

PC TODAY!

PC Technology, the longest surviving technology, has come of age! PC TODAY has unparalleled performance. Thanks to the emerging enhancements. Take any PC today; it talks of EDO RAM, ECP/EPP, PnP, DMI.....What do these terms mean? Why do we need these enhancements? What are the advantages of these enhancements? Let us take a closer look at them:

CPU TECHNOLOGY

Let us start with CPU architecture. The first thing that comes to our mind is CISC vs. RISC. Today, there is very little difference between the two. Both have Complex Instruction set & both can execute instructions faster. Recent Intel Processors that are classified as CISC processors have a number of RISC features. Some of the design features of the latest CPU from Intel, Pentium Pro, proves this point.

- ?? The new Pentium Pro has Integrated and closely coupled L2 cache with its innovative *dual-cavity pin-grid array* packaging. One can see two compartments on the chip. One cavity is the processor and other one 256KB secondary cache. With this, the Pentium Pro has the advantages of both L1 cache (On-chip cache) and L2 cache (Secondary on-board cache). (Our HP PA-RISC architecture has similar feature.) A dedicated 64-bit bus connects the L2 cache to the processor to ensure that it runs at the same speed as that of the processor. The bus has Error Correcting and Checking (ECC), and diagnostic features. This dramatically simplifies system design.
- ?? Pentium Pro heralds a new generation of processors with *three-way superscalar* and *14-stage super-pipelined* architecture. Three-way superscalar means that three to five x86 instructions can be processed in a single clock cycle as compared to two in Pentium. In Pentium, Super-pipelining was done in five stages while a Pentium Pro does it in 14 stages.
- ?? Another design feature of Pentium Pro is *Out-of-order execution*. So when a stream of instructions flows in, Pentium Pro, unlike a Pentium, picks up the easier ones for execution instead of taking instructions one by one. This again helps in throughput.
- ?? Pentium Pro has algorithm for *Branch prediction* for jump instructions. If the prediction is right, time is saved else the processor cleans the pipeline of instructions that got partially executed by mistake. That is not all; Pentium Pro goes a step beyond with *Speculative execution*. That means, it executes instructions even before the condition that specifies the branch is confirmed to have been satisfied. What if it turns out to be wrong? Nothing, just flush out the speculatively executed instructions. The hit rate is pretty high resulting in higher gain.
- ?? Another RISC feature on this new Intel chip is *Shadow registers*. When two instructions contend to use the same register, clones of physical registers are created and used. That is, virtual registers are created as and when they are needed. This technique is invaluable for 'out-of-order-execution' to be fully realised.

MEMORY TECHNOLOGY

EDO RAM (Extended Data-Out) is the latest buzz word that promises *exceptionally fast memory access*. EDO RAM reduces memory access time significantly compared with standard DRAM chips and costs only a little more to manufacture. The accesses to consecutive locations in memory is speeded up by (1) assuming that the next memory access will target an address in the same transistor row as the previous one and (2) latching data at the output of the chip so it can be read even as the inputs are being changed for the next memory location. With extra circuitry that fetches in advance of a data request, EDO RAM is good for systems running at 66 MHz and above. For future releases of Pentium Pro that execute three instructions per clock cycle, it will be necessary to feed the processor cache at a higher speed.

DIMM stands for *Double Inline Memory Module*. This is superior to SIMM based architecture in terms of data integrity. It gives 10% higher performance than SIMM.

Flash RAM is a new twist to EEPROM. Instead of requiring special, higher voltages to be erased, Flash RAM can be erased and reprogrammed using the normal voltages inside a PC. For system designers, that makes Flash RAM easy to use. Unfortunately, Flash RAM is handicapped by the same limitation as EEPROM -- its life is finite (although longer than ordinary EEPROM) and (in most but not all cases) it must be erased and reprogrammed as a block.

DISK TECHNOLOGY

MASTER ENHANCED IDE:

The Enhanced IDE was designed to overcome the limitations of standard IDE. It has many features: Support for other devices such as CD-ROM drives; Support for Multiple devices (up to four); Support for drives of up to 8.4 GB in size; Support data transfer rates of up to 16.7 MB per second.

Enhanced IDE does not allow multiple data transfers on the same bus. When a system with an IDE interface requests data from the disk, the disk interface cannot accept another request until the disk returns the data from the first. This may take 15 or 20 ms, which is a relatively long time in processor terms.

Master Enhanced IDE has been proposed to overcome this. Master IDE involves off-loading responsibility for data transfers from the processor, freeing up the processor during transfers and *allowing multiple simultaneous data transfers* when using multitasking operating systems.

SCSI UPDATE

SCSI has some flavours. SCSI-2 Standard was developed to allow more flexibility and higher performance. SCSI-2 extends the original specification to include support for devices such as CD-ROMs, scanners, communications devices, and optical memory devices.

SCSI-2 supports faster and wider bus than original. In addition to the standard 8-bit SCSI bus, you can have 16-or 32-bit buses. SCSI-2 provides for speeds of 10 Mbps on the 8-bit bus, which means you can get up to 40 Mbps on the 32-bit bus. SCSI-2 implementations that support the 10-Mbps/8bit speed are called *FAST SCSI*. Implementations that use the 16-or 32-bit bus are called *WIDE SCSI*. You may also find SCSI-2 with both (fast and wide SCSI) or neither (the slow and narrow 5-Mbps 8-bit bus). Ultra SCSI is a SCSI-2, FAST & WIDE interface delivering 40Mbps.

Work has already begun on its successor, the SCSI-3 specification, which will address some of SCSI-2's limitations. The plans include support for optical fiber, longer cables, and more than eight targets per bus. It will also be faster than the SCSI-2 bus, running at 20 Mbps.

PCI IMPLEMENTATIONS

PCI is an open, non-proprietary local bus standard that was originally proposed by Intel in late 1991 and later embraced by other leading companies. The PCI bus is a 32-bit bus running at up to 33 MHz. A transparent 64-bit extension of the 32-bit data and address buses is defined, effectively doubling the bus bandwidth and offering forward and backwards compatibility with 32-bit and 64-bit PCI peripherals.

The electrical characteristics of PCI are rigidly controlled which allows the bus to run extremely fast with very low noise. This puts limitations on the load that can be put on the bus. The solution is to extend the PCI bus, using one of the two implementations.

PCI-to PCI BRIDGE

In this implementation, the system bus is connected to the first PCI bus via a system PCI bridge. The first PCI bus is then connected to the second PCI bus over a PCI to PCI bridge chip. The EISA bus is connected to the first bus over a PCI-to-EISA bridge chip.

The advantage of this implementation is not only the extension of PCI bus, but also the economy since it requires only one chip to talk to the system bus. HP Netserver LH has this implementation.

DUAL PCI PEER

In this implementation two PCI buses are connected directly to the system bus. Apart from PCI bus extension, this also improves performance by removing the bridge latency problem and by doubling the bandwidth of available PCI-to-CPU transfers because it offers two pipes instead of one. This, for some users, outweighs the slight cost premium incurred. HP Netserver L employs this implementation.

I/O UPDATE

ENHANCED PARALLEL PORT

This bi-directional port conforms to the ECP/EPP standard developed by the IEEE 1284. This transforms a parallel port into an expansion bus that can handle up to 64 disk drives, CD-ROM drives, and other mass-storage devices. The maximum signal rate is 16 Mbps. With this faster port, external LAN cards are more viable. EPP is rapidly gaining acceptance as inexpensive means to connect portable drives.

ENHANCED SERIAL PORT

The new PC serial ports have UART 16550 that support a 16 byte FIFO. Older PC with UART 16450 has no FIFO. The absence of FIFO does not degrade throughput, but generates more interrupts at the CPU.

Let us take a typical situation to witness the impact of the FIFO. Consider a PC receiving data through the serial port and having high interrupt activity (e.g. mouse moving from one application window to another). If the system is running several applications, with many interrupts, the serial port will sometimes not be given the CPU's attention in time to transmit the characters to the screen. It will be the hardware cursor that monitors the mouse movement that will be given the priority.

INFRARED TALK

A notebook technology IRDA (InfraRed Data Association) standard has moved on to desktops to transfer data or print. With this wireless port, just point and shoot !The point-to-point data transfer rate of this point-and-squirt business is around 1150 kbps.

DMI

Desktop Management Interface, is the interface created to standardise communication to and from PCs. The DMI standard is the industry's response to the pressing need for PCs to communicate their intelligence through corporate network and to become fully integrated and manageable elements in the enterprise information system. The DMI interface is composed of a service layer plus two sets of application programming interfaces (APIs) - the Management interface and the Component interface. These interfaces allow software developers and hardware manufacturers to write software to a standard format that can be

read and relayed by the DMI service layer.

DMI-enabled PCs are able to work in concert with management applications to enable systems managers to provide better end-user service, greater up-time and more intelligent management (Software distribution, trouble tickets, automatic inventory) by providing real-time information about the specific PC and its elements (Hard disks, Word Processors, Motherboards, Operating Systems, etc). With DMI-enabled PCs there is no more guessing about what hardware or software is installed on a PC, where faults are on the PC and the system administrator can actually change elements of the PC. This means, for instance, that a user who needs help can be aided on-screen, that software can be distributed according to the PC's real configuration or that floppy and hard disk drives or serial ports.

The members of the Desktop Management Taskforce (DMTF) cover almost all major PC vendors and over 400 software developers as members. The DMTF was founded by DEC, HP, IBM, Intel, Microsoft, Novell, Sun and Synoptics in 1992 with the goal of creating DMI.

How does DMI differ from current network management software? DMI is an enabling technology (software) that allows management applications to manage PCs - it provides access to the information imbedded in a PC and allows certain changes to this information. It is not a management application.

The availability of DMI standard will allow management applications software companies to modify/develop DMI-complaint applications that can manage any DMI-complaint PC, provide more information defining the desktop environment, exchange information and instructions with any other DMI-complaint application.

This opens the management of PCs to all standard management applications; they merely need to call for information, or act upon it, through the DMI API.

Does DMI replace SNMP? DMI and SNMP are complimentary. SNMP (Simple Network Management Protocol) is the protocol designed to manage network infrastructure elements such as routers, hubs and bridges whereas DMI addresses desktop machines such as printers, personal computers and servers.

PLUG N PLAY

Plug and Play are a set of hardware and software design specifications. Plug n Play allow users to add new devices to their systems effortlessly. A system is said to be *Plug n Play* when all three system components- BIOS, OS and add-on cards, have Plug and Play functionality. In a Plug n Play compliant system, devices will be uniquely and positively identified, their resource requirements read, and their drivers loaded. If any change is made to the system configuration, the process will be repeated.

The Plug n Play architecture must provide compatibility with the installed base of existing non-plug and play systems and peripherals, and be extensible to new classes of devices that may arise in the future.