



Intel EtherExpress™ PRO/10 Adapter Technical Report

Reducing Bottlenecks for Maximum Performance

By 1995, the percentage of PCs connected to a LAN is expected to grow to over 75%. Applications ranging from spreadsheets and graphics to complex databases and multimedia are placing greater demands on networks than ever before. As the demand for high-performance desktop computing increases, the need to reduce the bottlenecks of the 10Mbps Ethernet network becomes more and more imperative. The new Intel EtherExpress™ PRO/10 adapter, the powerful LAN connection for Intel486™ and Pentium™ processor-based PCs, is designed to reduce 10Mbps bottlenecks, thereby maximizing performance.

Defining Performance

LAN adapter performance can be defined in a number of ways. For example, CPU performance can mean the number of instructions executed per second (MIPs) or the processor clock frequency (MHz). PC performance can refer to the amount of memory (RAM) available or the type of I/O bus (PCI, EISA, ISA) installed. LAN adapter performance is typically defined via two types of measurements - data throughput and host CPU utilization.

Throughput is a measurement of how much data can be sent or received by the adapter card during a specified time interval. Throughput is typically measured in megabits per second (Mbps) or frames per second. The higher the throughput, the better the adapter's performance. On a 10Mbps Ethernet network, it is important to maximize throughput in order to keep up with the rest of the PC's activities.

Host *CPU utilization* is a measurement of how much of the CPU's total time is spent processing network activities. CPU utilization is typically measured as a percentage. The lower the CPU utilization, the better the overall system performance, since the host CPU can spend more time processing other requests.

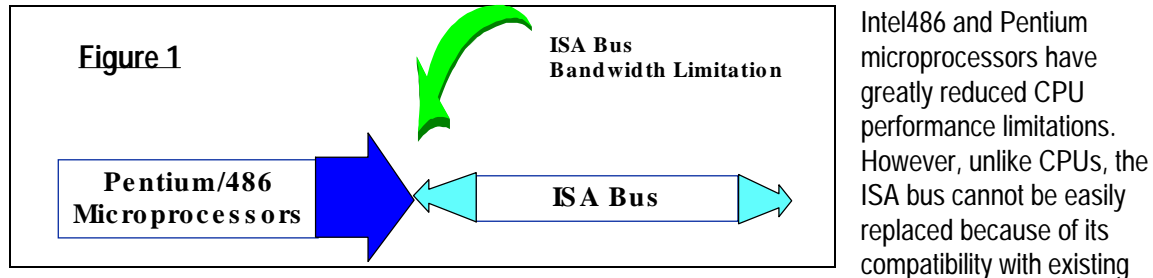
By overlaying these two sets of measurements, it becomes clear that the most high-performance adapters will be those that maximize throughput while maintaining a low CPU utilization.

Network PC Bottlenecks

Bottlenecks may occur when one or more critical data paths are much slower than the rest of the system. These bottlenecks can severely limit performance at a client or server, hampering the system's ability to keep up with network and host PC demands. Minimizing the bottlenecks that occur due to limitations in host system architecture and network adapter designs will yield the best performing system overall.

System Architecture Bottlenecks

Many components contribute to a PC's overall performance, including the host CPU, system memory and the I/O bus. As an integral part of the PC architecture, the I/O bus carries data between the host and various peripheral devices. Today, the ISA bus, originally designed by IBM for the 8088 processor-based PC, is the oldest and most prevalent I/O bus in the PC installed base. It also represents a bottleneck in system architecture (see figure 1).



Faster buses designed to overcome the bus bandwidth limitation, such as PCI and EISA, offer a solution to reducing the I/O bus bottleneck, but most networked PCs today still use the ISA bus for peripherals.

Background - Adapter Bottlenecks

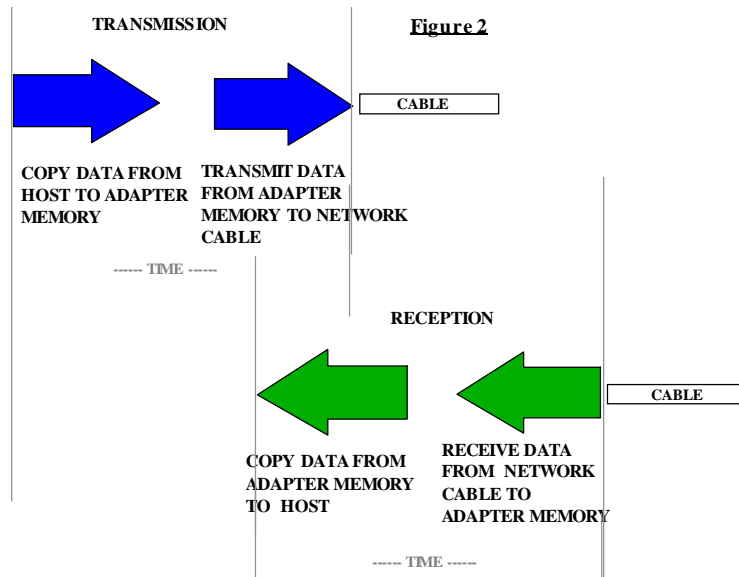
Before discussing LAN adapter bottlenecks, it is important to clarify the differences between bus-master and slave adapter architectures. Bus-master adapters, the more expensive of the two architectures, are generally targeted to servers because they can offload the host CPU. This architecture is ideal for multi-segmented servers which use multiple adapter cards. Slave adapters are generally installed in client PCs because of their excellent price/performance ratio. Many client adapters use efficient software/hardware interfaces, which enable them to perform at par with bus-master adapters in client PCs.

Slave adapters can be classified as either buffered or non-buffered. Buffered slaves (programmed I/O) use local adapter memory to store frames when transmitting or receiving. Non-buffered slaves (shared-memory) use PC memory which is shared between the adapter and the CPU. Many client adapters use buffered slave architecture.

Adapter Bottlenecks


Traditional slave adapters, whether buffered or non-buffered can't transmit or receive data without the assistance of the host CPU. A frame must be completely copied from host memory to the adapter's local memory before the adapter can transmit that frame onto the wire. Similarly, a frame must be completely received from the wire into the adapter's local memory before it can be copied to host memory, as shown in Figure 2.

This sequential process causes significant delays in transmitting and receiving frames, since one step must be completed before another can begin. As a result, a bottleneck forms and data throughput can't use maximum wire bandwidth.



Reducing Bottlenecks with the EtherExpress PRO/10 LAN Adapter

Intel's new EtherExpress PRO/10 Adapter, the powerful LAN connection for Intel486 and Pentium processor-based PCs, maximizes performance by reducing system architecture and adapter bottlenecks. It includes such features as Concurrent Processing with 32k memory and a single 32-bit data port (see figure 3). In addition, the EtherExpress PRO/10 LAN adapter contains a recyclable ring-buffer structure and a single 16-byte I/O address block for a simplified software interface. This enables efficient use of both software and hardware.

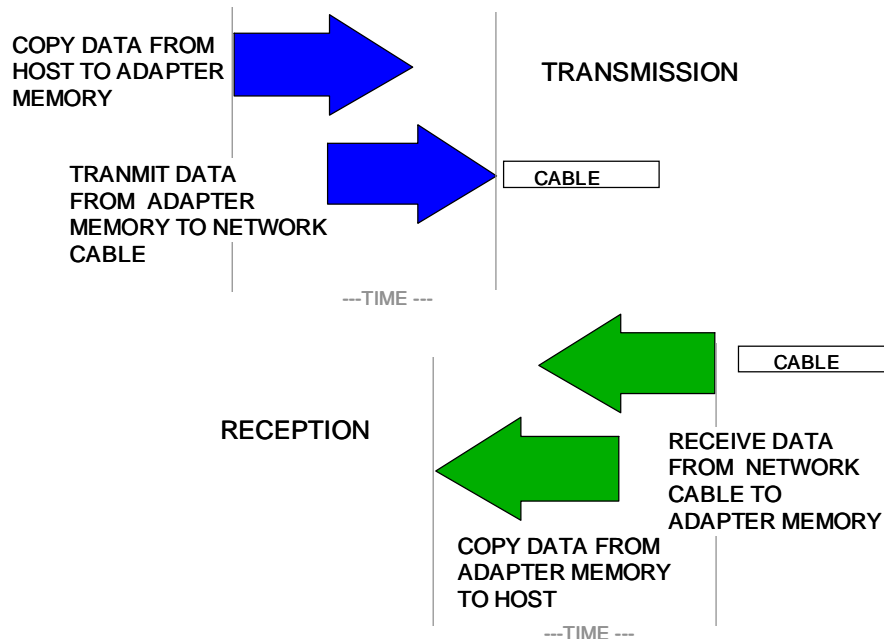


Minimizing the Adapter Bottleneck: Concurrent Processing with 32k Memory

Intel's Concurrent Processing is a new feature incorporated in the Intel EtherExpress PRO/10 LAN adapter to dramatically enhance network performance. It includes early interrupt technology that not only increases throughput, but also minimizes lost frames.

Concurrent Processing technology allows the EtherExpress PRO/10 adapter to begin transmitting a frame onto the wire before that frame is completely copied into local memory (Early Transmission). Similarly, when a frame is received from the wire, the EtherExpress PRO/10 adapter begins copying it to host memory before it is entirely received into local memory (Early Reception). This minimizes delays associated with transmitting and receiving frames. (See figure 4.)

Figure 4



As a result, the EtherExpress PRO/10 adapter can deliver significant performance improvements over traditional slave adapters that must follow a two-step sequential process to transmit or receive a frame.

Early Transmission

During data transmission, underruns can restrict network performance, especially when using early interrupts. A transmit underrun

occurs when the host can't copy data quickly enough to the LAN adapter's local memory. The cause can be a slow CPU or DMA unit, or a slow I/O bus. In the event of an underrun, the adapter reports an error to the network operating system, and then retransmits the frame.

To avoid transmit underruns and successfully transmit a frame onto the wire, several steps must occur during Early Transmission. During initialization, the software driver determines the CPU and I/O speed so the appropriate number of bytes is copied without an underrun. As shown in figure 4, the CPU copies a frame to be transmitted from the host system to the adapter's local memory. The exact size of the frame is calculated by the software driver. Then CPU then issues a transmit command to the adapter. The adapter begins to transmit the first part of the frame while the remaining part of the frame is copied to local memory from the host.

While transmitting data, Concurrent Processing ensures that maximum throughput is achieved. If the host CPU becomes loaded down and underruns persist, the driver may disable Early Transmissions altogether so that optimal performance is not jeopardized. Concurrent Processing also works to deliver maximum performance under various network operating systems - an operation that is done automatically and requires no user intervention.

Early Reception

During reception, receive overruns can also restrict network performance. A receive overrun occurs when data is copied to the adapter's local memory from the wire faster than it can be copied to the host. The cause can be a slow CPU or DMA unit, slow I/O bus, small buffer size, or heavy network traffic.

In order to prevent overruns from occurring, several steps must occur during Early Reception. Concurrent Processing uses predictive early interrupts to signal that key portions of a frame have been received. The first interrupt is generated after a few bytes have been received from the wire. The software driver can then read the frame header, and compute the exact number of bytes it must receive into local memory before it can start copying the frame to the host. The software driver computes several parameters including

EtherExpress PRO/10 Adapter Technical Report - Page 5

interrupt latency, I/O bus speed, CPU speed, and network utilization to precisely compute the appropriate number of bytes.

When most of the frame is in local memory, the second threshold interrupt is generated, and the rest of the data is copied to the host. All this occurs seamlessly to deliver the best performance possible.

The Benefits of 32k Buffer Memory

The advantage of concurrent processing or any other early interrupt technology is that it allows the adapter to transmit and receive data faster than traditional slave architectures. To exploit this advantage, however, it is especially important to minimize other delays. Long interrupt latencies generated by peripherals and multitasking operating systems (such as Windows NT or Chicago) may interfere with the adapter's ability to transmit and receive frames. Heavy network traffic can also cause receive overruns, decreasing adapter performance. The Intel EtherExpress PRO/10 adapter includes a 32K buffer memory - one of the largest available today - to ensure that packets will be processed, sent out and received without extra delays.

Transmission

To understand the performance advantages of 32k buffer memory, consider a network video application. Because video requires a large amount of intensive processing, the host CPU must constantly respond to interrupt requests generated by a video adapter or a software video application.

If the LAN adapter asks the host CPU to process a network-related request, the CPU might not be available for some time, and data transmission will be delayed. As soon as the CPU becomes available, data must be ready so that it can be sent out quickly. But, if the adapter card has a small buffer memory, most of the data will be waiting in host memory. The CPU will need to copy it to the adapter's local memory, adding extra time and burning critical CPU cycles.

By contrast, the EtherExpress PRO/10 adapter's large 32k buffer memory allows most of the data to wait in the adapter. As soon as the CPU is available, the adapter can simply transmit the frame.

Reception

The 32k buffer memory also offers performance advantages during data reception. Consider the case of a LAN adapter receiving a stream of video data from the wire. While it is receiving frames, the adapter attempts to issue an interrupt to the CPU. However, if the CPU is processing an active request by a peripheral device, the adapter can't copy data from its local memory to the host. The buffer then fills to capacity, and incoming packets will be dropped until the CPU processes the adapter's interrupt request.

A small buffer will result in many dropped packets, and the incoming video image will appear choppy. The EtherExpress PRO/10 adapter's 32k buffer ensures maximum performance without dropping frames, allowing images to project smoothly.

Minimizing System Architecture Bottlenecks: An Integrated 32-bit Data Port

Intel's EtherExpress PRO/10 adapter leverages 32-bit technology to maximize performance while maintaining low CPU utilization. To relieve the bottleneck imposed by 16-bit system architecture, Intel used its silicon expertise to integrate a 32-bit data port into the EtherExpress PRO/10 card's 82595TX LAN Controller. This data port enables the CPU to perform double-word, or back-to-back 16-bit access processes between host and local memory. In addition, 32-bit data instructions are incorporated into the initialization, transmission and reception portions of the driver code.


EtherExpress PRO/10 Adapter Technical Report - Page 6

When enabled, the 32-bit data port boosts performance by up to 8 percent. Since 32-bit data instructions need much less code than 16-bit instructions, the CPU won't have to execute as much code, and CPU utilization will decrease. Thus, despite operating as a 16-bit or 8-bit card, the EtherExpress PRO/10 adapter can exploit 32-bit technology to optimize throughput and CPU utilization.

Performance Comparison

To determine the total performance gain provided by the technologies described above, LANQuest Labs performed matched pair benchmark tests using Novell's Perform3 v1.61 with NetWare 4.01 and NETx (test adapter in both the server and a single client). The LANQuest tests yielded the results shown in figure 5.

**Figure 5 LANQuest Labs, July '94
Novell's Perform3 v1.61
Results**



For more information on the benchmark test, please see the July 1994 LANQuest Labs report on Network Adapter Card ISA Performance Benchmarks.

The Intel EtherExpress PRO/10 Adapter: The powerful LAN connection for Intel486 and Pentium processor based PCs

With its superior performance, the EtherExpress PRO/10 LAN adapter is ideally suited to Intel486 and Pentium processor-based PCs. Intel EtherExpress PRO/10 adapters reduce both system architecture and adapter bottlenecks, maximizing use of the 10Mbps Ethernet wire. For more information on the EtherExpress PRO/10 Adapter, please contact Intel or refer to the following Intel FaxBack documents:

EtherExpress PRO/10 adapter data sheet
Differences between EtherExpress PRO/10 and EtherExpress 16 adapters

document #6291
document #6299

Intel486, Pentium and EtherExpress are trademarks of Intel Corporation. Any other trademarks belong to the manufacturers or marketers of the products or processes with which the marks are associated.