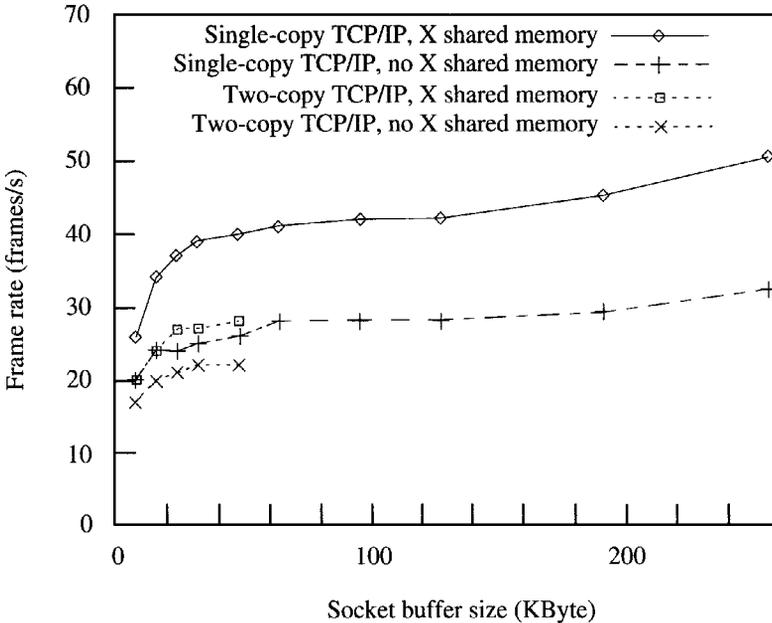


Figure 6. Measured effect of single-copy TCP/IP, X shared memory extensions, and socket buffer size on final throughput for the graphics visualization application.

bottleneck when subjected to intensive data traffic to and from main memory and system bus. This is not a problem for one way traffic between main memory and a peripheral device or vice-versa. However, it becomes a limitation in the case of networked multimedia applications where data flows in continuously from network to main memory and back out from main memory to graphics display. A possible solution is to transfer data directly from a network device to the frame buffer over the system bus, a mechanism commonly referred as *kernel-level streaming* [15]. Unfortunately, the SGC bus implementation in the machines used for our experiments does not allow the needed slave-slave bus transactions.

After applying all the above optimizations, Eq. (1) becomes:

$$\frac{1}{R} = \frac{1}{Z_{\text{Network-Application}}} + \frac{1}{Z_{\text{X server-Frame buffer}}} \quad (2)$$

Calculating the value of  $R$  from Eq. (2), using 200 Mbit/s for  $Z_{\text{Network-Application}}$ , and 230 Mbit/s for  $Z_{\text{X server-Frame buffer}}$ , we obtain an overall expected throughput of 107 Mbit/s. To verify the correctness of Eq. (2), we measured the average frame rate for the visualization application after we implemented all the above optimizations. The value obtained was 50 frames/second, which translates to a throughput of 105 Mbit/s. The difference between the measured and the expected values is smaller (2 Mbit/s) than that obtained in Section 4 (i.e., 4.7 Mbit/s). This is because of the elimination of the second term of Eq. (1) has also reduced overheads such as memory allocation by the X server.

## 6. Overall performance for multiple applications

To quantify the impact of running *both* video and visualization applications on overall performance, we used the following metrics: CPU usage, frame rate (visualization application), and percentage of frames dropped (video application). CPU usage was measured using *Glance* [9], a performance monitor tool that comes with standard HP-UX operating system. Table 2 summarizes the results obtained using the Jetstream network. All tests have been performed on a kernel that supports single-copy TCP/IP. For the visualization application, we used a socket buffer size of 256 KByte. Digital video playback was via NFS at 30 frames/second, the size of each frame being  $512 \times 380$  pixels and 24 bit color per pixel.

To better understand the degradation in performance when both applications are running concurrently, we first make some observations on their performance when executed on their own. When the video application is running by itself, there are no video frames dropped. For the visualization application, the frame rate without using X shared memory is 33 frames/second. However, when running both applications, there was a 39% drop of frames displayed with video and the frame rate observed for visualization decreased to 26 frames/second as shown in the Table 2. The degradation of video performance is due to the high percentage of CPU time (56%) spent in system mode. This is because the visualization application uses IPC to move data to the X server which involves multiple kernel-user interactions. The 34% of CPU cycles left for user mode are not sufficient for the demands of video, which requires 42% user-mode CPU time. On the other hand,

Table 2. Measurement of impact of CPU utilization on video and graphics visualization applications performance.

	VIDEO	GRAPHICS	VIDEO and GRAPHICS	GRAPHICS-SHARED-M	VIDEO and GRAPHICS-SHARED-M
% CPU in user mode	42	27	34	40	41
% CPU in system mode	48	66	56	54	50
GRAPHICS frame rate (frames/s)	—	33	26	50	35
% frames dropped for VIDEO	0	—	39	—	10

VIDEO video playback application.

GRAPHICS visualization application without X shared memory extensions.

GRAPHICS-SHARED-M visualization application with X shared memory extensions.

56% of CPU time spent in system mode is not sufficient for the needs of the visualization application which requires the CPU to spend 66% of its time in system mode. Thus, it is evident that the conflicting requirements of the two applications affect their overall performance.

From Table 2, we note that with X shared memory extension support, not only the visualization application has a higher frame rate on its own, but there is also an improvement in overall performance when both applications run. That is, only 10% of frames are dropped by the video application and the frame rate increased from 26 to 35 frames per second for the visualization application. The use of shared memory significantly reduces kernel-user interactions by eliminating data movement by IPC. As a result, less time is spent in system mode (50%) thereby increasing the availability of CPU for user mode (41%). This obviously benefits the video application. Also, the frame rate increase for the visualization application can be explained by the fact that with shared memory it requires 54% of CPU in system mode as opposed to 66% when not using X shared memory.

## 7. Conclusions and future work

In this paper, we demonstrate that to achieve high application-application throughput in a high-speed network environment, we need to solve bottlenecks at all levels: network, operating system, and application. We have shown how using various optimization techniques, it is possible to increase network-application performance. These techniques include the use of a Gbit/s network, single-copy schemes (including improved protocol processing), and X shared memory extensions. Figure 7 summarizes the throughput optimizations for the visualization application. At each level of the data path between network and application, we apply the optimizations from the previous level. Thus, the final throughput of 105 Mbit/s for Jetstream is the result obtained after applying optimizations at all levels. Compare this to the results shown in Table 1, where without any optimization, the throughput was 8.4 Mbit/s for Ethernet, 23.1 Mbit/s for FDDI, and 46.2 Mbit/s for Jetstream.

	Network level	Operating system level	Application level
Throughput optimization Mbit/s	Ethernet/FDDI 10/100	—	—
	Jetstream 800	Two-copy 90	—
		Single-copy 200	No shared memory 69.3
			Shared memory 105

Figure 7. Summary of all applied optimizations. Each level includes the optimizations from the previous level.

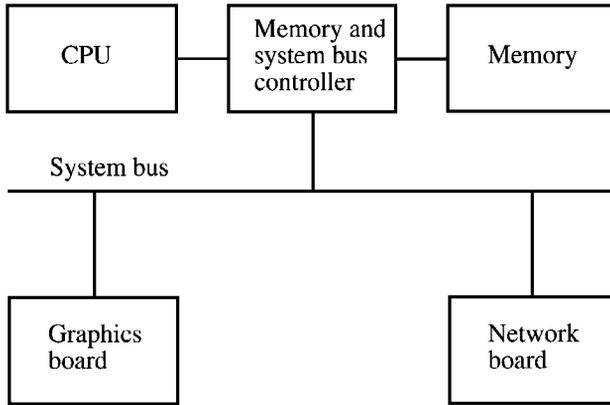


Figure 8. Simplified HP 9000 Series 700 architecture.

As pointed out in Section 4, the hardware architecture prevents the set up of a direct data path between network adapter and display for those network applications that require minimal or no data processing. Furthermore, as depicted in figure 8, the system and memory bus controller interconnects CPU, memory, and the I/O subsystem, thereby typically becoming the bottleneck during concurrent access or transfer of data between these components. This can limit the performance of network multimedia applications which involve *simultaneous* data transfer from network to main memory, and from memory to display device.

We are investigating new architectures which will better cope with the demands of multimedia applications in the context of high-speed networks. We anticipate that new switch-based bus architectures will allow greater flexibility in setting up different data paths between components of the system. Historically, switching logic and interconnect components were expensive thereby limiting their use in systems. However, progress in silicon and packaging technology has changed the relative cost of interconnections and now it is possible to build general purpose computer systems based on switched bus architectures [2]. In this context, the data path used to derive Eq. (1) no longer holds since in the case of these architectures, many paths can be used simultaneously to improve performance.

We believe that the next generation of networked multimedia applications will require more than just network displays: in addition, it should be possible to *manipulate* the multimedia data before storing, displaying or transmitting over the network. In this context, we are exploring a design space that will provide the user with the capability of setting

up data paths between devices and also the flexibility of selecting portions of multimedia data in transit (e.g., from network card to display) and performing any manipulation required.

## Acknowledgments

The authors wish to thank the many employees of Hewlett-Packard laboratories, Bristol, UK for their support and encouragement during the course of this project, in particular we are grateful to Aled Edwards for his valuable discussions on many aspects of this work. We thank Dr. Ulrich Neumann for his help in developing the visualization application. We also thank Kaleb Keithley of The X Consortium for his explanations on the X shared memory extensions. This work was supported by the Integrated Media System Center NSF grant EEC-9529-152 and the DARPA POLO consortium agreement MDA972-94-3-0038.

## References

1. A. Banerjee, D. Ferrari, B.A. Mah, M. Moran, D.C. Venna, and H. Zhang, "The tenet real-time protocol suite: Design, implementation, and experiences," *IEEE/ACM Transactions on Networking*, Vol. 4, No. 1, Feb. 1996.
2. A. Boxer, "Where buses cannot go," *IEEE Spectrum*, pp. 41–45, Feb. 1995.
3. J. Corbet and K. Packard, "The MIT shared memory extension," MIT Consortium, 1991.
4. C. Dalton, G. Watson, D. Banks, C. Calamvokis, A. Edwards, and J. Lumley, "Afterburner," *IEEE Network*, Vol. 7, No. 4, pp. 36–43, July 1993.
5. A. DeBaets and K. Wheeler, "Midrange PA-RISC workstations with price/performance leadership," *Hewlett-Packard Journal*, pp. 6–11, Aug. 1992.
6. Z. Dittia, J. Cox, Jr., and G. Parulkar, "Design of the APIC: A high performance ATM host-network interface chip," in *Proc. of IEEE INFOCOM 95*.
7. P. Druschel and L. Peterson, "Fbufs: A high-bandwidth cross-domain transfer facility," in *Proc. of Fourteenth Symposium on Operating System Principles 1993*, pp. 189–202.
8. C. Frink, R. Hammond, J. Dykstal, and D. Soltis, "High-performance designs for the low-cost PA-RISC desktop," *Hewlett-Packard Journal*, pp. 55–63, Aug. 1992.
9. Hewlett-Packard, *HP Visual User Environment 3.0 User's Guide*, Hewlett-Packard Company, 1992.
10. V. Jacobson, "Efficient protocol implementation," *ACM SIGCOMM Tutorial*, Sept. 1990.
11. V. Jacobson, R. Braden, and D. Borman, "TCP extensions for high performance," *RFC 1323*, May 1992.
12. R. Jones, "Netperf: A network performance benchmark," Revision 1.7 Information Networks Division, Hewlett Packard, March 1993.
13. H. Kanakia and D. Cheriton, "The VMP network adapter board (NAB): High-performance network communications for multiprocessors," in *Proc. ACM SIGCOMM 1988, Symposium on Communication Architectures and Protocols*, pp. 175–187.
14. R. Keller, W. Effelsberg, and B. Lamparter, "Performance bottlenecks in digital movie systems," in *Proc. of the 4th International Workshop on Network and Operating System Support for Digital Audio and Video*, Lancaster House, Lancaster, UK, 1993, pp. 163–174.
15. B.J. Murphy, S. Zeadally, and C.J. Adams, "An analysis of process and memory models to support high-speed networking in a UNIX environment," in *Proc. of Usenix Winter Technical Conference 1996*.
16. J.K. Ousterhout, "Why aren't operating systems getting faster as fast as hardware?," in *Proc. of Usenix Summer Conference 1990*, pp. 247–256.
17. Parallax Hardware Guide, "XVideo700, MultiVideo700, and PowerVideo700," Parallax Graphics Inc., Santa Clara, CA.

18. J. Pasquale, G. Polyzos, E. Anderson, and V. Kompella, "A digital video-conferencing experiment using DECstation 5000 workstation and an FDDI network," Internal Report, Department of Computer Science and Engineering, University of California, San Diego, CA, 1992.
19. J. Pasquale, E. Anderson, and P.K. Muller, "Container shipping—Operating system support for I/O intensive applications," *IEEE Computer*, Vol. 27, No. 3, pp. 84–93, March 1994.
20. D. Saha, D. Kandlur, T. Barzilai, Z. Shae, and M. Willebeek-LeMair, "A video conferencing testbed over ATM: Design, implementation and optimizations," in *Proc. of the International Conference on Multimedia Computing and Systems*, Washington, DC, May 1995, pp. 23–31.
21. G. Watson, D. Banks, C. Calamvokis, C. Dalton, A. Edwards, and J. Lumley, "AAL5 at a gigabit for a kilobuck," *Journal of High Speed Networks*, Vol. 3, No. 2, pp. 127–145, 1994.
22. C.L. Williamson, "Dynamic bandwidth allocation using loss-load curves," *IEEE/ACM Transactions on Networking*, Vol. 4, No. 6, Dec. 1996.
23. D.K.Y. Yau and S.S. Lam, "An architecture towards efficient OS support for distributed multimedia," in *Proc. of Multimedia Computing and Networking*, San Jose, CA, Jan. 1996.



**Sherali Zeadally** is a research assistant professor in the Electrical Engineering Department at the University of Southern California. He received the B.A. degree in Computer Science from University of Cambridge, England, in 1991, and the Ph.D. degree in Computer Science from University of Buckingham, England, in 1996. His current research interests include operating systems internals, distributed systems, high speed networks and multimedia.



**Grig Gheorghiu** is a Ph.D. candidate in the Computer Science Department at the University of Southern California. He received the B.S. degree in Computer Science from the University of Bucharest, Romania, in 1993. His research interests include distributed systems and multimedia applications over high speed networks.



**Anthony F.J. Levi** received the B.S. degree from the University of Sussex, England, in 1980 and the Ph.D. degree from University of Cambridge, England, in 1982. Dr. Levi worked at AT&T Bell Laboratories at Murray Hill, New Jersey, from January 1984 to mid-1993. He was promoted to Distinguished Member of Technical Staff in 1988. In mid-1993 he left AT&T to take up a position as Professor of Electrical Engineering at the University of Southern California. Professor Levi's research interests include scaling of photonic devices to sub-micron dimensions, integration of electronic and photonic devices for high-speed network applications, and optimization of networks for high-sustained throughput applications. To date, he has published over 150 refereed journal papers and holds 9 U.S. patents in these and related research subjects.