

The PCI Local Bus

A Technical Overview



inside front cover

Introduction

Over the past ten years, dramatic advances in microprocessor technology have enabled personal computers to achieve continually higher levels of performance and functionality. Hand-in-hand with these silicon improvements, software applications have become much more complex and sophisticated, providing users with an array of capabilities that previously could not be found in the personal computer platform. Compute-intensive applications such as sophisticated graphics, on-line transaction processing, local-area networking and real-time video now require vast amounts of data to be processed and moved faster than ever before between a personal computer's central processing unit (CPU) and peripheral units.

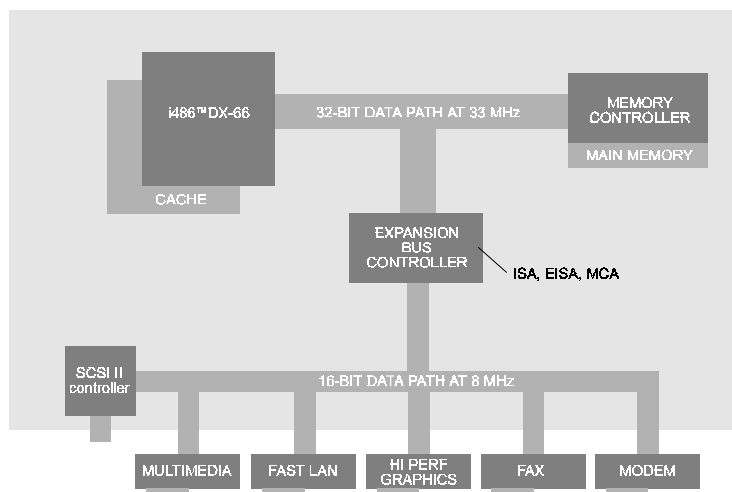
To illustrate this problem, consider the growing bandwidth requirements of graphics and real-time video playback applications. Increasing the resolution or frame size, number of colors and frame rate to achieve higher quality displays also increases the data rate.

| Graphics | | | | | |
|---------------------|---|----------------------|---|---------------------------|-----------------------|
| Resolution (pixels) | | Colors (bits/pixels) | | Redraw Rate (updates/sec) | Bandwidth (bytes/sec) |
| (640 x 480) | x | 8 | x | 10 | = 2.9 MB |
| (1024 x 768) | x | 16 | x | 10 | = 15 MB |
| (1280 x 1024) | x | 24 | x | 10 | = 37.5 MB |

| Real-Time Video Playback | | | | | |
|--------------------------|---|-----------------------|---|-------------------------|-----------------------|
| Frame Size (pixels) | | Colours (bits/pixels) | | Frame Rate (frames/sec) | Bandwidth (bytes/sec) |
| (160 x 120) | x | 8 | x | 15 | = 288 KB |
| (320 x 240) | x | 24 | x | 15 | = 3.5 MB |
| (640 x 480) | x | 24 | x | 30 | = 26.3 MB |

Unfortunately, the industry-standard 16-bit/8-MHz ISA expansion bus - the decade-old bus found in most computers today - cannot adequately transport the large amounts of data generated by these compute-intensive applications. Today's high-performance 32-bit microprocessors, operating at speeds of 33MHz and beyond, are forced to wait as hard disks, video boards and other peripherals send and receive data along a path that is not only narrow and slow, but inefficient as well.

TRADITIONAL BUS ARCHITECTURE



As a result, serious bottlenecks occur in even the most powerful 32-bit Intel486™ microprocessor-based systems as large files vie for the precious limited bandwidth offered by the ISA system bus. Over the years, enhanced bus standards such as EISA and Micro Channel* have been introduced in an attempt to alleviate these concerns, but neither has provided an adequate cost-effective, long-term solution. Users and system manufacturers need a standard bus architecture that not only solves today's problems, but also accommodates tomorrow's emerging applications and standards as well.

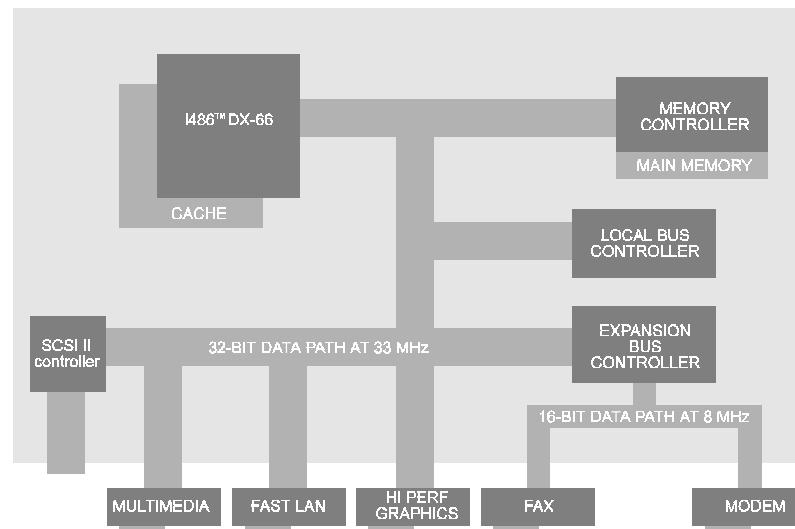
The solution to the technical dilemma is generally acknowledged to be an innovation called the local bus - so named because it brings peripheral functions closer to the microprocessor. The first section of this paper presents an overview of local bus architecture, while the second focuses on the advantages offered by the only forward-thinking of today's local bus solutions - the Peripheral Component Interconnect, commonly referred to as PCI.

Local Bus: A Technology Whose Time Has Come

The CPU bus provides a high speed data path that traditionally has been reserved for the CPU, cache and main memory subsystem. The expansion bus, by comparison, provides the lower speed data path and physical connection between all peripherals and the CPU bus. With the addition of a local bus, peripherals that require rapid transfers of large amounts of data are put in close connection to the CPU and memory subsystem. This allows high-performance peripherals such as video cards, network adapters and disk controllers to bypass the I/O bottlenecks of traditional system buses and take a short cut to the system CPU. In an i486 DX2-66 CPU-based system, for example, the local bus is 32 bits wide and runs at 33 MHz.

Not all peripherals require high data transfer rates in order to perform up to their potential. Lower-bandwidth peripherals such as FAX/modems, tape drives and printers are still well served by the standard expansion bus. Because of that, all of today's local bus systems are designed in conjunction with an expansion bus. In a typical local bus implementation, the ISA expansion bus controller feeds into the local bus as though it were another peripheral, providing the interface for lower performance add-in cards that are configured within the system.

LOCAL BUS ARCHITECTURE



.....

In the past, PC users have had two local bus options to choose from - each of which is integrated onto the motherboard alongside traditional ISA, EISA or Micro Channel buses. The first option consists of a small number of proprietary local bus systems available through individual PC manufacturers. These systems are usually designed for local bus graphics only, which is integrated directly on the motherboard. Without extra local bus slots, the system cannot accommodate multiple high-performance peripherals such as network adapters or multi-media add-in boards. Where slots are designed in, third-party suppliers may be reluctant to develop products not based on industry-standard systems, limiting the availability of compatible add-in cards for these proprietary designs.

The second option is the VESA local bus, or VL bus. This implementation overcomes some of the barriers imposed by OEM-specific proprietary local bus designs, but it also has a number of limitations. VESA's VL specification was developed by add-in board manufacturers to be a fast time-to-market solution. As a result not every VL bus card is compatible with every VL bus-based system. VL users need cards that are specific to both the speed of the microprocessor and the particular expansion bus architecture used in their system - a factor that has limited the availability of VL add-in cards. In addition, high-performance graphics acceleration techniques such as burst write capabilities are not supported on today's VL local bus systems. And lack of forward-thinking capabilities such as auto-configuration support will likely prevent the VL bus from ever becoming a long-term local bus standard.

In light of today's technological needs, a well-designed local bus must offer much more than high bandwidth. It must also allow peripheral functions to take full advantage of available processing power without being dependent on the CPU speed or architecture. It should provide users with ease-of-application features, while offering system designers an easy design path. And it must achieve wide acceptance as an industry standard that not only meets today's needs, but also provides built-in upgradability to accommodate future technical advances.

PCI Local Bus: Fueling the Future

Originally proposed by Intel Corporation's Architecture Lab in late 1991, the PCI local bus specification has risen to prominence in a very short period of time. In June of 1992, Intel and other leading companies in the computer industry formed the PCI Special Interest Group (SIG) to promote, oversee and enhance the development of PCI as an open, non-proprietary local bus standard. More than 160 companies - including a mix of semiconductor suppliers, BIOS vendors, computer and add-in board manufacturers such as Madge Networks - are now active members of the PCI SIG, reflecting the widespread support the specification is gaining throughout the industry.

PCI is an advanced high-performance local bus that supports multiple peripheral devices. Working as a processor-independent bridge between the CPU and high-speed peripherals, it accelerates data throughput by acting as a traffic controller between buses. As a highly integrated local bus architecture, PCI is optimized to take advantage of today's available microprocessor and personal computer technology. It ensures reliable operation between components, add-in cards and systems while maintaining compatibility with today's existing ISA/EISA/Micro Channel expansion buses.

.....

The PCI local bus specification offers a number of key benefits:

- High Performance
- Compatibility
- Processor Independence
- Platform Flexibility
- Cost Effectiveness
- Future Support

High Performance

Unlike other local buses - which may be configured only for speeding up graphics or video operations - the PCI local bus is a total system solution. It provides increased performance for network adapters, hard disk drives, full-motion video, graphics and the wide range of today's high-speed peripherals. Running at a clock speed of 33 MHz, the PCI local bus employs a 32-bit data bus that supports multiple peripheral components and add-in cards at a peak bandwidth of 132 MB/second - a substantial improvement over the 5 MB/second transfer rate of the standard ISA bus. Even in a 32-bit implementation, PCI supports the fast graphics data rates found in Pentium™ processor-based systems.

High bandwidth is not the only aspect that contributes to PCI's improved performance. The following features of the PCI specification have been designed to keep the bus filled with data, minimize CPU wait-states and allow multiple operations to be executed simultaneously:

Linear Bursts

PCI supports a method of transferring data called linear bursts which ensure that the bus is continuously filled with data. Peripheral devices expect to receive data from the system's main memory in linear address order. This linear addressing means that large amounts of data - or burst - are read from - or written to - a single address, which is then automatically incremented for the next byte in the stream. Linear bursts allow more bandwidth to be used to actually send data instead of needless addresses.

In addition, PCI is unique in that it supports both burst reads and burst writes. This is particularly important when using high-performance graphics accelerators, where more than 90% of CPU data accesses are writes to the frame buffer from main memory.

Low Access Latency


Devices designed to support PCI have low access latency, reducing by more than an order of magnitude the time required for a peripheral to be granted control of the bus after requesting access. For example, an Ethernet* controller card connected to a LAN may have large data files from the network coming into its buffer. Waiting for access to the bus, the Ethernet card is unable to transfer file's data to the CPU quickly enough to avoid a buffer overflow, forcing it to temporarily store the file's contents in extra RAM. Since PCI-compliant devices support faster access times, the Ethernet card can promptly send data to the CPU, eliminating the need for extra memory and reducing overall add-in board costs.

Bus Mastering and Concurrency

Performance enhancements are also achieved through PCI's support for bus mastering and concurrency. Found on most buses, bus mastering allows any one of a number of intelligent peripherals to take control of the bus in order to accelerate high-throughput, high-priority tasks. Unique to PCI, the concurrency capability ensures that the microprocessor operates simultaneously with these masters, instead of waiting for them. For example, concurrency allows the CPU to perform floating-point calculations on a spreadsheet application while operations continue between an Ethernet card and the local area network.

Compatibility

Because PCI has been designed as a supplement to the expansion bus standards, it is completely compatible with the ISA, EISA and MicroChannel buses. And although every system has a limited number of slots, the PCI local bus specification allows system manufacturers to offer 'shared slots' that can accommodate both a PCI and ISA, EISA or Micro Channel card connector. Finally, any PCI compatible add-in board will function on any PCI-compliant system, regardless of the type of expansion bus or microprocessor in use. These compatibility features protect user investments by allowing the continued use of existing add-in cards, providing extra slot options and easing the selection of new peripherals.



Processor Independence

PCI's processor-independent architecture features a unique intermediate buffer design between the CPU subsystem and the peripherals. Typically, adding more devices or components directly to the CPU bus degrades performance and decreases reliability. Buffering allows users to expand their systems by adding multiple peripheral devices without inducing performance variations at different clock speeds.

Processor-independent design ensures that changes in processor technologies do not render any one particular system design obsolete. It enables the PCI local bus to maintain compatibility with the entire Intel architecture, including the full array of Intel486 CPUs, the Pentium processor, the OverDrive™ processors and all future generations of the Intel architecture. This forward and backward compatibility benefits end users by reducing the costs facing system manufacturers - savings that ultimately, over time, are passed on to consumers.

Flexible Platform Support

The PCI local bus provides a cost-effective local bus solution not only for standard desktop caliber systems, but also for mobile computers and servers. It provides desktop graphics performance for portables and notebooks, while supporting 3.3-volt operation to pave the way to longer battery life and smaller, lighter computers. PCI local bus design conserves board real estate by minimizing package sizes and chip counts, enabling system designers to pack more functions into their products.

In the server environment, PCI's hierarchical peripheral support enables a single PCI interface to support a cascaded PCI local bus. This allows servers designed with multiple PCI buses to offer the added benefits of more I/O connections added expansion slots and the isolation of high bandwidth traffic from lower bandwidth traffic.

Cost Effectiveness

The PCI specification was originally developed to help reduce the overall costs associated with system design - an advantage that over the long term will ultimately be realized by end users. It uses highly integrated PCI chips to incorporate system functions such as DRAM and cache controllers, eliminating the costs and board space associated with extra support circuitry commonly referred to as glue logic. Unlike VESA's VL components, which have more than 80 pins, PCI components save circuit board real estate by multiplexing the address and data bus lines. This reduces to less than 50 the number of pins needed to interface with other components. Also, PCI reduces the costs incurred when add-in board manufacturers must design cards for every CPU speed or architecture.

Future Support

One of the most important attributes of the PCI specification is that it has been developed with an eye to the future. For one, it supports a multiplexed 64-bit data and address but for the emerging wave of high-performance peripherals that inevitably will require such a wide data path. This 64-bit extension of PCI's 32-bit data and address buses can double the system's overall bus


bandwidth to 264 MB/second. Equally as important, all communications between 32- and 64-bit peripherals are transparent to users. This transparency is provided by a connector that accepts both 32- and 64-bit cards, ensuring backward and forward compatibility.

Another forward-thinking advantage of PCI is its provision for auto-configuration capabilities. The PCI SIG is working with the industry to ensure that users can install a new peripheral board without having to manually configure jumpers and set DIP switches or interrupts. Configuration registers built into each PCI component are set up in the resident software that runs at the time of system initialization, automatically defining the address and interrupt setting the peripheral will use to communicate with the microprocessor. When a new peripheral is added, configuration software selects an unused interrupt, ensuring that no conflicts will occur when installing multiple add-in boards.

Performance Benchmarks


Making the distinction between bus bandwidth and throughput helps to illustrate the fact that a well-designed local bus is not just a wide and fast data path, but also an efficient one as well.

Bandwidth relates to how fast the bus can move data. Throughput gives an indication of how well the features of the bus match the problems the bus will encounter during I/O operations.



Torque is a performance measurement utility that illustrates how PCI's unique burst write and linear bursting capabilities substantially improve graphics-oriented operations. In a memory-to-screen test, it measures the rate at which pixels are transferred from main memory to screen. Although it measures the number of pixels/second, the test was performed in 256-color mode (8 bits/pixel) and is, therefore, equally represented in bytes/second.

Using Ziff Davis Lab's WinBench* benchmark v3.11, the PIC local bus achieved a Graphics WinMark* test score of 39.97. WinMarks, measured in millions of pixel/second, are widely



recognized by users as a measure of a system's graphics performance. WinMark test results are dependent on many factors, including the CPU, bus architecture, and the addition of graphics accelerator card with drivers optimized to run applications under Windows.* In these tests, Intel used the highest performing accelerator cards that were available from a leading graphics card manufacturer for the PCI, VESA and ISA buses.

All systems used during the testing were identically configured with an Intel486™ DX2-66 CPU, a 256K L2 cache, 4MB RAM, DOS 5.0/Windows 3.1 and IDE hard drive. Three of the four systems used a version of ATI's Mach32 graphics accelerator card while the fourth used a standard VGA controller. The PCI local bus system used Intel's PCI components.

PCI: A Standard for Today and Tomorrow

Because the PCI Local Bus specification accommodates today's technical requirements as well as satisfying tomorrow's, it is acknowledged as the most forward-thinking local bus architecture in the industry. PCI's high performance, efficient design, compatibility with existing standards and support for future requirements offer significant advantages when compared to all other expansion and local bus implementation in use today. Its well defined specifications and widespread industry support ensure that add-in board manufacturers have a clear design path to follow for the present and future. And PCI's processor independence opens the door to a range of designs spanning all segments, of the PC platform, including desktop computers, notebook PCs and server applications.

The industry is already beginning to see tangible product implementations of the PCI local bus standard. The first wave of PCI add-in cards and PCI-based computer systems will be appearing on the market by mid-1993, with many more soon to follow. As processing power increases and new compute-intensive advanced applications continue to emerge, PCI's forward thinking architecture will increasingly benefit users as the personal computer local bus standard of the 1990s and beyond.

Inside back cover

Amsterdam
 Atlanta
 Bangkok
 Berlin
 Boston
 Brussels
 Cape Town
 Chicago
 Cologne
 Copenhagen
 Dallas
 Denver
 Detroit
 Frankfurt
 Hong Kong
 Jakarta
 Johannesburg
 Kuala Lumpur
 London
 Los Angeles
 Madrid
 Minneapolis
 Munich
 Nashville
 New York
 Oslo
 Paris
 Philadelphia
 San Jose
 Seattle
 Singapore
 Stockholm
 Sydney
 Tokyo
 Toronto
 Vancouver
 Washington DC



Madge Europe

Madge Networks Ltd
 Loudwater
 High Wycombe
 Bucks HP10 9QZ
 England
 Tel: +44 1628 858000
 Fax: +44 1628 858011

Madge Americas

Madge Networks Inc
 2310 North First Street
 San Jose
 Calif. 95131-1011
 United States
 Tel: +1 408 955 0700
 Fax: +1 408 955 0970

Madge Asia

Madge International Ltd
 64-01 Central Plaza
 18 Harbour Road
 Wanchai
 Hong Kong
 Tel: +852 2593 9888
 Fax: +852 2519 8022

Madge Japan

Madge Japan KK
 Believe Mita
 43-16 Shiba 3-chome
 Minato-ku, Tokyo 105
 Japan
 Tel: +81 3 5232 3281
 Fax: +81 3 5232 3208

This piece has been reproduced with the kind permission of the Intel Corporation.

Madge, the Madge logo, Fastmac, Ringhub, Ringnode, Smartboot, SmartCAU, SmartLAM, Smartrom and Straight Blue are registered trademarks of Madge Networks Limited. Bridgenode, Plus Madge, Ringbridge, RingRunner, Smart Networking, TechNote, TrueView and WireSense are trademarks of Madge Networks Limited. All trademarks and registered trademarks acknowledged.

Copyright © 1994 Madge Networks Limited