



**FACULTAD DE INGENIERIA U.N.A.M.
DIVISION DE EDUCACION CONTINUA**

**CENTRO DE INFORMACIÓN Y DOCUMENTACIÓN
“ ING. BRUNO MASCANZONI “**

El Centro de Información y Documentación Ing. Bruno Mascanzoni tiene por objetivo satisfacer las necesidades de actualización y proporcionar una adecuada información que permita a los ingenieros, profesores y alumnos estar al tanto del estado actual del conocimiento sobre temas específicos, enfatizando las investigaciones de vanguardia de los campos de la ingeniería, tanto nacionales como extranjeras.

Es por ello que se pone a disposición de los asistentes a los cursos de la DECFI, así como del público en general, los siguientes servicios:

- 
- **Préstamo interno.**
 - **Préstamo externo.**
 - **Préstamo interbibliotecario.**
 - **Servicio de fotocopiado.**
 - **Consulta a los bancos de datos: librunam, seriunam en cd-rom.**

Los materiales a disposición son:

- **Libros.**
- **Tesis de posgrado.**
- **Publicaciones periódicas.**
- **Publicaciones de la Academia Mexicana de Ingeniería.**
- **Notas de los cursos que se han impartido de 1988 a la fecha.**

En las áreas de ingeniería industrial, civil, electrónica, ciencias de la tierra, computación y, mecánica y eléctrica.

El CID se encuentra ubicado en el mezzanine del Palacio de Minería, lado oriente.

El horario de servicio es de 10:00 a 14:30 y 16:00 a 17:30 de lunes a viernes.



**FACULTAD DE INGENIERIA U.N.A.M.
DIVISION DE EDUCACION CONTINUA
A LOS ASISTENTES A LOS CURSOS**

Las autoridades de la Facultad de Ingeniería, por conducto del jefe de la División de Educación Continua, otorgan una constancia de asistencia a quienes cumplan con los requisitos establecidos para cada curso.

El control de asistencia se llevará a cabo a través de la persona que le entregó las notas. Las inasistencias serán computadas por las autoridades de la División, con el fin de entregarle constancia solamente a los alumnos que tengan un mínimo de 80% de asistencias.

Pedimos a los asistentes recoger su constancia el día de la clausura. Estas se retendrán por el periodo de un año, pasado este tiempo la DECFI no se hará responsable de este documento.

Se recomienda a los asistentes participar activamente con sus ideas y experiencias, pues los cursos que ofrece la División están planeados para que los profesores expongan una tesis, pero sobre todo, para que coordinen las opiniones de todos los interesados, constituyendo verdaderos seminarios.


Es muy importante que todos los asistentes llenen y entreguen su hoja de inscripción al inicio del curso, información que servirá para integrar un directorio de asistentes, que se entregará oportunamente.

Con el objeto de mejorar los servicios que la División de Educación Continua ofrece, al final del curso deberán entregar la evaluación a través de un cuestionario diseñado para emitir juicios anónimos.


Se recomienda llenar dicha evaluación conforme los profesores impartan sus clases, a efecto de no llenar en la última sesión las evaluaciones y con esto sean más fehacientes sus apreciaciones.

**Atentamente
División de Educación Continua.**

PALACIO DE MINERIA

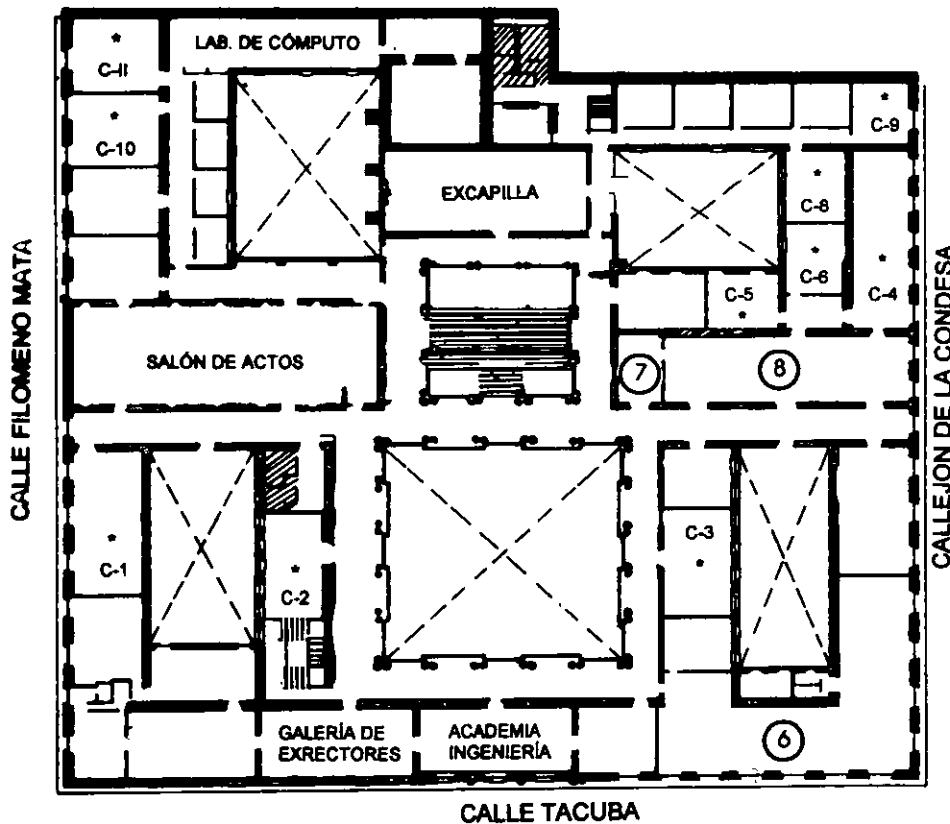


PLANTA BAJA



MEZZANINNE

PALACIO DE MINERÍA



1er. PISO

GUÍA DE LOCALIZACIÓN

1. ACCESO
2. BIBLIOTECA HISTÓRICA
3. LIBRERÍA UNAM
4. CENTRO DE INFORMACIÓN Y DOCUMENTACIÓN "ING. BRUNO MASCANZONI"
5. PROGRAMA DE APOYO A LA TITULACIÓN
6. OFICINAS GENERALES
7. ENTREGA DE MATERIAL Y CONTROL DE ASISTENCIA
8. SALA DE DESCANSO

SANITARIOS

* AULAS



DIVISIÓN DE EDUCACIÓN CONTINUA
FACULTAD DE INGENIERÍA U.N.A.M.
CURSOS ABIERTOS

DIVISIÓN DE EDUCACIÓN CONTINUA





**DIVISION DE EDUCACION CONTINUA
FACULTAD DE INGENIERIA, UNAM
CURSOS ABIERTOS**

DIVISION DE EDUCACIÓN CONTINUA



CURSO: CC36 MANTENIMIENTO DE PC'S Y PERIFERICOS (PARTEII)

FECHA: 5 al 16 de junio del 2000

EVALUACIÓN DEL PERSONAL DOCENTE

(ESCALA DE EVALUACIÓN: 1 A 10)

CONFERENCISTA	DOMINIO DEL TEMA	USO DE AYUDAS AUDIOVISUALES	COMUNICACIÓN CON EL ASISTENTE	PUNTUALIDAD
ING. JUAN CARLOS MAGAÑA C.				

Promedio _____

EVALUACIÓN DE LA ENSEÑANZA

CONCEPTO	CALIF.
ORGANIZACION Y DESARROLLO DEL CURSO	
GRADO DE PROFUNDIDAD DEL CURSO	
ACTUALIZACIÓN DEL CURSO	
APLICACIÓN PRACTICA DEL CURSO	

Promedio _____

EVALUACIÓN DEL CURSO

CONCEPTO	CALIF
CUMPLIMIENTO DE LOS OBJETIVOS DEL CURSO	
CONTINUIDAD EN LOS TEMAS	
CALIDAD DEL MATERIAL DIDÁCTICO UTILIZADO	

Promedio _____

Evaluación total del curso _____

Continúa...2

1. ¿Le agradó su estancia en la División de Educación Continua?

SI NO

Si indica que "NO" diga porqué:

2. Medio a través del cual se enteró del curso:

Periódico <i>La Jornada</i>	<input type="checkbox"/>
Folleto anual	<input type="checkbox"/>
Folleto del curso	<input type="checkbox"/>
Gaceta UNAM	<input type="checkbox"/>
Revistas técnicas	<input type="checkbox"/>
Otro medio (Indique cuál)	<input type="checkbox"/>

3. ¿Qué cambios sugeriría al curso para mejorarlo?

4. ¿Recomendaría el curso a otra(s) persona(s) ?

SI NO

5. ¿Qué cursos sugiere que imparta la División de Educación Continua?

6. Otras sugerencias:



FACULTAD DE INGENIERIA U.N.A.M.
DIVISION DE EDUCACION CONTINUA

MANTENIMIENTO DE PC'S Y PERIFERICOS

(PARTE II)

PRESENTACIÓN

JUNIO 2000

MANTENIMIENTO DE PC'S/XT'S Y PERIFERICOS

PRESENTACION

Temas

- 1 ARQUITECTURA DE UNA COMPUTADORA
- 2 CARACTERISTICAS Y DIFERENCIAS ENTRE LOS SISTEMAS
- 3 TALLER DE CONFIGURACIONES
- 4 TALLER DE DIAGNOSTICOS
- 5 TALLER DE ENERGIA
- 6 MANTENIMIENTO PREVENTIVO LOGICO
- 7 MANTENIMIENTO CORRECTIVO BASICO
- 8 CONTRATOS DE MANTENIMIENTO

Virtualmente las necesidades del mundo actual, fácilmente han convertido al ciudadano común en usuario de una computadora personal, ya que cualquier empresa u organismo por pequeños o grandes que sean, requieren ya, de esta tan indispensable herramienta. Lo mismo está sucediendo en aquellos hogares donde el estudiante o el hombre de negocios engrosa el "mundo de usuarios de PC's"

La aparición de las "Pcs" a partir de las XT 8088, se puede decir que es un hecho relativamente reciente, (sólo 15 años), la dinámica evolución de los mismos pasando por las AT's (80286, 80386, hasta el popular Pentium), no ha bastado para que gran parte de las comunidades de usuarios, haya desarrollado los hábitos adecuados para el uso y mantenimiento de sus PC's. En la actualidad es necesario que estos hábitos se generen casi paralelamente al mayor conocimiento de sus equipos, a efecto de que puedan prolongar la vida útil de tan valiosa herramienta.

El auge en el uso de las computadoras "PC's" compatibles sigue en aumento y se está extrapolando hacia los equipos de mayor capacidad y mejor tecnología, como: (80486, PENTIUM, RISC, ALPHA, POWER PC, ETC.). Es evidente que los equipos PC's, XT (8088, 8086), y AT (80286,80386) aun "están vivos" y se les debe prestar la atención que estas herramientas merecen.

La mayoría de los problemas que disminuyen la vida útil de una computadora y de sus

periféricos, pueden ser prevenidos, si se ejerce un mantenimiento continuo y adecuado. Casi cualquier usuario con un poco de experiencia, puede cubrir este ejercicio de mantenimiento, si su capacitación sin ser la más exigente, es la mínima y bien orientada. y con base en ella, sepa cuándo debe acudir a un especialista.

La aseveración anterior es uno de los objetivos de este curso teórico-práctico, mismo que después de la introducción, muestra la arquitectura básica de una microcomputadora, luego menciona las características y diferencias físicas entre los equipos -tanto microcomputadoras como sus periféricos, más populares que existen actualmente en el mercado, así como las arquitecturas y tecnologías de punta, como ISA, MCA, EISA, LBUS, VESA, PCI, etc Para lograr un mejor conocimiento del mantenimiento adecuado, se establece el "A,B,C," del mantenimiento preventivo, complementando esta materia, con la práctica y aplicaciones de los programas de diagnóstico existentes y se especifica el juego de herramientas básico tanto físicas como lógicas que se requieren. No debiendo soslayar la parte administrativa-legal, que genera un servicio de mantenimiento contratado, para finalizar, se analiza con sus pros y contras, algunos de los diferentes contratos de mantenimiento que se ofrecen en el mercado

MANTENIMIENTO DE PC'S/XT'S Y PERIFERICOS

2000

Temas

- 1** ARQUITECTURA DE UNA COMPUTADORA
- 2** CARACTERÍSTICAS Y DIFERENCIAS ENTRE LOS SISTEMAS
- 3** TALLER DE CONFIGURACIONES
- 4** TALLER DE DIAGNOSTICOS
- 5** TALLER DE ENERGIA
- 6** MANTENIMIENTO PREVENTIVO LOGICO
- 7** MANTENIMIENTO CORRECTIVO BASICO
- 8** CONTRATOS DE MANTENIMIENTO

OBJETIVOS

Lograr que los participantes después del curso, puedan alargar la vida útil de sus equipos, aplicando las acciones y cuidados del mantenimiento preventivo, evitando así problemas en potencia. Destruir los frecuentes mitos de que una computadora es "intocable por profanos" y que sólo los "GURUS" son los elegidos. Ayudar al usuario a decidir cuando puede resolver un problema por si mismo y cuando debe acudir a un especialista, y en general, dotarlo de seguridad en el manejo y cuidado de su PC y equipo.

A QUIEN VA DIRIGIDO

A profesionales, técnicos y usuarios en general, de microcomputadoras compatibles que por sus necesidades, requieran dar por si mismos, mantenimiento a sus equipos

REQUISITOS

Dado que es un curso especializado es deseable que el aspirante tenga el perfil adecuado:

Profesional o técnico en electrónica

Manejo teórico-práctico de sistemas digitales

Manejo de instrumentos de medición como multímetros, osciloscopios, fuentes de poder, etc.

Amplia experiencia en el uso de equipo de computo

Dominio de sistemas operativos como MS-DOS y Windows.

MANTENIMIENTO DE PC'S/XT'S Y PERIFERICOS

(Duración 40 hrs.)

Objetivo:

Lograr que los participantes después del curso, puedan alargar la vida útil de sus equipos, aplicando las acciones y cuidados del mantenimiento preventivo, evitando así problemas en potencia. Destruir los frecuentes mitos de que una computadora es "intocable por profanos" y que sólo los "GURUS" son los elegidos. Ayudar al usuario a decidir cuando puede resolver un problema por si mismo y cuando debe acudir a un especialista, y en general, dotarlo de seguridad en el manejo y cuidado de su PC y equipo.

Temas

- 1 ARQUITECTURA DE UNA COMPUTADORA
- 2 CARACTERISTICAS Y DIFERENCIAS ENTRE LOS SISTEMAS
- 3 TALLER DE CONFIGURACIONES
- 4 TALLER DE DIAGNOSTICOS
- 5 TALLER DE ENERGIA
- 6 MANTENIMIENTO PREVENTIVO LOGICO
- 7 MANTENIMIENTO CORRECTIVO BASICO
- 8 CONTRATOS DE MANTENIMIENTO.

TEMARIO

1. – ARQUITECTURA DE UNA COMPUTADORA
 - ~ Diagrama a bloques
 - ~ Sistema de reloj
 - ~ ISA
 - ~ ISA
 - ~ PCI
 - ~ Configuración practica en base a la arquitectura
2. – CARACTERISTICAS Y DIFERENCIAS ENTRE LOS SISTEMAS
 - ~ XT's, AT's, PS/2, Etc
 - ~ Evolucion desde el 8088 al Pentium III MMX
3. – TALLER DE CONFIGURACIONES
 - ~ Físicas y lógicas
 - ~ Interrupciones
 - ~ Mapas de memoria
 - ~ Mapas de puertos
 - ~ Canales de DMA, etc
4. – TALLER DE DIAGNOSTICOS
 - ~ Equipo de medición
 - ~ Herramientas físicas
 - ~ Herramientas lógicas.
 - ~ Programas de diagnostico
5. – TALLER DE ENERGIA
 - ~ Supresores de pico, reguladores y UPS's
 - ~ Fuentes
 - ~ Niveles de voltajes.
 - ~ Familias TTL y CMOS
6. – MANTENIMIENTO PREVENTIVO LOGICO
 - ~ En la computadora
 - ~ En el monitor
 - ~ En las impresoras
7. – MANTENIMIENTO CORRECTIVO BASICO
 - ~ MotherBoard
 - ~ Floppy's
 - ~ Monitores
 - ~ Impresoras
8. – CONTRATOS DE MANTENIMIENTO

CURSO DE MANTENIMIENTO DE PC's Y PERIFERICOS

EL "A,B,C," DEL MANTENIMIENTO DE PC'S Y PERIFÉRICOS

A).- RECEPCIÓN, REGISTRO Y PROGRAMA DEL SERVICIO

- 1.- Pruebas de funcionamiento. Inicialización
- 2.- No responsabilidades por información contenida

B).- MICROCOMPUTADORAS

- 1.- Limpieza externa y reportes de deterioros
- 2.- Limpieza interna y reportes de daños a la vista
- 3.- Chequeo de la calidad de la energía de la toma (voltaje y limpieza de la línea).
En domicilio
- 4.- Aplicación de programas del diagnóstico
- 5.- Chequeo de voltaje que da la fuente de poder, a los Drivers
- 6.- Limpieza de cabezas lectoras
- 7.- Limpieza de circuitería y tarjetas
- 8.- Chequeo de alineación de cabezas y velocidad de los motores
- 9.- Limpieza del teclado
- 10.- Reporte de pieza con deterioro y/o daños
- 11.- Aplicación de programas de diagnósticos, para prueba final y reporte de fallas
- 12.- Lacreo y sellado con etiquetas fechadas
- 13.- Tiempo empleado (programación del próximo servicio)

C) MONITORES

- 1.- Autoprueba
- 2.- Limpieza externa y reporte de deterioro
- 3.- Descarga del CRT (cinescopio)
- 4.- Limpieza interna y reporte de daños a la vista
- 5.- Chequeo y/o ajuste de señales de vídeo
- 6.- Prueba final de funcionamiento y reporte de fallas
- 7.- Lacreo y sellado con etiquetas fechadas

D) IMPRESORES Y GRAFICADORES

- 1.- Autopruebas
- 2.- Limpieza externa y reporte de deterioros
- 3.- Limpieza interna y reporte de daños a la vista
- 4.- Chequeo de voltaje
- 5.- Limpieza de partes robóticas (cabezas o brazos)
- 6.- Lubricación
- 7.- Prueba final de funcionamiento y reporte de fallas
- 8.- Lacreo y sellado con etiquetas fechadas
- 9.- Tiempo empleado

CURSO DE MANTENIMIENTO DE PC's Y PERIFERICOS

E) ENTREGA DEL EQUIPO

- 1.- Informe del servicio y actualización del control de mantenimiento
- 2.- Costo del servicio, refacciones y facturación
- 3.- Actualización del stock de refacciones y proveedores
- 4.- Salida de laboratorio con prueba para entrega
- 5.- Entrega física y recomendaciones

F) CONTRATOS

- 1.- Edad del equipo, calidad y lugar en el mercado
- 2 - Control de existencia en el mercado, de partes y refacciones
- 3 - Pólizas de servicio.- A domicilio o laboratorio
- 4.- Actualización de costos por inflación
- 5.- Actualización y firma de convenio

G) SEGURIDAD

- 1.- Seguros contra incendio, robos, etc.
- 2.- Guardar N° de series y facturas
- 3.- Etiquetar internamente, o marcar con rotulador de seguridad (sólo se ve con luz ultravioleta)
- 4 - Guardar fotografías del equipo
- 5.- Extintor con BCF (Bromoclorodifluorometano), gas Alone o Bióxido de carbono, que son los productos de extinción para emplearse sin peligro en equipos electrónicos
- 6.- Evitar en la instalación, ventanas y puertas por donde circula vientos, brizas y polvos

**SECUENCIA POST (POWER ON-SELF TEST) DEL
MICROPROCESADOR iX86**

- 1.- Recibe la señal de **POWER GOOD**.
- 2.- Genera señal de **RESET** al pin 21 del **8088** o del procesador.
- 3.- Los registros **DS, ES, SS e IP** son puestos a cero, **CS** apunta a **OFFFOh**.
- 4.- En **OFFFOh** se ejecuta la autoprueba.
- 5.- Se prohíben las interrupciones.
- 6.- Se colocan las banderas a 1, y se escribe y lee en los registros.
- 7.- Se ejecuta el **CHECKSUM** del **BIOS**.
- 8.- Se inicializan los **8253 y 8237**.
- 9.- Si el arranque es **WARM** se salta la prueba de memoria.
- 10.- Se prueban **16KB** de **RAM** con 5 secuencias diferentes.
- 11.- Se inicializa el **8259** (quién ordena las interrupciones).
- 12.- Se verifica el **8253**.
- 13.- Se inicializa el controlador de vídeo (p.e. **6845**) y se prueba la memoria de vídeo.
- 14.- Aparece el cursor.
- 15.- Comprueba si existe algo en las ranuras de expansión y comprueba el bus de datos y de direcciones.
- 16.- Comprueba el resto de la **RAM**.
- 17.- Comprueba el teclado (que no este pegada alguna tecla, que esté vacío el buffer y coloca el vector de interrupciones).
- 18.- Busca **ROM** opcional.
- 19.- Busca las unidades de disco (mediante los **DIPS SW**).
- 20.- Comprueba los puertos **LPTn y COMn**, almacena las direcciones validas.
- 21.- Habilita las **NMI**.
- 22.- Sonido por altavoz.
- 23.- Llama a **INT 19h** (unidad de disco).
- 24.- Stand by.




FACULTAD DE INGENIERIA U.N.A.M.
DIVISION DE EDUCACION CONTINUA

MANTENIMIENTO DE PC'S Y PERIFERICOS

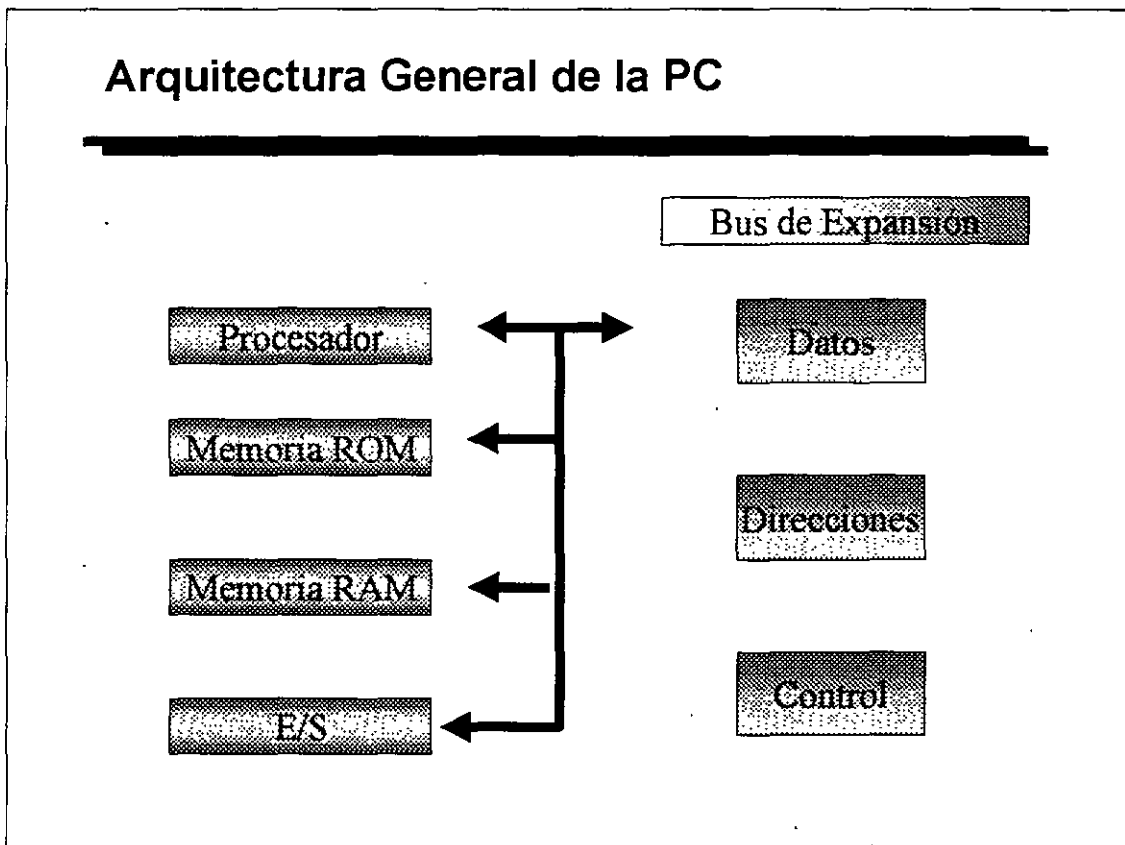
(PARTE II)

CARACTERÍSTICAS Y DIFERENCIAS ENTRE LOS
SISTEMAS, TECNOLOGÍAS Y TALLER DE
CONFIGURACIONES


JUNIO 2000




Arquitectura General de la PC



Notas:





Características de un Microprocesador


Un microprocesador es un circuito combinacional y secuencial que interactúa con otros circuitos para formar en conjunto un sistema digital de computo.

Funciones Principales:

- Provee las señales de tiempo y control para todos los elementos del sistema
- Busca instrucciones y datos desde la memoria.
- Transfiere datos desde y hacia dispositivos de Entrada /Salida.
- Decodifica Instrucciones
- Realiza operaciones lógicas y aritméticas solicitadas a través de instrucciones.
- Responde las señales de control de E/S, tales como RESET e INTERRUPT.




Notas:






Niveles de Interrupción XT

N°	Causa
NMI	Error de Paridad
0	Contador
1	Teclado
2	Reservado
3	Comunicación puerto serie(COM2), SDLC ó BSC (secundaria)
4	Comunicación/Puerto paralelo(COM1), FDLC ó BSC (primaria)
5	Disco Duro
6	Controlador de Disco Flexibles
7	Puerto Paralelo.



Notas:



Niveles de Interrupción AT

Nº	FUNCION
0	Timer del Sistema de Salida 0
1	Salida del teclado Buffer Lleno
2	Interrupción del Control 2 (Niveles 8-15).
3	Puerto serial 2
4	Puerto serial 1 .
5	Puerto paralelo 2
6	Controlador de disco
7	Puerto paralelo 1
8	Reloj de tiempo real
9	Redireccionado via Software al INT 0Ah
10	Reservado
11	Reservado
12	Reservado
13	80287
14	Disco duro
15	Reservado



Notas:




Memoria ROM (Read Only Memory)


Funciones Principales:

- ▣ Inicialización del Sistema.
- ▣ Diagnostico de Encendido y Revisión del Sistema.
- ▣ Determinación de la Configuración del Sistema.
- ▣ Manejo del Dispositivos de Entrada/Salida-BIOS.
- ▣ Cargado del Sistema Operativo.
- ▣ Patrones de Bits para los 1os. 128 Caracteres ASCII.



Notas:






Memoria RAM (Random Acces Memory)

Características Principales:

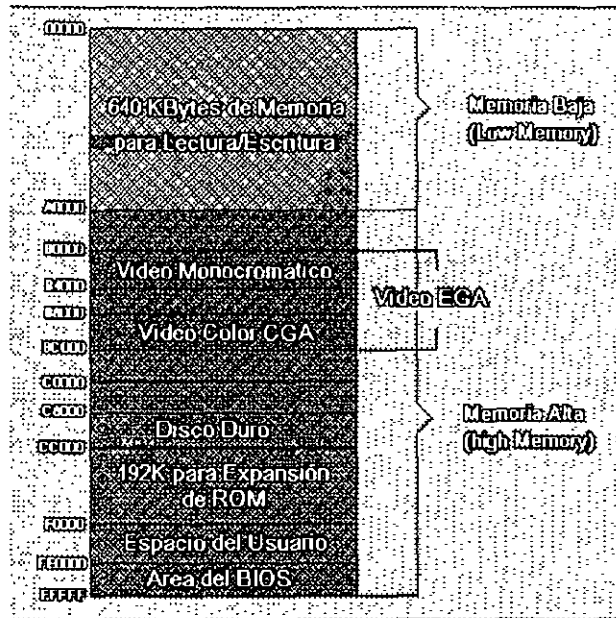
- ☞ Lectura/Escritura.
- ☞ Acceso Aleatorio
- ☞ Espacio Disponible al Usuario y sus Aplicaciones.
- ☞ Tamaño Limitado por el Número de Bits de Direcciones del Microprocesador.
- ☞ Se Dirección a través de un Mapa de Memoria Predefinido.
- ☞ Tiempo de Acceso de 150 a 60 nanosegundos.



Notas:



Mapa de Memoria



Notas:

Mapa de Puertos

Hex/Porte	Devices	Usage
00000	RAM Controller 1	System
00001	Memory Controller 1	System
00002	Timer	System
00003	UART (Hardware)	System
00004	Real-time clock with Mask	System
00005	DMA Controller	System
00006	Memory Controller 2	System
00007	DRAM Controller 2	System
00008	Cache (RAM Expansion) Bus	System
00009	Board on a Chip Controller	System
00010	RAM Expansion	System
00100	Power SW	IO
00101	Power SW	IO
00102	Power SW part 2	IO
00103	Power SW part 2	IO
00104	Power SW part 2	IO
00105	Power SW	IO
00106	Power SW	IO
00107	Power SW part 1	IO
00108	Power SW	IO
00109	Power SW	IO
00110	Power SW	IO
00111	Power SW	IO
00112	Power SW	IO
00113	Power SW	IO
00114	Power SW	IO
00115	Power SW	IO
00116	Power SW	IO
00117	Power SW	IO
00118	Power SW	IO
00119	Power SW	IO
00120	Power SW	IO
00121	Power SW	IO
00122	Power SW	IO
00123	Power SW	IO
00124	Power SW	IO
00125	Power SW	IO
00126	Power SW	IO
00127	Power SW	IO
00128	Power SW	IO
00129	Power SW	IO
00130	Power SW	IO
00131	Power SW	IO
00132	Power SW	IO
00133	Power SW	IO
00134	Power SW	IO
00135	Power SW	IO
00136	Power SW	IO
00137	Power SW	IO
00138	Power SW	IO
00139	Power SW	IO
00140	Power SW	IO
00141	Power SW	IO
00142	Power SW	IO
00143	Power SW	IO
00144	Power SW	IO
00145	Power SW	IO
00146	Power SW	IO
00147	Power SW	IO
00148	Power SW	IO
00149	Power SW	IO
00150	Power SW	IO
00151	Power SW	IO
00152	Power SW	IO
00153	Power SW	IO
00154	Power SW	IO
00155	Power SW	IO
00156	Power SW	IO
00157	Power SW	IO
00158	Power SW	IO
00159	Power SW	IO
00160	Power SW	IO
00161	Power SW	IO
00162	Power SW	IO
00163	Power SW	IO
00164	Power SW	IO
00165	Power SW	IO
00166	Power SW	IO
00167	Power SW	IO
00168	Power SW	IO
00169	Power SW	IO
00170	Power SW	IO
00171	Power SW	IO
00172	Power SW	IO
00173	Power SW	IO
00174	Power SW	IO
00175	Power SW	IO
00176	Power SW	IO
00177	Power SW	IO
00178	Power SW	IO
00179	Power SW	IO
00180	Power SW	IO
00181	Power SW	IO
00182	Power SW	IO
00183	Power SW	IO
00184	Power SW	IO
00185	Power SW	IO
00186	Power SW	IO
00187	Power SW	IO
00188	Power SW	IO
00189	Power SW	IO
00190	Power SW	IO
00191	Power SW	IO
00192	Power SW	IO
00193	Power SW	IO
00194	Power SW	IO
00195	Power SW	IO
00196	Power SW	IO
00197	Power SW	IO
00198	Power SW	IO
00199	Power SW	IO

Notas:




Bus de Expansión

Funciones Principales:

- Conectar los componentes funcionales al microprocesador
- Esta formado por:
 - Bus de datos
 - Bus de direcciones
 - Bus de control
- Además de señales de:
 - Tiempo
 - IRQs
 - DMA




Notas:




Dispositivos Inteligentes

DMA (Direct Memory Access)

- Ventajas:
 - Velocidad en el dispositivo
 - No distrae al microprocesador
 - Transferencia de información rápida.



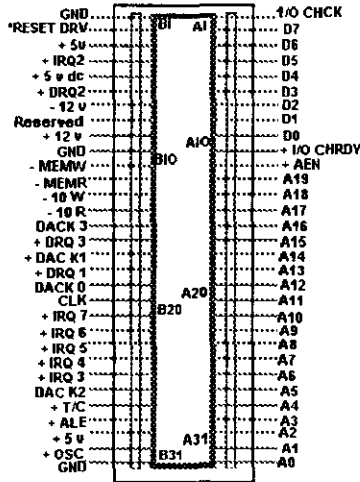
Notas:



Bus 8 Bits XT



REAR PANEL



I/O CHANNEL SLOT




Notas:


Extensión Bus AT (16 BITS)

-MEMCS16	D1	C1	SBHE
I/OCS 16	D2	C2	LA23
IRQ16	D3	C3	LA22
IRQ11	D4	C4	LA20
IRQ12	D5	C5	LA19
IRQ15	D6	C6	LA18
IRQ14	D7	C7	LA17
-DACK0	D8	C8	-MEMR
DRQ0	D9	C9	-MEMW
-DACK5	D10	C10	SD08
DRQ5	D11	C11	SD09
-DACK6	D12	C12	SD10
DRQ6	D13	C13	SD11
-DACK7	D14	C14	SD12
DRQ7	D15	C15	SD13
+5Vdc	D16	C16	SD14
-MASTER	D17	C17	SD14
GND	D18	C18	SD15

Notas:



Microprocesadores Familia "Intel"




8088



80286



386SX



386DX



486SX




486DX




486DX2



486DX4




Notas:



Arquitectura de la Microcomputadoras

Especificaciones IBM Personal Computer

- Fuente de poder de 63.5 Watts
- Microprocesador 8088 de 4.77 Mhz.
- 5 Slots de expansión (62 pins)
- Memoria de RAM Base de 16K-64K
- Bocina
- Unidad de Disco Flexible de 320 o 360K de 5_{1/4}
- Teclado de 83 Teclas.




Notas:




Arquitectura de la Microcomputadoras


Especificaciones IBM Personal Computer XT

- Fuente de poder de 130 Watts
- Microprocesador 8088 de 4.77 Mhz.
- 8 Slots de expansión (Socket62 pins)
- Memoria de RAM Base de 256K
- Disco Duro (en algunos modelos)
- Adaptador de comunicaciones asincronas (en algunos modelos)
- Teclado de 83 Teclas.




Notas:







Microprocesador Intel 8088

Frecuencia de Operación	4.77-12 Mhz	} Memoria = 1 MB	640 Kbytes Usuario 384 Kbytes Sistema
Tamaño del Bus de Expansión	16/8		
Tamaño del Bus de direcciones	20 -----		
Modos de Operación	Real		



Notas:






Microprocesador Intel 8086

Frecuencia de Operación	4.77-12 Mhz	
Tamaño del Bus de Expansión	16/16	
Tamaño del Bus de direcciones	20	Memoria = 1 MB { 640 Kbytes Usuario 384 Kbytes Sistema
Modos de Operación	Real	




Notas:



Arquitectura de las Microcomputadoras

Especificaciones IBM Personal Computer AT



- Fuente de poder 192 Watts switchable a 115 o 230
- Microprocesador 80286 a 6 Mhz
- 8 Slots de Expansión: 6 c/Socket de 62 pins y uno de 32
2 c/Socket de 62 pins únicamente
- Memoria RAM base de 256K
- Memoria RAM de tipo CMOS (semiconductor complementario de Oxidos Metálicos) para mantener la configuración del Setup del equipo.
- Batería para mantener activa la memoria CMOS cuando el equipo este apagado
- bocina
- Disco Duro
- Unidad de discos flexible de 5₁₄ de 1.2 MB
- Teclado con 83 Teclas

Notas:




Microprocesador Intel 80286

Frecuencia de Operación	6 - 25 Mhz	
Tamaño del Bus de Expansión	16/16	
Tamaño del Bus de direcciones	24	Memoria = 16MB { 640 Kbytes Usuario 384 Kbytes Sistema
Modos de Operación	Real/Protegido	




Notas:



Microprocesador Intel 80286

Modo REAL




Modo Protegido


- 16 MB Memoria RAM
- Multitareas
- Multiprocesamiento
- Memoria Virtual




Notas:




Modo Protegido de Operación 80286




Notas:







Microprocesador Intel 80386

Frecuencia de Operación	16 - 32 Mhz	
Tamaño del Bus de Expansión	32/32	
Tamaño del Bus de direcciones	32	Memoria = 4GB { 640 Kbytes Usuario 384 Kbytes Sistema
Modos de Operación	Real/Protegido/Virtual	



Notas:






Microprocesador Intel 80386/Sx

Frecuencia de Operación	16 - 33 Mhz	
Tamaño del Bus de Expansión	32/16	
Tamaño del Bus de direcciones	32 ----- Memoria = 4GB	} Limitante tecnología (256 Mbytes) Usuario ' 1 MByte Sistema'
Modos de Operación	Real/Protegido/Virtual 8086	



Notas:

Modos de Operación 80386



Notas:




Arquitectura 80386



- Micro Channel Adapter IBM MCA
- Flex Compaq
- Smartslot AST Research



Notas:



EISA




Miembros del Consorcio EISA

- AST Research
- Compaq
- Hewlett Packard
- NEC
- Zenith Data Systems
- Epson
- Olivetti
- Tandy
- Wyse_Technology




Notas:






Tecnologías de las Microcomputadoras



Notas:




Características Principales de las Diversas Arquitectura

Características	MCA	EISA	ISA
Amplitud máxima de datos	32 bits	32 bits	16 bits
Permite el uso de periodicos inteligentes y bus de 32 bits	SI	SI	NO
Promedio máximo de datos: DMA CPU	20 MB/seg 14 MB/seg	33 MB/seg 16 MB/seg	2 MB/seg 8 MB/seg
Sopórt e para memoria direccionable	16 MB	4 GB	16 MB
Compatibilidad	Ninguna	ISA	Ninguna

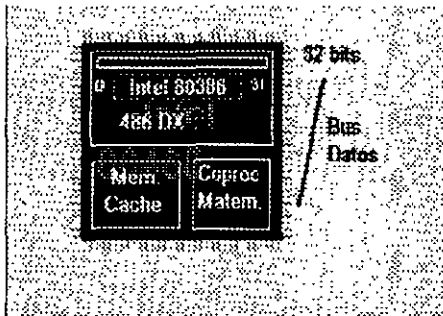


Notas:







Microprocesador Intel 80486




Características Similares es al 80386
Incluye Coprocesador Matematico
Incluye Memoria Caché



Notas:




Microprocesador Intel 80486/SX



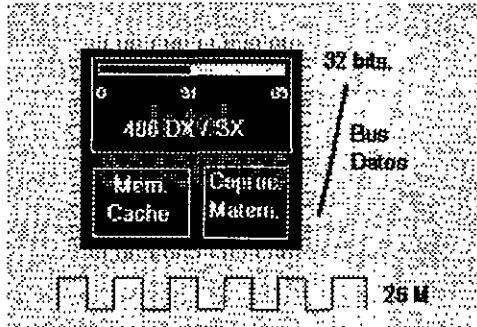
Características Similares 80386/sx
Incluye Memoria Caché




Notas:




Microprocesador 80486 DX2




80486 DX2 (Preinstalado)
CHIP Over Drive (Por Scket)



Notas:





La Memoria



Procesador



Memoria





Estrategias

- Simple DRAMS
- Simple SRAMS
- Interleaved RAM
- Page Mode
- Caching

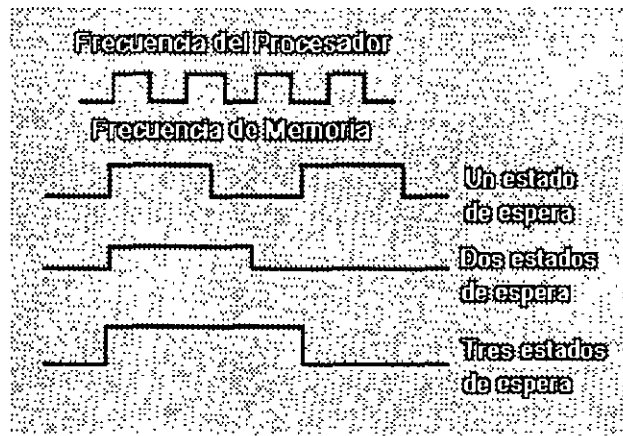


Notas:







Estados de Espera "Wait States"




Notas:



Tipos de Controladores

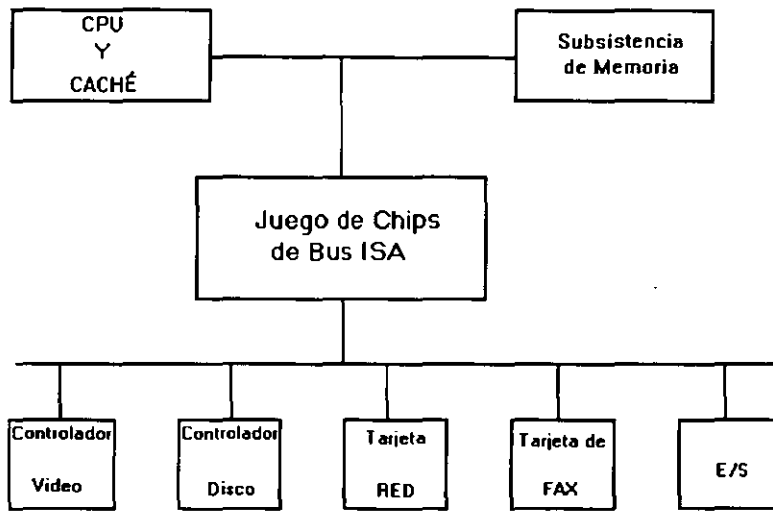


Notas:






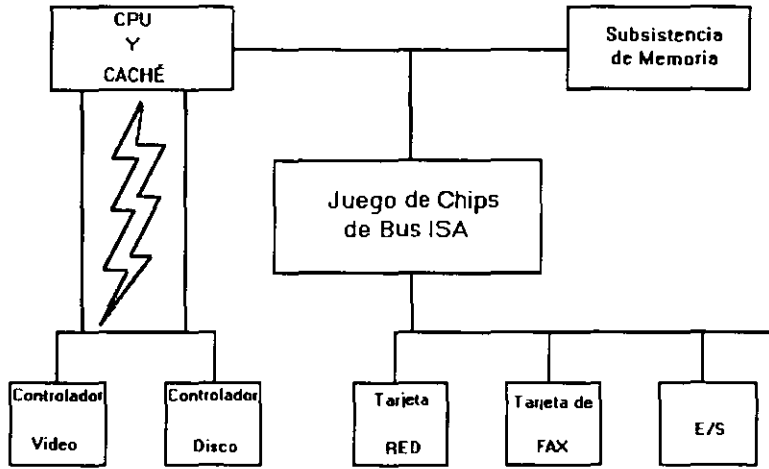
Bus Estándar AT




Notas:



Bus Local




Notas:







Bus Local

Desempeño del Bus



Notas:





Microprocesador Intel 80486


Existen versiones: SX, DX, DX2, DX4

Frecuencia de operación: 25 - 100 MHz

Tamaño del Bus de Datos: 32/16 SX
32/32 DX
64/32 DX2
128/32 DX4

Tamaño del Bus de Direcciones: 32 --Memoria 4 GB


Modos de Operación: Real / Protegido /Virtual 8086



Limitante tecnología
(2 Gbyte) Usuario
1 Mbyte Sistema



Notas:



Microprocesador Intel 80586 (Pentium)

Frecuencia de operación: 75 - 230 MHz


Frecuencia interna: 66 MHz

Tamaño del Bus de Datos: 32/64

Tamaño del Bus de Direcciones: 32 -- Memoria 0 4 GB

Modos de Operación: Real / Protegido / Virtual 8086 / MSS


Tecnologías asociadas: PCI, Flash Bios




Limitante tecnología
(2 Gbyte) Usuario
1 Mbyte Sistema




Notas:





Microprocesador Intel 80586 (Pentium)





3.3 millones de transistores
 Voltaje Máximo 3.8 V.
 Sistema de arquitectura por niveles
 con soporte a tecnología MMX

Existen versiones: Pentium,
 Pentium Pro,
 Pentium II
 Celeron

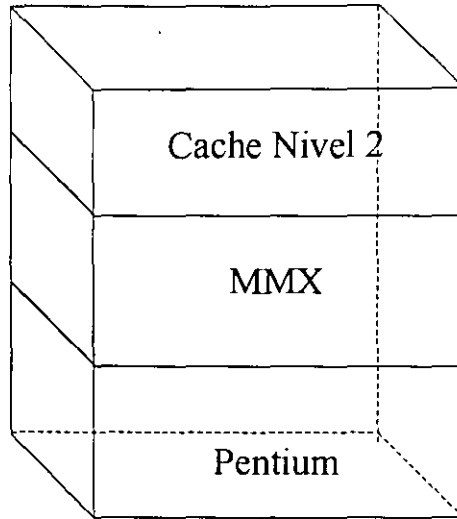


Notas:





Estructura del Pentium



Notas:



FACULTAD DE INGENIERIA U.N.A.M.
DIVISION DE EDUCACION CONTINUA

MANTENIMIENTO DE PC'S Y PERIFERICOS

(PARTE II)

TALLER DE ENERGÍA

JUNIO 2000

The Power Supply

One of the most failure-prone components in any computer system is the power supply. The power supply is a critical component in a PC, as it supplies electrical power to every component in the system. Because of its importance to proper and reliable system operation, you should understand both the function and limitations of a power supply, as well as its potential problems and their solutions.

Power Supply Function and Operation

The basic function of the power supply is to convert the type of electrical power available at the wall socket to that which is usable by the computer circuitry. The power supply in a conventional desktop system is designed to convert the 120-volt, 60Hz, AC current into something the computer can use—specifically, both +5- and +12-volt DC current. Usually, the digital electronic components and circuits in the system (motherboard, adapter cards, and disk drive logic boards) use the +5-volt power, and the motors (disk drive motors and any fans) use the +12-volt power. You must ensure a good, steady supply of both types of current so that the system can operate properly.

If you look at a specification sheet for a typical PC power supply, you see that the supply generates not only +5v and +12v, but also -5v and -12v. Because it would seem that the +5v and +12v signals power everything in the system (logic and motors), what are the negative voltages used for? The answer is not much! In fact, these additional negative voltages are not really used at all in modern systems.

Although -5v and -12v are supplied to the motherboard via the power supply connectors, the motherboard uses only the +5v. The -5v signal is simply routed to the ISA Bus on pin B5 and is not used in any way by the motherboard. It was originally used by the analog data separator circuits found in older floppy controllers, which is why it was supplied to the bus. Because modern controllers do not need the -5v, it is no longer used but is still required because it is part of the ISA Bus standard. Note that power supplies in systems with a Micro Channel Architecture (MCA) Bus do not have -5v. This power signal was never needed in these systems, as they always used a more modern floppy controller design.

The Power Supply

Both the +12v and -12v signals also are not used by the motherboard logic and instead are simply routed to pins B9 and B7 of the ISA Bus, respectively. These voltages can be used by any adapter card on the bus, but most notably they are used by serial port driver/receiver circuits. If the motherboard has serial ports built in, the +12v and -12v signals can be used for those ports. Notice that the load placed on these voltages by a serial port would be very small. For example, the PS/2 Dual Async adapter uses only 35mA of +12v and 35mA of -12v (0.035 amps each) to operate two ports.

Most newer serial port circuits no longer use 12v driver/receiver circuits, but instead now use circuits that run on only 5v or even 3.3v. If you have one of these modern design ports in your system, the -12v signal from your power supply is likely to be totally unused by anything in the system.

The main function of the +12v power is to run disk drive motors. Usually a large amount of current is available, especially in systems with a large number of drive bays, such as in a tower configuration. Besides disk drive motors, the +12v supply is used by any cooling fans in the system, which of course should always be running. A single cooling fan can draw between 100mA to 250mA (0.1 to 0.25 amps); however, most newer ones use the lower 100mA figure. Note that although most fans run on +12v, most portable systems use fans that run on +5v or even 3.3v instead.

In addition to supplying power to run the system, the power supply also ensures that the system does not run unless the power being supplied is sufficient to operate the system properly. In other words, the power supply actually prevents the computer from starting up or operating until all the correct power levels are present. Each power supply completes internal checks and tests before allowing the system to start. The power supply sends to the motherboard a special signal, called Power_Good. If this signal is not present, the computer does not run. The effect of this setup is that when the AC voltage dips and the power supply becomes over-stressed or overheated, the Power_Good signal goes down and forces a system reset or complete shutdown. If your system has ever seemed dead when the power switch is on and the fan and hard disks are running, you know the effects of losing the Power_Good signal.

IBM originally used this conservative design with the view that if the power goes low or the supply is overheated or over-stressed, causing output power to falter, the computer should not be allowed to operate. You even can use the Power_Good feature as a method of designing and implementing a reset switch for the PC. The Power_Good line is wired to the clock generator circuit (an 8284 or 82284 chip in the original PC/XT and AT systems), which controls the clock and reset lines to the microprocessor. When you ground the Power_Good line with a switch, the chip and related circuitry stop the processor by killing the clock signal and then reset the processor when the Power_Good signal appears after you release the switch. The result is a full hardware reset of the system.

Instructions for installing such a switch in a system not already equipped can be found later in this chapter.

Power Supply Form Factors

The shape and general physical layout of a component is called the *form factor*, and items that share form factor are generally interchangeable. When a system is designed, the designers can choose to use one of the popular standard form factors, or they can “roll their own.” Choosing the former means that a virtually inexhaustible supply of inexpensive replacements is available in a variety of quality and power output levels. Going the custom route means that the supply will be unique to the system and available only from the original manufacturer in only the model or models they produce. If you cannot tell already, I am a fan of the industry-standard form factors!

The form factor of the power supply that a particular system uses is based on the case design. Six popular case and power supply types can be called “industry standard.” The different types are

- PC/XT style
- AT/Desk style
- AT/Tower style
- Baby AT style
- Slim style
- ATX style

Each of these supplies are available in numerous different configurations and output levels.

When IBM introduced the XT, it used the same basic power supply shape as the original PC, except that the new XT supply had more than double the power output capability (see fig. 8.1). Because they were identical in both external appearance and the type of connectors used, you could easily install the better XT supply as an upgrade for a PC system. Because of the tremendous popularity of the original PC and XT design, a number of manufacturers began building systems that mimicked their shape and layout. These clones, as they have been called, could interchange virtually all components with the IBM systems, including the power supply. Numerous manufacturers have since begun producing these components, and nearly all follow the form factor of one or more IBM systems.

When IBM later introduced the AT desktop system, it created a larger power supply that had a form factor different from the original PC/XT. This system was rapidly cloned as well, and to this day still represents the basis for most IBM-compatible designs. The power supply used in these systems is called the AT/Desktop style power supply (see fig. 8.2). Hundreds of manufacturers now make motherboards, power supplies, cases, and so on that are physically interchangeable with the original IBM AT. If you are buying a compatible system, I recommend those that have form factors compatible with the IBM AT, because you will have numerous motherboards and power supplies from which to choose.



The Power Supply




Figure 8.1
PC/XT-form factor power supply

The compatible market has come up with a couple of other variations on the AT theme that are popular today. Besides the standard AT/Desktop type power supply, we also have the AT/Tower configuration, which is basically a full-sized AT-style desktop system running on its side. The power supply and motherboard form factors are basically the same in the Tower system as in the Desktop. The tower configuration is not new, in fact even IBM's original AT had a specially mounted logo that could be rotated when you ran the system on its side in the tower configuration. The type of power supply used in a tower system is identical to that used in a desktop system, except for the power switch location. Most AT/Desktop systems required that the power switch be located right on the power supply itself, while most AT/Tower systems use an external switch attached to the power supply through a short 4-wire cable. A full sized AT power supply with a remote switch is now called an AT/Tower form-factor supply (see fig. 8.3).

Another type of AT based form factor that has been developed is the so called Baby AT, which is simply a shortened version of the full sized AT system. The power supply in these systems is shortened on one dimension; however, it matches the AT design in all other respects. These Baby AT-style power supplies can be used in both Baby AT chassis and the larger AT-style chassis; however, the full size AT/Tower power supply does not fit in the Baby AT chassis (see fig. 8.4).




Figure 8.2. AT/Desktop form factor power supply.

The fifth type of form factor that has developed is the Slimline (see fig. 8.5). These systems use a different motherboard configuration that mounts the slots on a "riser" card that plugs into the motherboard. The expansion cards plug into this riser and are mounted sideways in the system. These types of systems are very low in height, hence the name "Slimline." A new power supply was specifically developed for these systems and allows interchangeability between different manufacturers' systems. Some problems with motherboard interchanges occur because of the riser cards, but the Slimline power supply has become a standard in its own right.

The slimline power supply is by far the most popular power supply design in use today. Despite how it might sound, even most full sized AT Desktop and Tower cases today are designed to accept the slimline form factor power supply.

The newest standard on the market today is the ATX form factor (see fig. 8.6). This describes a new motherboard shape, as well as a new case and power supply form factor. The ATX supply is based on the slimline or low-profile design, but has several differences worth noting.

System Components

The Power Supply





Figure 8.3

AT/Tower form factor power supply.

One difference is that the fan is now mounted along the inner side of the supply, blowing air across the motherboard and drawing it in from the outside at the rear. This flow is opposite most standard supplies, which blow air out the back of the supply and also have the fan positioned at the back. The reverse flow cooling used in the ATX supply forces air over the hottest components of the board, such as the CPU, SIMMs, and expansion slots. This eliminates the need for the notoriously unreliable CPU fans that have unfortunately become common today. Another benefit of the reverse flow cooling is that the system will remain cleaner and free from dust and dirt. The case is essentially pressurized, so air will push out of the cracks in the case, the opposite of what happens in non-ATX systems. For example, if you held a lit cigarette in front of your floppy drive on a normal

Power Supply Function and Operation

system, the smoke would be inhaled through the front of the drive and contaminate the heads! On an ATX system with reverse flow cooling, the smoke would be blown out away from the drive because the only air intake is the single fan vent on the power supply at the rear. Systems that operate in extremely harsh environments could add a filter to the fan intake vent, which would ensure even further that all air entering the system is clean and dust free.



System Components

Figure 8.4 Baby-AT form factor power supply

8.2.2 The Power Supply




Figure 8.5 Slimline/Low Profile form factor power supply.

The ATX system format was designed by Intel in 1995, but became popular in the new Pentium Pro based PCs in '96. The ATX form factor takes care of several problems with the Baby AT or slimline form factors, and where the power supply is concerned this covers two main problems. One is that the traditional PC power supply since the original one used in the IBM PC has two connectors that plug into the motherboard. The problem is that if you insert these connectors backwards or out of their normal sequence, you will fry the motherboard! Most responsible system manufacturers will have the motherboard and power supply connectors keyed so they cannot be installed backwards or out of sequence, but many of the cheaper system vendors do not feature this keying on the boards or supplies they use.

To solve the potential for disaster that awaits those who might plug their power supply connectors incorrectly, the ATX form factor includes a new power plug for the motherboard. This new connector features 20 pins, and is a single keyed connector. It is virtually impossible to plug it in backwards, and since there is only one connector instead of two almost identical ones, it will be impossible to plug them in out of sequence. The new connector also can optionally supply 3.3v, eliminating the need for voltage regulators on the motherboard to power the CPU and other 3.3v circuits. Although the 3.3v signals are labeled as optional in the ATX specification, they should be considered mandatory in any ATX form factor power supply you purchase. Many systems will require this in the future.




Figure 8.6 ATX form factor power supply.

Besides the new 3.3v signals, there is one other set of signals that will be found on the ATX supply not normally seen on standard supplies. They are the Power_On and 5v_Standby signals, which are also called Soft Power. Power_On is a motherboard signal that can be used with operating systems like Windows 95 or Windows NT, which support the ability to power the system down with software. This will also allow the optional use of the keyboard to power the system back on, exactly like the Apple Macintosh systems. The 5v_Standby signal is always active, giving the motherboard a limited source of power even when off.

The other problem solved by the ATX form factor power supply is that of system cooling. Most of the high end Pentium and Pentium Pro systems have active heat sinks on the processor, which means there is a small fan on the CPU designed to cool it. These small fans are notoriously unreliable, not to mention expensive when compared to standard passive heatsinks. In the ATX design, the CPU fan is eliminated, and the CPU is mounted in a socket right next to the ATX power supply, which has a reverse flow fan blowing onto the CPU. The ATX supply draws air from outside and pressurizes the system case instead of the other way around. Essentially the airflow is backwards from before, which results in far better cooling for the processor and other system components.

You will find it easy to locate supplies that fit these industry-standard form factors. Several vendors who manufacture PC power supplies in all these form factors are listed later in this chapter. For proprietary units, you will likely have to go back to the manufacturer.



The Power Supply

Power Supply Connectors

Table 8.1 shows the pinouts for most standard AT or PC/XT-compatible systems. Some systems may have more or fewer drive connectors. For example, IBM's AT system power supplies have only three disk drive power connectors, although most of the currently available AT/Tower type power supplies have four drive connectors. If you are adding drives and need additional disk drive power connectors, "Y" splitter cables are available from many electronics supply houses (including Radio Shack) that can adapt a single power connector to serve two drives. As a precaution, make sure that your total power supply output is capable of supplying the additional power.

Connector	AT Type	PC/XT Type
P8-1	Power_Good (+5v)	Power_Good (+5v)
P8-2	+5v	Key (No connect)
P8-3	+12v	+12v
P8-4	-12v	-12v
P8-5	Ground (0)	Ground (0)
P8-6	Ground (0)	Ground (0)
P9-1	Ground (0)	Ground (0)
P9-2	Ground (0)	Ground (0)
P9-3	-5v	-5v
P9-4	+5v	+5v
P9-5	+5v	+5v
P9-6	+5v	+5v
P10-1	+12v	+12v
P10-2	Ground (0)	Ground (0)
P10-3	Ground (0)	Ground (0)
P10-4	+5v	+5v
P11-1	+12v	+12v
P11-2	Ground (0)	Ground (0)
P11-3	Ground (0)	Ground (0)
P11-4	+5v	+5v
P12-1	+12v	—
P12-2	Ground (0)	—
P12-3	Ground (0)	—
P12-4	+5v	—
P13-1	+12v	—
P13-2	Ground (0)	—
P13-3	Ground (0)	—
P13-4	+5v	—

Notice that the Baby AT and Slimline power supplies also use the AT/Desktop or Tower pin configuration. The only other type of industry standard power supply connector is found on the new ATX form factor power supply. This is a 20-pin keyed connector with pins configured as shown in table 8.2.

Signal	Pin	Pin	Signal
3.3v*	11	1	3.3v*
-12v	12	2	3.3v*
GND	13	3	GND
Pwr_On	14	4	5v
GND	15	5	GND
GND	16	6	5v
GND	17	7	GND
-5v	18	8	Pwr_Good
5v	19	9	5v_Stby
5v	20	10	12v

* = Optional signal

Notice that the ATX supply features several signals not seen before, such as the 3.3v, Power_On, and 5v_Standby signals. Because of this, it will be difficult to adapt a standard slimline or low-profile form factor supply to work properly in an ATX system, although the shapes are virtually identical.

Although the PC/XT power supplies do not have any signal on pin P8-2, you can still use them on AT-type motherboards, or vice versa. The presence or absence of the +5v signal on that pin has little or no effect on system operation. If you are measuring voltages for testing purposes, anything within 10 percent is considered acceptable, although most manufacturers of high-quality power supplies specify a tighter five percent tolerance (see table 8.3). I prefer to go by the five percent tolerance, which is a tougher test to pass.

Desired Voltage	Loose Tolerance		Tight Tolerance	
	Min. (-10%)	Max. (+8%)	Min. (-5%)	Max. (+5%)
+/-5.0v	4.5v	5.4v	4.75	5.25
+/-12.0v	10.8v	12.9v	11.4	12.6

The Power_Good signal has tolerances different from the other signals, although it is nominally a +5v signal in most systems. The trigger point for Power_Good is about +2.5v, but most systems require the signal voltage to be within the tolerances listed in table 8.4.

System Components

The Power Supply

Signal	Minimum	Maximum
Power_Good (+5v)	3.0v	6.0v

A power supply should be replaced if the voltages are out of these ranges. A later section in this chapter details how to measure the power supply voltage and where to get replacement supplies.

Power Switch Connectors. The AT/Tower and Slimline power supplies use a remote power switch. This switch is mounted in the front of the system case and is connected to the power supply through a standard type of 4-wire cable. The ends of the cable are fitted with spade connector lugs, which plug onto the spade connectors on the power switch. The switch is usually a part of the case, so the power supply comes with the cable and no switch.

The cable from the power supply to the switch in the case contains four color coded wires. There may also be a 5th wire supplying a ground connection to the case as well. These wires carry 110v wall current, so be careful as you can be electrocuted if you touch them with the power supply plugged in.

The remote power switch leads carry 110v AC current at all times. You could be electrocuted if you touch the ends of these wires with the power supply plugged in! Always make sure the power supply is unplugged before connecting or disconnecting the remote power switch.

The four or five wires are color coded as follows; the Brown and Blue wires are the live and neutral feed wires from the 110v power cord to the power supply itself. These wires are always hot when the power supply is plugged in. The Black and White wires carry the AC feed from the switch back to the power supply itself. These leads should only be hot when the power supply is plugged in and the switch is turned on. Finally, there is often a green wire or a green wire with a yellow stripe, which is the ground lead. It should be connected somewhere to the PC case, and helps to ground the power supply to the case.

On the switch itself, the tabs for the leads are usually color coded, if not they can still be easily connected. If there is no color coding on the switch, then plug the Blue and Brown wires onto the tabs that are parallel to each other, and the Black and White wires to the tabs that are angled away from each other. See figure 8.7 as a guide.

As long as the Blue and Brown wires are on the one set of tabs, and the Black and White leads are on the other, the switch and supply will work properly. If you incorrectly mix the leads, you will likely blow the circuit breaker for the wall socket, as you can create a direct short circuit.




Figure 8.7 Power supply remote switch connections.

Disk Drive Power Connectors. The disk drive connectors are fairly universal with regard to pin configuration and even wire color. Table 8.5 shows the standard disk drive power connector pinout and wire colors.

Pin	Wire Color	Signal
1	Yellow	+12v
2	Black	Gnd
3	Black	Gnd
4	Red	+5v

This information applies whether the drive connector is the larger Molex version or the smaller mini-version used on most 3.5-inch floppy drives. In each case, the pinouts and wire colors are the same. To determine the location of pin 1, look at the connector carefully. It is usually embossed in the plastic connector body; however, it is often tiny and difficult to read. Fortunately these connectors are keyed and therefore are difficult to insert incorrectly. The following figure shows the keying with respect to pin numbers on the larger drive power connector.




Figure 8.8 A disk drive female power supply cable connector.

System Components

The Power Supply

Notice that some drive connectors may supply only two wires—usually the +5v and a single ground (Pins 3 and 4)—because the floppy drives in most newer systems run on only +5v and do not use the +12v at all.

Physical Connector Part Numbers. The physical connectors used in industry-standard PC power supplies were originally specified by IBM for the supplies used in the original PC/XT/AT systems. They used a specific type of connector between the power supply and the motherboard (the P8 and P9 connectors), as well as specific connectors for the disk drives. The motherboard connectors used in all the industry-standard power supplies have not changed since 1981 when the IBM PC appeared. With the advent of 3.5-inch floppy drives in 1986, however, a new smaller type of drive power connector appeared on the scene for these drives. Table 8.6 lists the standard connectors used for motherboard and disk drive power.

Connector Description	Female (on Power Cable)	Male (on Component)
Motherboard P8/P9	Burndy GTC6P-1	Burndy GTC 6RI
Disk Drive (large style)	AMP 1-480424-0	AMP 1-480426-0
Disk Drive (small style)	AMP 171822-4	AMP 171826-4

You can get these raw connectors through the electronics supply houses (Allied, Newark, Digi-Key, and so on) found in the vendor list. You also can get complete cable assemblies including drive adapters from the large to small connectors, disk drive "Y" splitter cables, and motherboard power extension cables from a number of the cable and miscellaneous supply houses such as Cables To Go, the Cable Connection, Ci Design, and Key Power.

The Power_Good Signal

The Power_Good signal is a +5v signal (+3.0 through +6.0 is generally considered acceptable) generated in the power supply when it has passed its internal self tests and the outputs have stabilized. This normally takes anywhere from 0.1 to 0.5 seconds after you turn on the power supply switch. This signal is sent to the motherboard, where it is received by the processor timer chip, which controls the reset line to the processor.

In the absence of Power_Good, the timer chip continuously resets the processor, which prevents the system from running under bad or unstable power conditions. When the timer chip sees Power_Good, it stops resetting the processor and the processor begins executing whatever code is at address FFFF:0000 (usually the ROM BIOS).

If the power supply cannot maintain proper outputs (such as when a brownout occurs), the Power_Good signal is withdrawn, and the processor is automatically reset. When proper output is restored, the Power_Good signal is regenerated and the system again begins operation (as if you just powered on). By withdrawing Power_Good, the system never "sees" the bad power because it is "stopped" quickly (reset) rather than allowed to operate on unstable or improper power levels, which can cause parity errors and other problems.



In most systems, the Power_Good connection is made via connector P8-1 (P8 Pin 1) from the power supply to the motherboard.

A well-designed power supply delays the arrival of the Power_Good signal until all voltages stabilize after you turn the system on. Badly designed power supplies, which are found in *many* low-cost compatibles, often do not delay the Power_Good signal properly and enable the processor to start too soon. The normal Power_Good delay is from 0.1 to 0.5 seconds. Improper Power_Good timing also causes CMOS memory corruption in some systems. If you find that a system does not boot up properly the first time you turn on the switch but subsequently boots up if you press the reset or Ctrl-Alt-Del warm boot command, you likely have a problem with Power_Good. This happens because the Power_Good signal is tied to the timer chip that generates the reset signal to the processor. What you must do in these cases is find a new high-quality power supply and see whether it solves the problem.

Many cheaper power supplies do not have proper Power_Good circuitry and often just tie any +5v line to that signal. Some motherboards are more sensitive to an improperly designed or improperly functioning Power_Good signal than others. Intermittent startup problems are often caused by improper Power_Good signal timing. A common example occurs when somebody replaces a motherboard in a system and then finds that the system intermittently fails to start properly when the power is turned on. This ends up being very difficult to diagnose, especially for the inexperienced technician, because the problem appears to be caused by the new motherboard. Although it seems that the new motherboard might be defective, it usually turns out to be that the original power supply is poorly designed and either cannot produce stable enough power to properly operate the new board, or more likely has an improperly wired or timed Power_Good signal. In these situations, replacing the supply with a high-quality unit is the proper solution.

Adding a Hardware Reset Switch. A switch that applies a full reset to your system keeps power moving to the system and rescues you from a system lockup. A reset switch saves you much time, as well as some of the wear and tear on your unit from using the power switch as a reset button. IBM and most vendors of compatibles have built reset circuitry into the motherboard and added reset switches to the front of the computer case. If your machine doesn't already have a reset switch, however, the following section teaches you how to add one. (The hardest part of adding a reset switch to your system is figuring out where to mount it.)


Adding a reset button is possible on any system, including all IBM systems, because it has a power supply that provides a Power Good signal. On most IBM-compatible computers, the Power Good signal is on the connector that plugs into the rearmost power-supply connectors. In PC and XT systems, this signal traces through the motherboard to the 8284a chip at Pin 11. When the line is shorted to ground and returned to normal, the 8284a (82284, in an AT) clock-timer chip generates a reset signal on Pin 10. The reset signal is sent to the 8088 at Pin 21, and the boot process begins. In other systems that have different processors and timer chips—for example, AT and PS/2 systems—the Power Good signal also initiates a reset if the signal is grounded and returned to normal, although the wiring details vary.




ATX form factor power supply.



Notas:








8.35"x5 9"x5 9"

AT/Tower form factor power supply.


Notas:






PC/XT-form factor power supply.

Notas:



Signal	Pin	Pin	Signal
3.3v*	11	1	3.3v*
-12v	12	2	3.3v*
GND	13	3	GND
Pwr_On	14	4	5v
GND	15	5	GND
GND	16	6	5v
GND	17	7	GND
-5v	18	8	Pwr_Good
5v	19	9	5v_Stby
5v	20	10	12v

* = Optional signal



Notas:



FACULTAD DE INGENIERIA U.N.A.M.
DIVISION DE EDUCACION CONTINUA

MANTENIMIENTO DE PC'S Y PERIFERICOS
(PARTE II)

MANTENIMIENTO CORRECTIVO BÁSICO

JUNIO 2000

8 UNIDADES DE DISCOS FLEXIBLES

DESCRIPCIÓN Y FUNCIONES DE LAS UNIDADES DE DISCOS FLEXIBLES 5¼" Y 3½"


Las unidades de discos son los dispositivos de las computadoras que permiten leer, grabar y borrar información en discos flexibles en forma digital.

Las primeras unidades de disco fabricada por Tandon y Control Data Corporation fueron diseñadas de altura completa (fig 4.1), posteriormente varias compañías, además de las anteriores fabricaron unidades de media altura, es decir de la mitad de altura de la primera, esto trajo beneficios para la industria y el usuario pues se comenzaba a compactar el tamaño de los periféricos y microcomputadoras en general, cabe mencionar que estos parámetros de altura se manejan en unidades de 5¼" de diámetro.


En cuanto a las unidades de 3½", desde su aparición en el mercado en 1987 aproximadamente, se han diseñado con las dimensiones que hasta la fecha rigen y es muy notorio el avance tecnológico, pues el tamaño nuevamente se redujo casi a la mitad.

A pesar de estas variaciones unidades de 5¼" y 3½" todas manejan el estándar de las señales de interfaz.






Es necesario hacer una pequeña descripción de los disco flexibles que emplean las unidades de 5 1/4", así como los utilizados por las de 3 1/2"



El disco está hecho de un material llamado Mylar (marca registrada por la compañía Dupont) Este plástico se reviste de una fina capa de material magnético (óxido de hierro) por ambos lados, a la vez que se pule Después se corta en círculos y se pule nuevamente; posteriormente se coloca dentro de la cubierta protectora, que en su parte interna tiene una capa suave y fina de tela nylon que evita que el disco se raye. A través de este proceso de fabricación se hace una serie de pruebas a los discos, de tal manera que al final de ella se le asigna la etiqueta de: un solo lado, dos lados, densidad simple, densidad doble, alta densidad, etc.





**Muesca de
protección
contra
escritura**

DISCO FLEXIBLE

El centro del disco está reforzado con un anillo de plástico que le permite soportar mayor presión por parte del eje de la unidad y puede hacerlo girar en forma apropiada. Muy cerca del hoyo central tiene uno más pequeño llamado índice.

Por otro lado, la cubierta protectora tiene un corte ovalado que permite el contacto necesario de las cabezas de lectura/escritura con la superficie del disco.

En un lado se encuentra un muesca la cual se conoce como muesca de protección de escritura. Estas características tienen los disco de 5¼".


En cuanto a los discos de 3½" cumplen con las misma características de fabricación, su variación consta, aparte del tamaño físico y lógico en cuanto a capacidad máxima de almacenamiento de información, el plástico que protege es más resistente, el corte ovalado donde se posiciona la cabeza de lectura/escritura está cubierta con un metal para evitar dañar el disco y al momento de entrar a la unidad se desplaza por un mecanismo propio de ésta. Uno de los factores más importantes de la confiabilidad de las unidades radica en el buen trato y manejo adecuado de los discos, para que la totalidad de los datos grabados no se altere. Un disco dañado o contaminado puede deteriorar o impedir la recuperación de los datos, pudiendo además, dañar las cabezas de lectura/escritura de la unidad. A continuación se enumera una lista con las recomendaciones necesaria para un mejor cuidado y manejo de los discos:

- 1) Manténgalos alejados de fuentes magnéticas, tales como bocinas, teléfonos, cualquier máquina de grabación, motores, etc.
- 2) No lo flexione o doble
- 3) No toque las superficies magnéticas del disco con los dedos
- 4) Inserte el disco con cuidado dentro de la unidad hasta donde lo permita el topo




5) Coloque el disco en su funda cuando no se use

La siguiente figura muestra un diagrama a bloques de los 7 grupos funcionales de que consta la unidad, mismos que se describen a continuación.



▣ FORMADOR DE PULSOS DE ÍNDICE


El circuito formador de pulsos de índice consiste de un led, un fototransistor y una red formadora de pulsos. El led y el fototransistor forman el sensor del pulso de índice, el primero como fuente luminosa y el segundo como receptor de esa luminosidad. Esta combinación led/fototransistor está colocada de tal manera que cuando el pulso de índice del disco para por ella, la luz de led índice en el fototransistor haciendo que conduzca y éste a su vez satura al transistor Q1 generando un nivel alto en TP7, esta señal junto con la señal DRIVE SELECT (nivel alto), produce un nivel bajo en la línea 8 de la interfaz.



▣ SENSOR DE PROTECCIÓN CONTRA ESCRITURA

La unidad cuenta con un sensor de protección contra escritura, constituido por un interruptor mecánico. El interruptor cierra cuando se inserta un disco con una etiqueta cubriendo la muesca de protección de escritura, deshabilitando la electrónica de escritura una señal de salida (nivel bajo) que indica este estado se proporciona en línea 28 de la interfaz.






SENSOR DE PISTA 00

El sensor de pista 00 está constituido por un interruptor que cierra cuando el carro alcanza una posición tal que las cabezas de lectura/escritura se encuentren en la pista 00

El nivel en la línea 26 de la interfaz 8N TRACK 00) es función de la posición del carro de las cabezas magnéticas, y de la fase del motor de pasos, como lo muestra el circuito de la siguiente figura cuando la cabeza se posiciona en la pista 00 y la fase del motor de pasos es 0 (líneas 4C-6 y 4C8 en niveles altos) se genera un nivel bajo en línea 26 de la interfaz



SISTEMA CONTROLADOR DE LA VELOCIDAD DEL MOTOR CD

El sistema que hace girar el disco consiste de un motor de CD (con tacómetro integrado), y un circuito que mantiene la velocidad a 300 r.p.m. en discos de baja densidad y 360 r.p.m. en discos de alta densidad como se muestra en la siguiente figura. El circuito cuenta con un limitador de corriente que deshabilita el motor cuando la corriente a través de él excede 1.3 amperes.

Cuando el nivel es bajo en la línea 16 de la interfaz, corresponde al habilitador del motor, éste comienza a girar hasta su velocidad nominal.

CONTROL DE POSICIÓN

El sistema de control de posición del carro de las cabezas de lectura/escritura, está constituido por un motor de pasos de cuatro fases y la electrónica que maneja dicho motor. Para avanzar una pista en el disco, el motor cambia una fase que provoca una rotación de un paso del motor y éste a su vez provoca un movimiento lineal del carro de las cabezas. Cuenta además con una compuerta que inhibe al circuito posicionador durante una operación de escritura. La dirección del carro de las cabezas depende del nivel de la señal DIR (línea 18 de la interfaz), ya que esto cambia la secuencia de fases del motor de pasos. Si el nivel es bajo el carro se moverá hacia la pista 39, si es alto se moverá hacia la pista 00



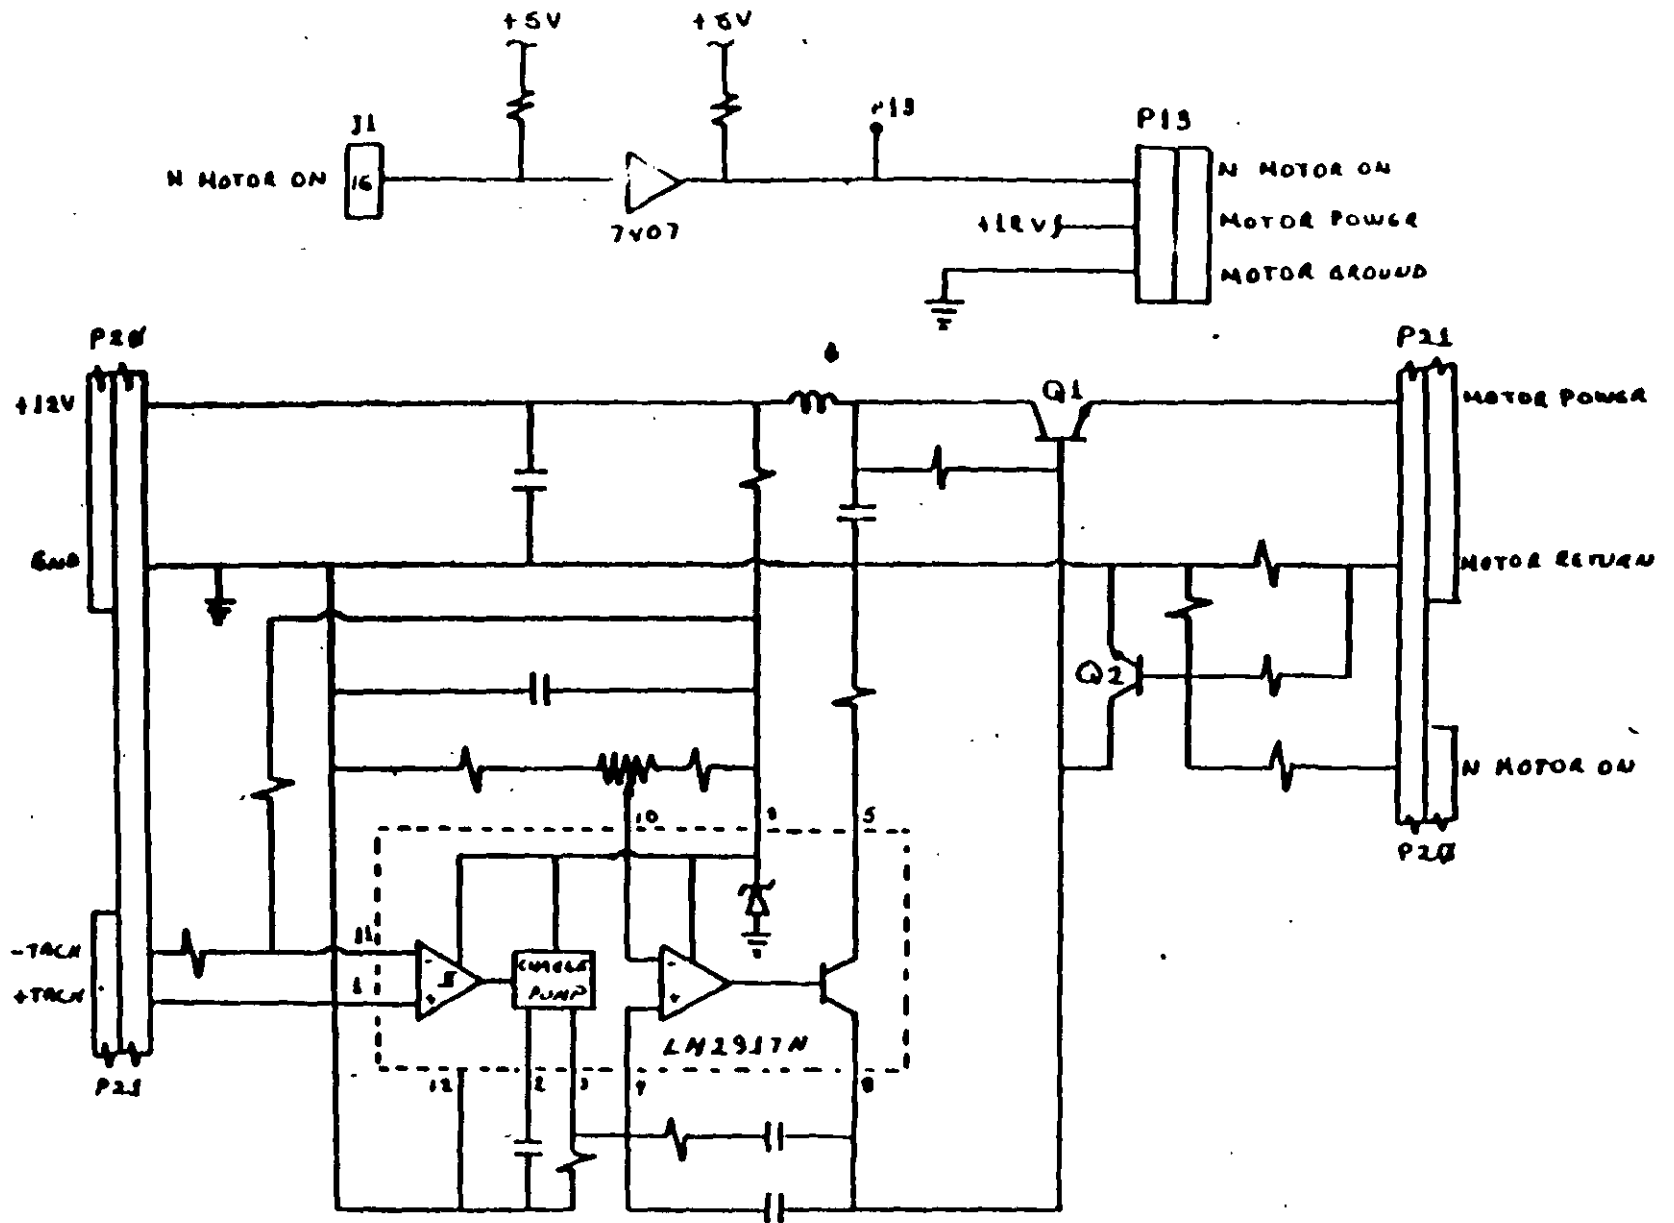


Figura 4.8

8-7
4-8

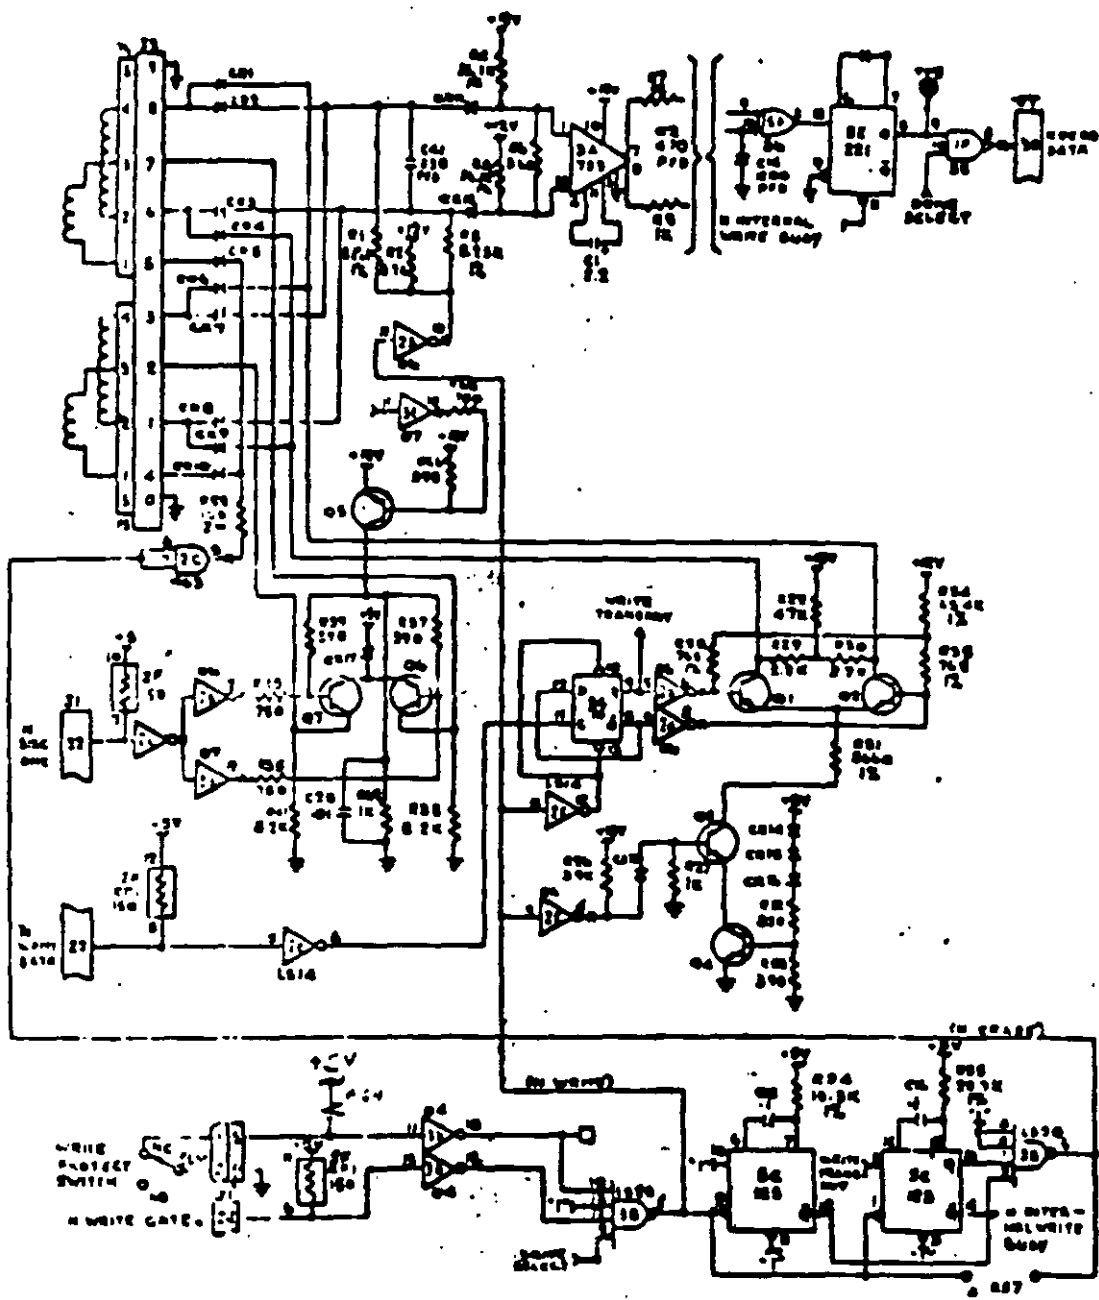


Figure 4.11



TABLA DE VERDAD DE FASES

Hacia la pista 39					Hacia la pista 00						
Pin N°	Fase					Pin N°	Fase				
	0	1	2	3	0		0	3	2	1	0
4C-5	0	0	1	1	0	4C-5	0	1	1	0	0
4C-6	1	1	0	0	1	4C-6	1	0	0	1	1
4C-8	1	0	0	1	1	4C-8	1	1	0	0	1
4C-9	0	1	1	0	0	4C-9	0	0	1	1	0

CONTROL DE ESCRITURA

La electrónica de escritura consiste de una fuente de corriente de escritura, un generador de forma de onda de escritura, una fuente de corriente de borrado, y la lógica de selección de cabeza como se muestra en la siguiente figura:

Las condiciones requeridas para un proceso de escritura, son las siguientes:

- 1) Estabilización de la velocidad de 300 o 360 r.p.m.
- 2) Estabilización del motor de pasos en la pista correcta
- 3) Habilitador de escritura
- 4) Escritura de datos
- 5) Selección de lado

CONTROL DE LECTURA

La electrónica de recuperación de datos incluye un amplificador de lectura, un diferenciador, un filtro, un detector de cruce por cero y un circuito de digitalización

- ↳ La señal pequeña a la salida de la cabeza de lectura/escritura se aumenta por un amplificado de lectura filtrada para quitar el ruido
- ↳ La señal de salida después del filtro se pasa la diferenciador, cuyo detector de cruce por cero genera una forma de onda correspondiente a picos de la señal de lectura
- ↳ Esta señal pasa después a un comparador y al circuito digitalizador que generan un pulso por cada pico de la señal de lectura. Por último esta señal pasa a la línea 30 de la interfaz.

PRUEBAS BÁSICAS PARA ALINEAR UNA UNIDAD DE DISCOS

HERRAMIENTAS NECESARIAS

- ↳ Un osciloscopio de dos canales



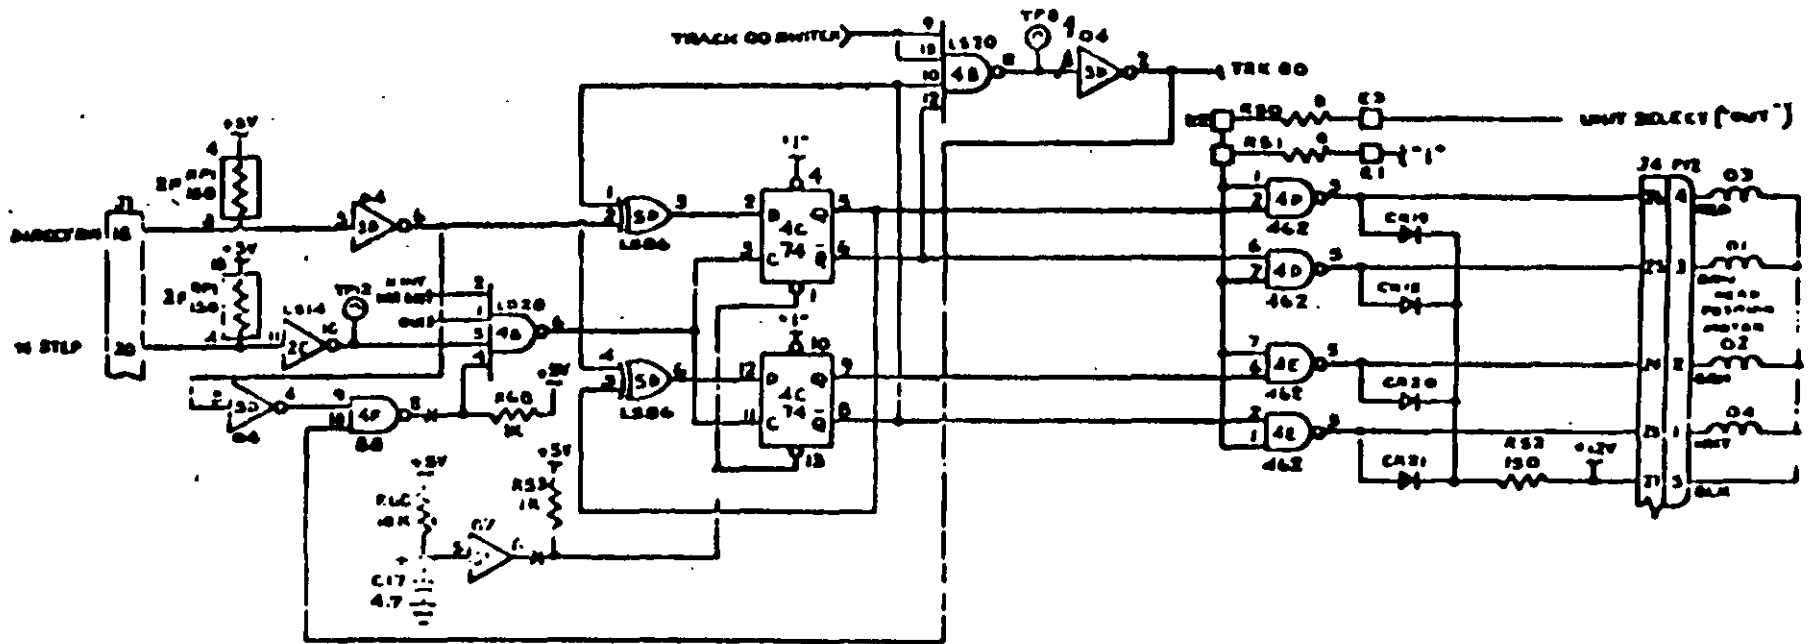


Figura 4.9



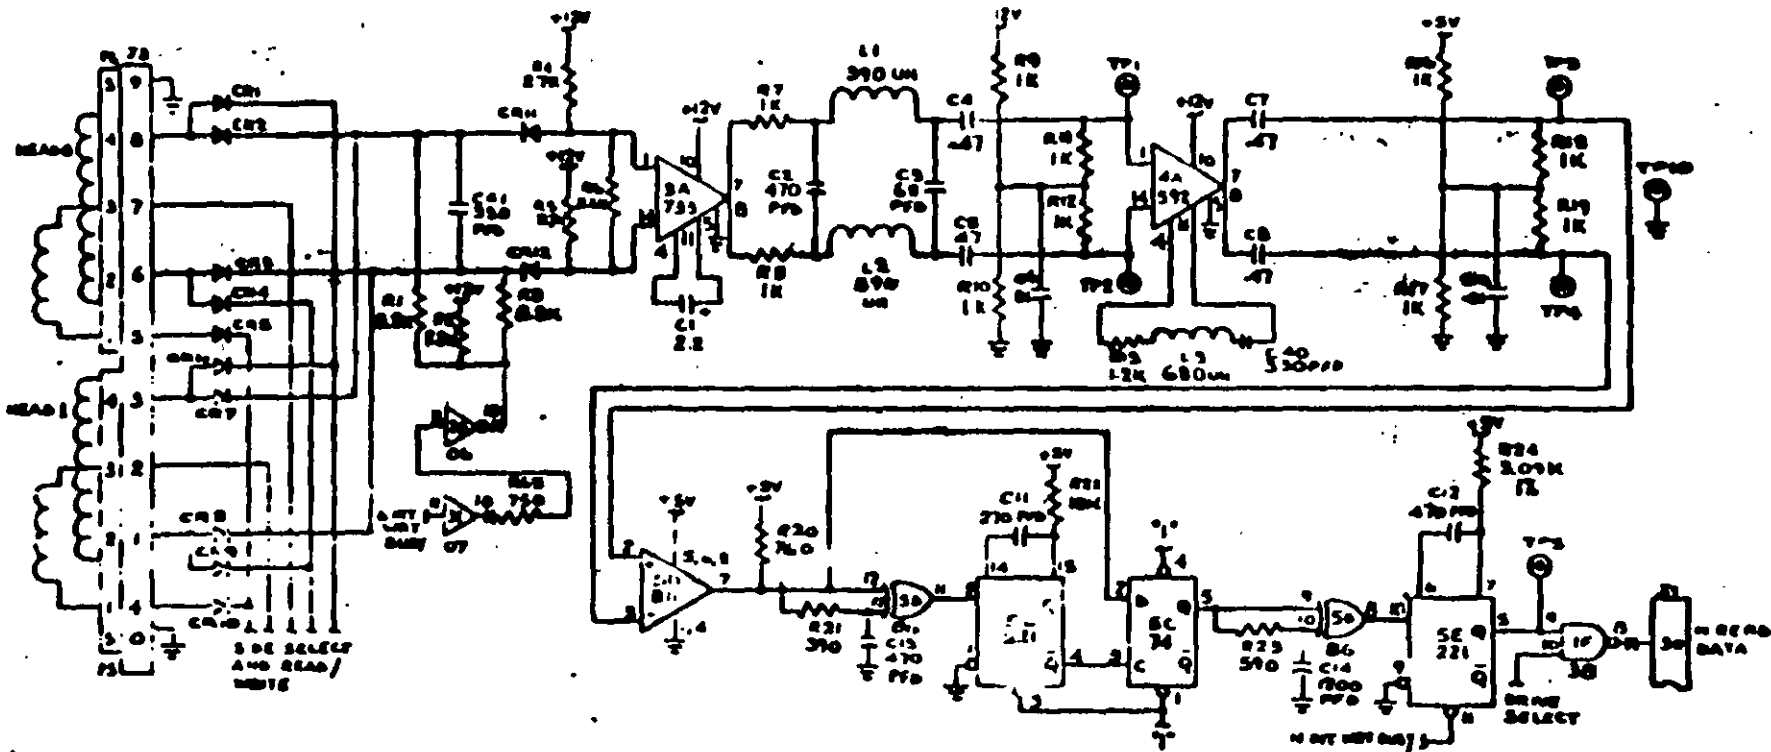


Figura 4.12

- ↳ Un programa capaz de seleccionar cualquier cabeza y colocarla en cualquier pista del disco
- ↳ Un juego de desarmadores
- ↳ Un disco de trabajo
- ↳ Software de alineamiento Dyssan o alguno equivalente

▣ UTILIZANDO COMO HERRAMIENTA EL OSCILOSCOPIO

- ↳ Comprobación y ajuste de la velocidad del motor

La velocidad del motor que hace que gire el disco debe mantenerse en 300 o 360 r.p.m. dependiendo el cada con una tolerancia de 5 r.p.m. Para comprobarlo se hace lo siguiente:

- 6) Verificar la alineación que entrega la fuente a la unidad, la cual debe ser:


$$+ 12 \text{ VDC} = 0.6 \text{ VDC}$$

$$+ 5 \text{ VDC} = 0.25 \text{ VDC}$$

- 7) Insertar un disco de trabajo


- 8) Con el programa de posicionamiento de cabezas habilitar el impulsor y observar bajo luz fluorescente las barras colocadas en la polea mayor que hace girar el disco situada en la parte inferior del impulsor.

VISTA INFERIOR DEL IMPULSOR TANDON



- 4) Ajustar la velocidad del motor con el potenciómetro R4 localizado en la tarjeta del servo hasta que el patrón permanezca inmóvil (50 Hz las interiores y 60 Hz las exteriores).

LOCALIZACIÓN DE R4



- 5) Hay que aclarar que este método de ajuste únicamente se utilizaba en las primeras unidades que salieron al mercado, actualmente los fabricante argumentan que esta velocidad de rotación del motor no tiene variaciones.

B) Comprobación y ajuste del alineamiento radial de las cabezas de lectura/escritura mediante el patrón de los "ojos de gato".

Esta prueba verifica que la cabeza de lectura/escritura esté en la distancia radial apropiada de la línea central del eje que hace girar el disco, asegurando así la colocación correcta de las cabezas en la pista deseada

COMPROBACIÓN

Coloque el osciloscopio como sigue.

- ↳ Canal A al TP1, canal B al TP 2 y tierra al TP10
- ↳ Lectura : A más B, B invertida
- ↳ Base de tiempo: 20 mseg Por división
- ↳ Disparo. Extremo con flanco positivo al TP7
- ↳ Insertar el disco de alineamiento
- ↳ Seleccionar la cabeza 0 (al inferior)
- ↳ Leer la pista 16 del disco para observar los "ojos de gato" en el osciloscopio Como se muestra en la siguiente figura.
- ↳ Verificar que uno de los "ojos de gato" no sea menor que el 75% de amplitud del otro
- ↳ Pase las cabezas a la pista 26 o una mayor, después regréselas a la pista 16 y compruebe nuevamente.
- ↳ Cambie la cabeza 1 (la superior) y repita los últimos 4 pasos



↪ Si cumple con lo anterior el alineamiento radial es aceptable, si no, continúe con lo siguiente.

▣ AJUSTE


↪ Afloje ½ vuelta los tornillos que sujetan el módulo de las cabezas

↪ Observe los "ojos de gato" de la cabeza que esté más desalineada

↪ Gire el tornillo CAM como se muestra en la siguiente figura Hasta que los "ojos de gato cumplan las condiciones del porcentaje de diferencia entre los tamaños

↪ Apriete los tornillos sujetadores y verifique nuevamente el patrón

PATRÓN DE LOS " OJOS DE GATO "



ESCALA DE TIEMPO A 20 MSEG

TORNILLOS QUE SUJETAN EL MODULO DE LAS CABEZAS

C) Comprobación y ajuste del pulso de índice

El ajuste del sensor de índice cambia el periodo de tiempo de pulso de índice al comienzo de los datos. Este ajuste debe realizarse después de alineamiento radial o cuando ocurre errores en la lectura de información.

▣ COMPROBACIÓN

↪ 0 Compruebe que la velocidad del motor sea correcta

↪ Conecte el osciloscopio como sigue:

↪ Canal A al TP1, canal B al TP2 y tierra al TP10

↪ Disparo Extremo con flanco positivo , al TP7

↪ Lectura: a mas B, B invertida

↪ Base de tiempo: 50 mseg. Por división



↪ Seleccione la cabeza 0

↪ Inserte el disco de alineamiento o al comienzo del primer pulso, como se muestra en la siguiente figura., que debe ser 200 mseg 100msg


↪ Compruebe que la cabeza 1 cumpla las mismas especificaciones que la cabeza 0 Si no las cumple, ajuste el sensor de índice

▣ AJUSTE

↪ Gire el tornillo que sujeta al sensor de índice en sentido contrario a las manecillas del reloj ¼ de vuelta como se muestra en la figura anterior

↪ Ajuste el sensor de índice con un desarmador hasta que el pulso comience 200 mseg 100mseg. Desde el punto de disparo

↪ Apriete el tornillo



Señal de índice al pulso de dato.
Base de tiempo: 50 pseg/div

D) Ajuste del interruptor de protección contra escritura

↪ Sin alimentación desconecte el conector P8 y compruebe que hay continuidad del interruptor

↪ Inserte un disco no protegido y habilite la unidad, compruebe que no hay continuidad entre los alambres de conector P8 y que hay un nivel alto en la línea 28 de la interfaz

↪ Inserte un disco protegido, compruebe que hay continuidad entre los alambres del conector P8 y que hay un nivel bajo en la línea 28 de la interfaz

↪ Para ajustar el interruptor afloje el tornillo que lo sujeta al soporte. Mueva el interruptor hacia arriba y hacia abajo hasta satisfacer las condiciones anteriores

▣ UTILIZANDO COMO HERRAMIENTA SOFTWARE DE DIAGNOSTICO



Por la amplitud de este tema el curso se ha diseñado en base a prácticas, utilizando software de diagnóstico para alineación de cabezas de lectura/escritura.

A lo largo del curso veremos este tipo de herramientas, quedando al juicio del capacitado que fabricantes de este tipo de software cumplen con los requerimientos del ingeniero de servicio.

TABLA DE ASIGNACIÓN DE PINES DEL CONECTOR DE INTERFAZ DEL IMPULSO

CONTROLLER-	TO-DISK	DRIVE
Ground	Signal	Mnemonic Description
1	2	Connector Clamp
3	4	Spura
5	6°	Select 3 (NDS3)
9	10	Select 0 (NDS0)
11	12	Select 1 (NDS1)
13	14	Select 2 (NDS 2)
15	16	Drive Motor Enable (N MOTOR ON)
17	18	Direction (DIR)
19	20	Step (N STEP)
21	22	Write Data (N WRITE DATA)
23	24	Write Gate (N WRITE ENABLE)
31	32	Side select (N SIDE SELECT)
33	34	Connector Clamp
7	8	Index (N INDEX/SECTOR)
25	26	Track 00 (N TRK 00)
27	28	Write Protect (N WRITE PROTECT)
29	30	Read Data (N READ DATA)

MANTENIMIENTO DE DISCO DURO. Un disco duro es un dispositivo que hace más versátil el uso de una computadora, reduce los tiempos de espera en carga y/o grabado de datos, aumenta la capacidad de almacenamiento y permite una mejor administración de los recursos del sistema. Así, es uno de los dispositivos que requieren mayor cuidado. La falla de un accesorio tan importante traer para el usuario grandes problemas, por lo que, en un servicio de mantenimiento, el tiempo de respuesta debe ser lo más pequeño posible.

El disco duro es en realidad un dispositivo combinado parte mecánica y parte electrónica; electrónicamente la función del dispositivo es transformar pulsos eléctricos en campos magnéticos permanentes, estos se llevan a cabo por medio de la utilización de "electroimanes", llamados en forma general cabezales de lectura/escritura, estas cabezas, alinean la polaridad de las partículas magnéticas contenidas en los discos o platos del dispositivo.

El mecanismo de un disco típico es relativamente simple, contiene pocas partes móviles. El elemento básico es un conjunto de uno o más platos que giran unidos a un eje conjunto común, llamado "spindle", este eje unido a un eje directamente a un servo-motor de la rotación controlada.

La parte adicional con movimientos dentro del mecanismo es el cabezal del sistema. Por lo general existe una cabeza por cada lado del disco, las que a su vez están conectadas a un brazo



en ensamble común a todas ellas, lo que permite el movimiento del conjunto de cabezas en unidad.

El ensamble de cabezas está unido al llamado actuador de cabezas (Motor o bobina), que permite el deslizamiento radial de la cabeza sobre los platos o discos.

Al contrario de la forma operación de los discos flexibles, el disco duro, permanece en rotación todo el tiempo que se mantiene encendida la máquina, dado que el tiempo requerido para vencer la inercia del sistema de disco es en promedio 20 segundos, el tiempo de acceso a información sería demasiado grande.

Debido a la constante rotación del sistema de discos, existen algunos inconvenientes como son, la constante pérdida de energía en forma de calor, el constante consumo de energía eléctrica y el desgaste de las partes por fricción

La capacidad de estos dispositivos, está dada en función del número de discos y cabezas dentro de el ensamble, por lo que con diferentes combinaciones de discos-cabezas, se puede tener capacidades iguales de almacenamiento, esta capacidad será o no completa (sin partes dañadas o inutilizables), dependiendo de la calidad del materia magnético de los discos y de la tecnología utilizada en su construcción.

Existen dos "tipos" reconocidos por una gran mayoría de personas en cuanto a discos duros se refiere, el llamado estandar-XT y el estándar - AT estas normalizaciones, están dadas de manera principal en los tiempos de acceso a información, par AT se requiere un máximo de 40 ms, en cambio para una máquina tipo XT, el tiempo de acceso puede ser de 100 ms o más. En general, un disco con menor tiempo de acceso, será un disco con mayor rendimiento.

Los platos de los discos rígidos son hechos usualmente de aluminio con recubrimiento de material que puede ser magnetizado, tradicionalmente, los discos duros eran realizado con una composición de óxido de hierro, de manera similar a las cintas y discos flexibles.

Recientemente un medio alternativo es una fina partícula magnética formada por una microscópica capa de material metálico puro o aleaciones de material unido a la superficie rígida. La ventaja de la película fina sobre la mezcla óxido-ferro, es principalmente que la película fina puede ser depositada en la superficie rígida de una manera más estrecha, proporcionando mayor capacidad de almacenamiento por menos espacio, además, el óxido de hierro llega a desprenderse de los platos, debido al recorrido radial de las cabezas, la película fina en cambio es realmente rígida lo que permite un ambiente más limpio. Para evitar los daños ocasionados por el golpeteo de las cabezas sobre los platos del disco, los modelos más recientes contiene seguros de cabezas, que alejan las cabezas a lugar seguro cuando el sistema se apega, reduciendo así la probabilidad de información.

No importando el sistema mecánico/magnético que tenga un disco duro, la nomenclatura utilizada para denominar la manera en que la información es grabada, resulta similar a la de los discos flexibles. Así, un track será un círculo formado por el movimiento circular de el disco en una posición fija de la cabeza, dado que el ensamble de la cabeza es un conjunto unido a un solo actuador, el giro de los discos sobre la posición fija del conjunto de cabezas forma un "cilindro" imaginario. Típicamente los discos duros para sistemas personales contienen un número entre 312 y 1,024 cilindros o tracks por disco



Cada track generalmente se divide en pequeños arcos de circunferencia, llamados sectores (17 usualmente). dos sectores se marcan en forma magnética con un programa de formateo inicial (inicializador). Esta operación de inicializador, permite el buen funcionamiento de un disco duro solamente si es realizado con los parámetros exactos de cilindros/cabezas, de otra manera el dispositivo podrá o no estar capacitado para manipular información, dependiendo si los parámetros son similares a los específicos por el fabricante del disco duro.

Dentro de la inicialización existe un factor importante llamado de entrelazado (interleave factor), que fuerza al sistema a leer un cierto sector, para luego saltar algunos antes de leer/escribir el siguiente, el número de sector saltando es el especificado durante esta inicialización es especificado por el factor. El valor de este factor influye en el rendimiento del sistema, pero generalmente en máquinas tipo XT es de 6 y en el tipo AT es de tres, dependiendo de la rapidez del sistema algunos discos manejan 1 ó 2 como factor de entrelazado.

Otro factor importante para maximizar el rendimiento de un disco duro, es la utilización de espacios intermedios de memoria (BUFFERS), que evitan el requerimiento constante de lectura al disco, esto se logra utilizando la opción BUFFERS, dentro del archivo de configuración de sistema CONFIG.SYS. Un valor recomendado para esta opción estará entre 10 y 20 dependiendo de la aplicación, es necesario además tomar en cuenta que esta opción toma espacio de memoria RAM, por lo que un número elevado en la opción requerirá más espacio de memoria de sistema operativo residente.

En lo referente al mantenimiento de disco duro, este se enfoca, básicamente, a la utilización de utilerías especiales para estos. El servicio realizado directamente en los discos es muy poco factible de realizarse por las condiciones propias de los mismos y por la casi nula disponibilidad de componentes en el mercado. Por lo anterior, el único mantenimiento directo a realizar es la limpieza de partes (conectores, sensores, etc) utilizando limpiadores y desengrasantes propios para equipo electrónicos.

Entre algunas de las utilerías especiales para disco duro podemos mencionar las siguientes:

- ↳ DEBUG: El DEBUG realiza la inicialización de la tarjeta controladora con disco duro esta inicialización es realizada por el fabricante.
- ↳ DIAGNOSTICO AVANZADO(PARA MAQUINA TIPO XT Y AT). Los diagnósticos avanzados tienen la capacidad de realizar pruebas, formatos de bajo nivel y, así mismo, la identificación de sectores defectuosos en discos duros.
- ↳ SPEEDSTOR: Esta utilería integra virtualmente cualquier tarjeta controladora de disco duro en una PC compatible XT o AT. Proporciona, además, un poderoso sistema de diagnósticos que permite identificar rápidamente los problemas en el disco y su controladora. Soporta diferentes capacidades (desde 10 MB hasta 320 MB) y tipos de discos tiene capacidad de realizar hasta ocho particiones de DOS y rutina de estacionamiento de cabezas.
- ↳ DISK MANAGER. Utilería que proporciona un sistema de diagnósticos que soporta diferentes tipos de discos. Tiene capacidad de realizar inicializaciones, particiones de DOS y formateo de bajo y alto nivel.

La primera regla en el mantenimiento a discos duros es el respaldo. Se deberá tener un respaldo de la información contenida en el disco antes de aplicar cualquier utilería de servicio



porque se podría incurrir en errores o fallas de potencia durante la ejecución de estos programas que puedan tener efectos fatales en los datos.

Por varias razones involucradas con el mismo medio magnético, bits de datos individuales- y algunas veces algunos bloques de ellos - pueden funcionar mal en el disco. Estos datos no tienen cambios extraños, sino que algunas áreas en el disco pierden su capacidad de almacenamiento, lo que cambia el espacio normal utilizable de disco.

El sistema operativo DOS no hace nada acerca de estos problemas. Pero todavía, hasta el programa FORMAT puede proporcionar una tabla de sectores dañados errónea y decirnos que algunos sectores son utilizables estando defectuoso, sirviendo como trampa para algunos datos.

Para prevenir tales sorpresas, pero no solucionar el problema periódicamente se deberá probar el disco para detectar sectores defectuosos. La utilización de algunas de las utilerías descritas anteriormente (diagnósticos avanzados, y DISK MANAGER) es recomendable.

Las pruebas del disco están habilitadas para localizar sectores defectuosos que el programa FORMAT, así que nos dan una respuesta más real acerca del disco. Algunos fabricantes de disco recomiendan que esto se deberá realizarse al menos cada mes, pero esto depende de cada disco y su utilización.

Las utilerías mencionada anteriormente, pueden realizar formateos de bajo nivel. Así como, en algunos casos, reparticiones el disco y, en seguida realizar formateos de alto nivel (FORMAT), esto es de gran ayuda para detectar tipo de problemas

☐ **MONITORES**

INTRODUCCIÓN.

Debido a la amplia variedad de monitores que existen en el mercado, describiremos en esta parte, en forma muy general, la teoría de operación y las características de los mismos. Dicha descripción deberá tomar en cuenta para asegurar su óptima operación.

- ☞ **INSTALACIÓN** El monitor está equipado con un conector de AC polarizado. Esta característica de seguridad le permite conectar la alimentación de la computadora al monitor.
- ☞ Deberá operar su monitor desde una fuente de alimentación indicada en las especificaciones del mismo. Si no está usted seguro de que tipo de alimentación tiene en el lugar ubicará su sistema consulte con el personal apropiado para solucionar estos casos.
- ☞ Las cargas y extensiones son peligrosas, así como los conectores defectuosos y conectores rotos. Pueden ocasionar un corto circuito. Verifique que estas partes estén en óptimas condiciones.
- ☞ No use su monitor cerca de agua.
- ☞ Los monitores están provistos de aberturas para ventilación en el gabinete para permitir la liberación del calor generado durante la operación. Si estas aberturas son bloqueadas, el calor puede causar fallas.

☐ **TEORÍA DE OPERACIÓN.**



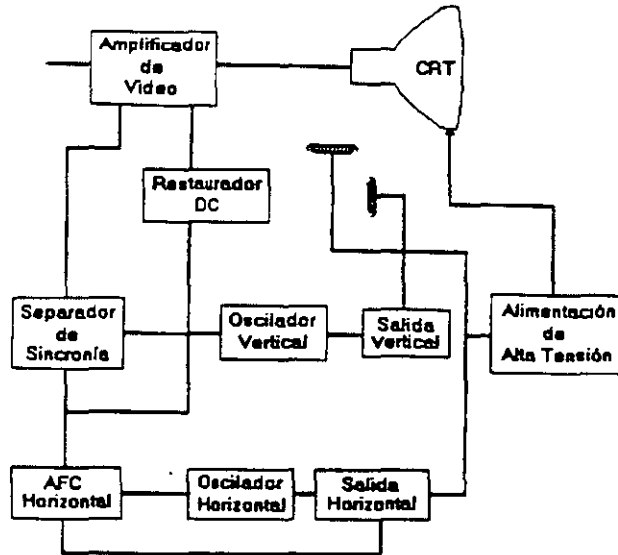
APLICACIÓN DE VÍDEO

- 1) La señal de vídeo compuesto es aplicada en el conector de entradas de vídeo, localizado en la parte trasera del monitor.

La señal es acoplada en amplificador de vídeo, el cual nos proporciona la salida de vídeo y la entrada al separador se sincroniza

- 2) CRT (tubo de rayo catódicos).

La señal de salida de vídeo es aplicada a un amplificador de alta ganancia, la cual nos proporciona la señal que va hacia el cátodo del CRT. Alta tensión es aplicada al CRT proporcionándole tensión adecuada y es del orden de miles de volts.



El flyback nos proporciona las tensiones para las rejillas de aceleración y foco del CRT.

- 3) SEPARADOR DE SINCRONÍA.

La señal de entrada al separador de sincronía obtenida del simplificador de vídeo entrada al separador de sincronía donde la señal positiva es recortada de la señal de vídeo.

La amplitud de los pulsos de sincronía es de 11 volts más o menos. Estas señales contienen ambos pulsos, el de horizontal y el de vertical.

- 4) Circuito de restauración de CD

Los pulsos de sincronía son retardados en este circuito para colocar en cada pulso retardado, en un punto fijo, el nivel obscuro de la señal de entrada de vídeo y todas las tonalidades de gris de la gama del negro al blanco aparecerán correctamente.

- 5) Oscilador vertical.



Los pulsos de sincronía vertical son acoplados al sistema de deflexión vertical de bajo ruido, el cual incorpora todas las funciones proporcionadas al yugo del monitor con la corriente requerida para deflexión vertical.

6) oscilador horizontal y alimentación de alta tensión .

El circuito del procesador horizontal realiza las siguientes funciones:

- ↳ Separador de ruido de sincronía horizontal
- ↳ Separador de ruido de sincronía vertical
- ↳ Oscilador horizontal.
- ↳ Comparador de frases enteras pulsos sincronizados y pulsos del oscilador.
- ↳ Comparador de frases entre pulsos del transformador flyback y pulsos del oscilador.

El transformador flyback genera los voltajes altos necesarios para el CTR y tiene un embobinado secundario para suministrar tensiones a las rejillas de aceleración y enfoque

☐ **PROBLEMAS GENERALES.**

Para realizar el servicio en un monitor, este deberá ser abierto y colocado en tal forma que la parte baja del circuito impreso sea accesible y una fuente de vídeo, tal como un sistema, esté conectado a la unidad en prueba

Los controles de contraste y brillantez deberá ser ajustado al máximo.

Cuando se enciende una unidad se deberá detectar un sonido de alta frecuencia, esto significa que el alto voltaje está siendo generado. Si no se detecta este sonido se deberá revisar el circuito del oscilador horizontal

Si hay alto voltaje y no hay rastreo en la pantalla, habrá que revisar el circuito del oscilador vertical

Si existe alto voltaje y rastreo pero el problema persiste, el paso a seguir es revisar el circuito de vídeo.

Si la unidad está completamente muerta, alguno de los fusibles probablemente esté fundido o existe algún problema en la fuente de alimentación.

Este tipo de seguimientos de fallas es muy general, se recomienda que para un monitor de un modelo determinado se profundice en el estudio de los circuitos particulares del mismo, así como las fallas generales presentadas en él .

☐ **SEÑALES DE ENTRADA**

Las señales de vídeo que nos proporcionan una tarjeta controladora son las siguientes.



N° DE TERMINAL DEL CONECTOR	SEÑAL
1	Tierra
2	Tierra
*3	Rojo
*4	Verde
*5	Azul
6	Intensidad
7	Vídeo Compuesto
8	Sincronía Horizontal
9	Sincronía Vertical

*Usado solamente para tarjeta Color Graphics.

El tipo de conector usado para realizar la conexión de la tarjeta controladora al monitor es de tipo D con 9 terminales o, para monitores monocromáticos, se utilizan también conectores de tipo A el cual traslada información de vídeo compuesto.

PROGRAMAS DE DIAGNOSTICOS

↳ Autoprueba de encendido (POWER-ON SELF TEST)

Cada vez que se enciende una PC, la microcomputadora realiza un diagnóstico rápido para asegurarse de que todas sus partes estén trabajando apropiadamente. Este diagnóstico toma aproximadamente 30 seg.

La respuesta normal después de una autoprueba es cuando el cursor esta parpadeando en la pantalla, seguido de un beep corto y, enseguida, la pantalla mostrará el sistema operativo o el software disponible en el impulsor o el sistema operativo en la unidad de disco duro (sise cuenta con el).

Se alguna parte tiene problema, la autoprueba dará una respuesta audible y se plegará en la pantalla un código de error que guiará al usuario acerca de la parte que está funcionando mal.

Los códigos de error se muestran en la siguiente tabla:

INDICACIÓN AUDIBLE	PROBLEMA
No hay beep (no hay despliegue)	Alimentación
Beep continuo	Alimentación
1 Beep largo y 1 beep corto	Alimentación
1 Beep largo y 2 Beeps cortos	Tarjeta Principal
1 Beep corto sin despliegue	Monitor
1 Beep corto sin despliegue	Monitor
1 Beep corto y BASICA en pantalla	Impulsor de discos flexibles



↳ Disquete de Diagnósticos

El disquete de diagnósticos está diseñado para dar una visión de los problemas que existen en una máquina y para hacer pruebas periódicas de las mismas.

Los diagnósticos son similares a la autoprueba de encendido y también utiliza códigos. La tabla siguiente muestra dichos códigos.

↳ Código de error de los Diagnósticos

Nota: si los dos últimos dígitos son cero, el sistema probado está operando correctamente.

La diferencia de los diagnósticos con autopruebas de encendido consiste en que los primeros son más poderosos y realizan las pruebas con más detalle.

Por otra parte, los diagnósticos no chocan dispositivos externos, tales como módem.

Cuando el disquete de diagnósticos han sido cargado, la pantalla despliega el siguiente menú:

```
The IBM personal computer DIAGNOSTICS
```

```
Versión 2.03 (c) Copyright IBM 1981,1983.
```

```
SELECT AND OPTION
```

```
0.- RUN DIAGNOSTICS ROUTINES.
```

```
1.- FORMAT DISKETTES.
```

```
2.- COPY DISKETTES.
```

```
3.- PREPARE SYSTEM FOR RELOCATION.
```

```
9.- EXIT TO SYSTEM DISKETTE.
```

```
ENTER THE ACTION DESIRED.
```

Donde el significado de cada opción es:

0.- RUN DIAGNOSTICS. comienza el procedimiento de prueba del sistema (ir al ,menú 2)

1.- FORMAT DISKETTE. Formatea un disquete para ser usado con los diagnósticos solamente.

2.- COPY DISKETTE. Copia el disquete de diagnóstico a otro disquete.

3.- PREPARE SYSTEM FOR RELOCATION. Coloca las cabezas del disco duro de estacionamiento para poder mover el sistema.

9.- EXIT TO SYSTEM DISKETTE. Carga el programa desde el disquete en el impulsor A.

después de teclear un 0, la pantalla desplegada deberá ser similar al menú 2 dependiendo de los dispositivos instalados.



THE INSTALLED DEVICES ARE:

- 1-S SYSTEM BOARD.
- 18-S EXPANSIÓN OPTION.
- 2-S XXXKB MEMORY
- 3-S KEYBOARD.
- 4-S MONCHOROME & PRINTER ADAPTER.
- 5-S COLOR/GRAPHICS MONITOR ADAPTER.
- 6-S X DISKETTE DRIVE(S) & ADAPTER
- 9-S PRINTER ADAPTER.
- 11-S ASYNC COMUNICACION ADAPTER.
- 12-S ALT ASYNC COMUNICACION ADPT.
- 13-S GAME CONTROL ADAPTER.
- 15-S SDLC COMUNICATIONS ADAPTER.
- 14-S MATRIX PRINTER.

IS THE LIST CORRECT (Y/N).

Este primer paso verifica que la computadora reconozca cuales dispositivos tiene conectados. Una de las maneras de hacer esto es checando los bancos de interruptores DIP (dual in package) dentro de la computadora. Si el despliegue en la pantalla es diferente a lo que se tiene conectado se deberán checar los cables y conexiones, así como la disposición de los interruptores DIP.

Si todo resulta compatible, teclee "Y" para indicar que las cosas están instaladas aparecen en la pantalla.

De esta manera pasamos a la siguiente pantalla (menú 3).

SYSTEM CHECKOUT.

- 0.- RUN TEST ONE TIME.
- 1.- RUN TEST MÚLTIPLE TIME.
- 2.- LOG UTILITIES.
- 9.- EXIT DIAGNOSTICS.

ENTER THE ACTION DESIRED.

Opciones 0 y 1, las rutinas de diagnósticos probarán los dispositivos del sistema y sus opciones uno por uno, empezando por la tarjeta de sistema (100) y a través de todas las opciones hasta concluir con el adaptador BSC (2100) si usted tiene dicho adaptador. Cuando una unidad es aprobada, la pantalla mostrará a usted que la unidad está funcionando correctamente,, (mostrando dos ceros en la parte última de código), o que



está funcionando incorrectamente (mostrando un código con algo diferente a dos ceros en su parte final).

Si los diagnósticos muestran una falla particular en alguna unidad, anote el código de error y continúe con los diagnósticos. puede ocurrir en otras unidades.

La opción 0 requiere una respuesta de usted en muchos pasos. La prueba de teclado, por ejemplo, le dice que presione cada tecla y vea en la pantalla el símbolo correcto. Si la unidad probada está funcionando bien, deberá teclear una "Y". Si la pantalla marca un error usted deberá marcar "N" y mostrará un código de error.

En la opción 1 no se requiere de una respuesta suya durante los diagnósticos. Solo deberá estar al pendiente de las pruebas. Con esta opción se pueden detectar problemas intermitentes. Si usted elige esta opción deberá decirle a la máquina cuantas veces quiere correr las pruebas y la opción para parar las mismas en cada error encontrado.

Opción 2. Los diagnósticos avanzados y estándar le dan a usted la oportunidad de grabar los mensajes de error que ocurran. Estos pueden realizarse a través de una impresora, al disquete de diagnósticos o a una unidad de cinta,, esto se hace posible con la opción 2 (LOG UTILITIES).

Registrar los errores a disquete requiere que usted tenga una copia de el disco de diagnóstico en el impulsor A y que no esté protegido contra escritura.

En resumen, estas son las características más importantes de los diagnósticos. Se recomienda hacer uso de los mismos para familiarizarse con funcionamiento.





FACULTAD DE INGENIERIA U.N.A.M.
DIVISION DE EDUCACION CONTINUA

MANTENIMIENTO DE PC'S Y PERIFERICOS

(PARTE II)

CONTRATOS DE MANTENIMIENTO

JUNIO 2000

CONTRATO No. _____

CONTRATO DE PRESTACIÓN DE SERVICIO DE MANTENIMIENTO que celebran por una parte: _____ quien en lo sucesivo se denominará el _____ y por otra parte _____ quien en lo sucesivo se denominará _____

.Ambas partes convienen en celebrar el contrato al rubro siguiente:

DECLARACIONES

1.0.- DECLARA EL CLIENTE por conducto de su representante.

1.1.- Que es una sociedad mercantil constituida de acuerdo a las leyes Mexicanas, que es propietario del equipo descrito en el anexo 1, que se agrega como parte de éste y que requiere de _____ el servicio de mantenimiento a sus equipos de cómputo.

1.2.- Que tiene su domicilio en _____

1.3.- Que está representado para este acto por _____

1.4.- Que su representante tiene las facultades necesarias para celebrar este contrato.

1.5.- Que el equipo objeto de este contrato descrito en el anexo 1, está instalado en _____

2.0.- DECLARA _____ por conducto de su representante:

2.1.- Que es una sociedad mercantil constituida de acuerdo a las Leyes Mexicanas.

2.2.- Que tiene su domicilio en _____
Tel: _____

2.3.- Que está representada en este acto por el _____

2.4.- Que tiene capacidad jurídica y técnica para contratar y prestar los servicios objeto de este contrato, y que cuenta con los recursos humanos y materiales necesarios, para el debido cumplimiento del mismo y los derivados de las relaciones personales con su personal

2.5.- Que puede acreditar el legal funcionamiento de _____ con la siguiente documentación.

- Escritura Pública No. _____

Registros:

- RFC. _____
- CED. _____
- CONACO. _____
- SPP (PROVEEDOR A GOB.): _____
- SPP (CONTRATISTA A GOB.): _____



3.0.- Declaran ambas partes protestando decir verdad, que su manifestación contenida en este convenio es voluntaria, libre de : dolo, mala fe, error voluntario o vicio alguno que pudiera nulificarlo todo o en partes.

CLAU S U L A S .

NATURALEZA DEL CONTRATO.

I).- A solicitud del _____, el contrato será de MANTENIMIENTO PREVENTIVO que _____ acepta prestar al equipo descrito en el anexo uno, de acuerdo a la siguiente cláusula, y con los cargos indicados en VI.

la).- _____ delega en _____ las funciones conducentes para efectos de saludable coordinación en los servicios materia de este contrato.

II).- Se define como MANTENIMIENTO PREVENTIVO, la limpieza externa e interna del equipo, la lubricación, los diagnósticos y los ajustes menores necesarios.

III).- EL MANTENIMIENTO PREVENTIVO será BIMESTRAL y programado con base a las necesidades específicas de c/u de las partes del equipo, según lo determine _____.

IV) .- La duración del presente contrato es de _____, debiendo avisar por escrito 30 días antes de su vencimiento, cualquiera de las partes la rescisión y/o actualización del mismo, de lo contrario, se considerará convenida por ambas partes, la renovación automática por otro período igual.

C A R G O S . -

VI).- Los cargos convenidos por el servicio, son por la cantidad de _____, que serán cubiertos por el _____ a _____ en moneda nacional, a la firma del presente contrato.

VII).- Si el equipo está fuera de la ciudad de México, los viáticos, traslados y gastos inherentes al caso, son con cargo al _____.

VII).- _____ prestará el servicio en el lugar mencionado, en 1.5 de las DECLARACIONES, en horas y días hábiles con personal propio.

IX .- _____ cuando así lo solicite podrá dar de alta en este contrato a equipos adicionales, previa negociación en costo y ajuste de tiempo, acompañando a su solicitud, un complemento del anexo 1.

X .- Un vez que ambas partes hayan convenido los cargos por nuevas altas en este contrato, y después de presentada y cobrada por _____ la factura correspondiente, se dará por aceptada la inclusión.



XI.- Quedan fuera de este contrato todos aquellos servicios que no estén contemplados en el MANTENIMIENTO PREVENTIVO como:

a) Reparación de las partes que no estén funcionando en el momento del primer servicio de mantenimiento, en el entendido de que este contrato se firma bajo el supuesto de que los equipos están trabajando normalmente.

b) Reparación de daños productos de accidentes, siniestros o negligencia en el uso del equipo, por efectos de humedad o cualquier otra causa distinta al uso normal.

c) En caso de interacción del equipo, conectado mecánica,, eléctrica o electrónicamente a otra máquina o mecanismos y no prevista esta situación en cláusula específica.

d) Los servicios de ingeniería de sistemas, operación y/o programación de cualquier tipo.

e) La reposición de cartuchos, de cintas de respaldo, cabezas de impresión partes de plástico,, CRT y discos.

f) Los trabajos externos del sistema, pintura y/o retocados de los equipos, aire acondicionado, instalaciones eléctricas etc.

XII.- Se define como MANTENIMIENTO CORRECTIVO las reparaciones y/o reemplazo de partes que resulten dañadas durante la operación normal del equipo y la mano de obra que se requiera.

XIII.- EL MANTENIMIENTO CORRECTIVO será por evento, previa solicitud telefónica o escrita del _____, a la que atenderá un técnico de _____ dentro de las siguientes seis horas hábiles.

a) Queda a juicio del Ingeniero de servicio si la reparación es posible resolverla en sitio, previo a un diagnóstico de fallas y daños detectados.

b) Si la reparación ha de resolverse en laboratorio, _____ recogerá el equipo después de contar con la orden respectiva de salida.

XIV.- La mano de obra calificada que se requiera será sin cargo alguno al _____

XV.- Las refacciones y PC-partes que resulten necesarias serán con cargo al _____.

XVI.- Si el servicio requiere de más de dos días hábiles, _____ dentro de su disponibilidad, facilitará al _____ un equipo lo más semejante al que esté en reparación.

XVII.- En fallas de disco duro sólo se dará servicio a circuitos y tarjetas externas, ya que por su tecnología, no permite su reparación aquí en México.

XVIII.- _____ no se hace responsable por la información contenida en los discos duros no obstante será su preocupación conservarla. Por seguridad, el _____ deberá respaldar su información periódicamente.



XIX.- Para efectos de prestar correctamente un servicio, el _____ se obliga en todo tiempo a dar las facilidades necesarias al personal de _____ quien respetará las normas y medidas de seguridad que indique el _____.

XX.- Si el _____ no da las facilidades necesarias para que se presente el servicio objeto de este contrato, cesará toda responsabilidad de _____.

XXI.- Si cualquier tipo de servicio de mantenimiento _____ considera que las falla o daño detectados fueron por negligencia de operación o mal uso del equipo el costo total de la reparación será con cargo al _____.

XXII.- Durante la vigencia de este contrato sólo el personal de _____ es el único autorizado para reparar, ajustar modificar o dar el servicio de mantenimiento al equipo materia de este convenio por lo que si cualquier persona ajena manipula con esa intención al equipo, el contrato quedará anulado.

XXIII.- _____ se obliga a presentar el servicio en los términos del presente contrato excepto en los casos de fuerza mayor que se lo impidan, tales como entrega tardía en los suministros de PC-partes, por problemas de importación, huelga en los organismos proveedores o por causas ajenas a _____.

XXIV.- Si el _____ modifica por su cuenta la configuración del equipo especificada en este contrato sin dar aviso, _____ se reserva el derecho de modificar las tarifas convenidas o apagares al artículo siguiente.

XXV.- Cualquier cambio en la Ingeniería del Hardware del equipo y/o de las instalaciones donde funcionará el mismo, será supervisado por _____ a efecto de ofrecer máxima seguridad. Si en estas acciones el _____ actúa sin acordado con _____, este contrato quedará sin efecto y todo servicio posterior, será con cargo adicional al _____.


XXIV.- En caso de violación de alguna cláusula de este contrato por cualquiera de las partes, la otra podrá exigir el cumplimiento o la rescisión del mismo.

XXVII.- Cualquier comunicación de una parte a la otra deberá hacerse por escrito, excepto los reportes telefónicos.

XVIII.- Para la interpretación o cumplimiento del presente contrato en caso de disputa, las partes se someten expresamente a los tribunales de la Ciudad de México.

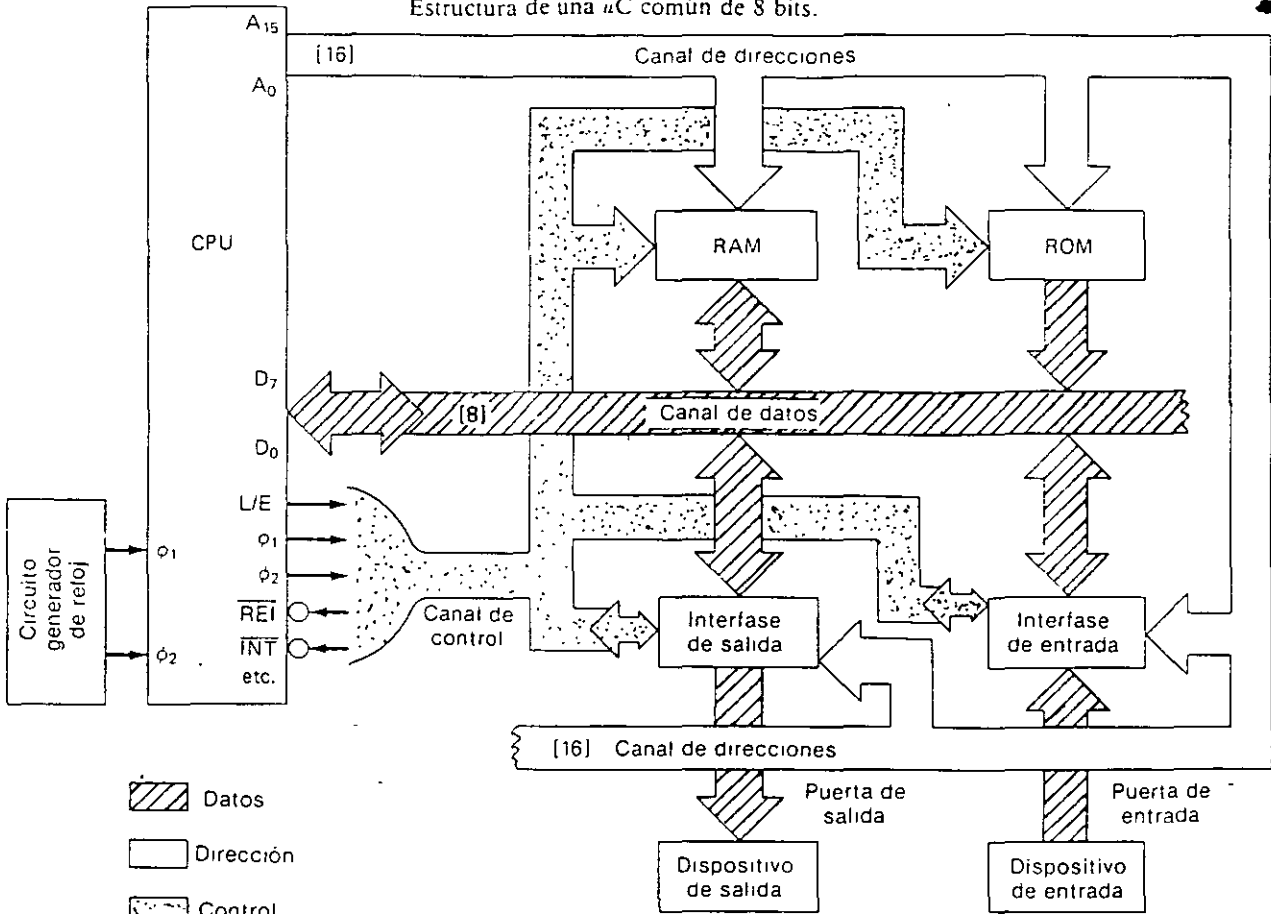
Para los efectos legales del caso, firman de acuerdo las partes que intervienen en este contrato, en la ciudad de México, a los _____.






Microprocesador



Estructura de una μ C común de 8 bits.



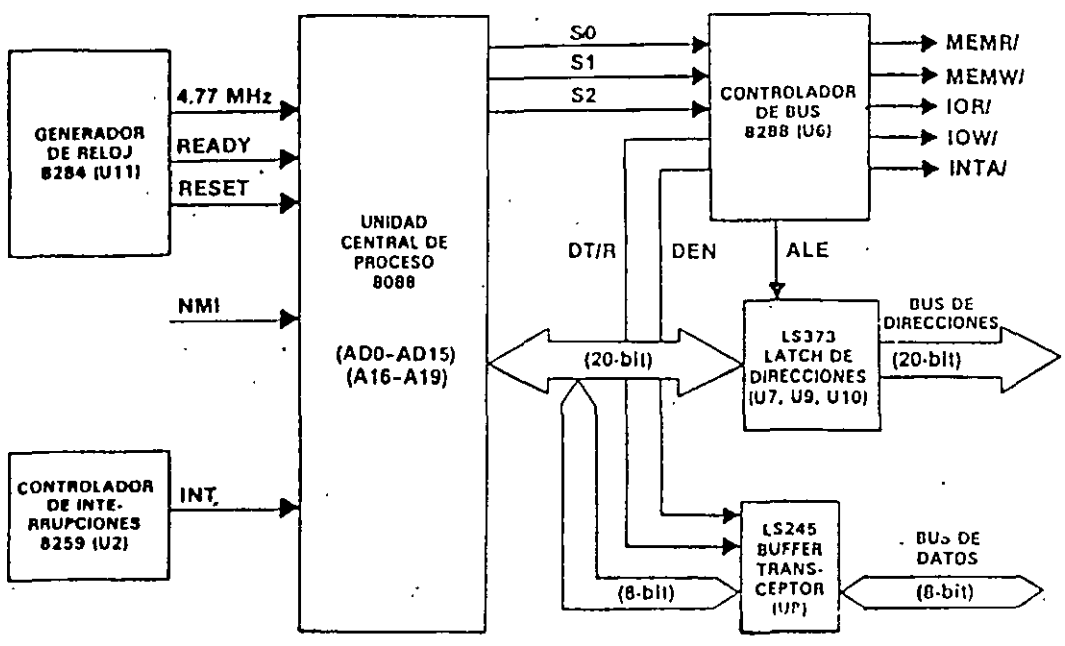

- Datos
- Dirección
- Control

Notas:






Notas:



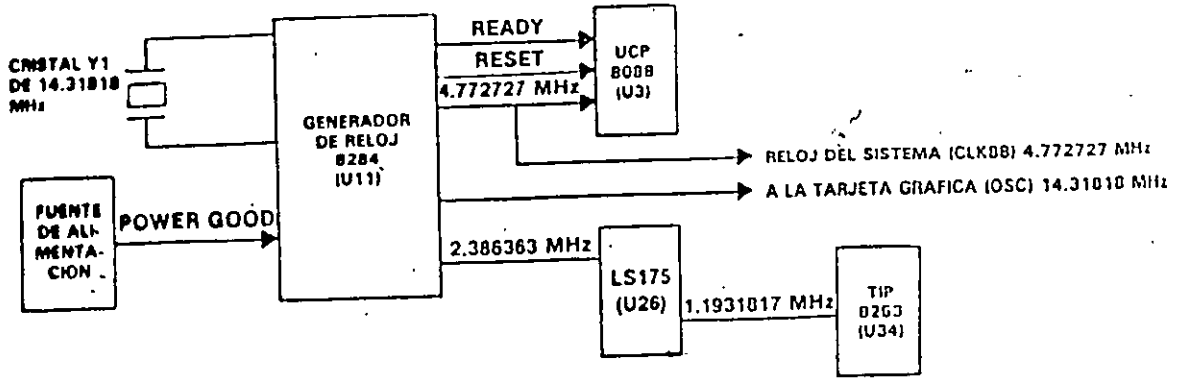

La UCP 8088 y el circuito asociado.



Notas:



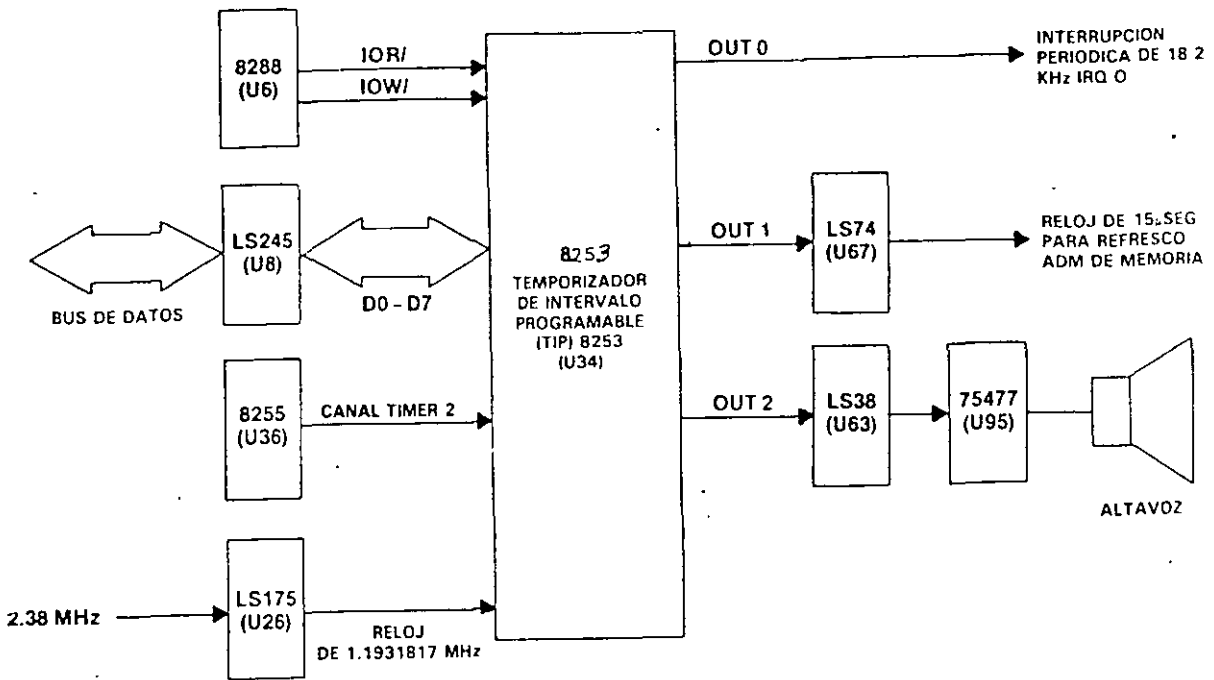
A



El circuito de reloj del IBM PC.




Notas:




El circuito del temporizador de intervalo programable (TIP) 8253.

Notas:




Socket Number	No. of Pins	Pin Layout	Voltage	Supported Processors
Socket 1	169	17x17 PGA	5v	SX/SX2, DX/DX2*, DX4 OverDrive
Socket 2	238	19x19 PGA	5v	SX/SX2, DX/DX2*, DX4 OverDrive, 486 Pentium OverDrive
Socket 3	237	19x19 PGA	5v/3.3v	SX/SX2, DX/DX2, DX4, 486 Pentium OverDrive
Socket 4	273	21x21 PGA	5v	Pentium 60/66, Pentium 60/66 OverDrive
Socket 5	320	37x37 SPGA	3.3v	Pentium 75-133, Pentium 75+ OverDrive
Socket 6**	235	19x19 PGA	3.3v	DX4, 486 Pentium OverDrive
Socket 7	321	37x37 SPGA	VRM	Pentium 75-200, Pentium 75+ OverDrive
Socket 8	387	dual pattern SPGA	VRM	Pentium Pro


*DX4 also can be supported with the addition of an aftermarket 3.3v-regulator adapter.
 **Socket 6 was a paper standard only and was never actually implemented in any systems.
 PGA = Pin Grid Array
 SPGA = Staggered Pin Grid Array
 VRM = Voltage Regulator Module




Notas:




SOCKETS




Notas:




SOCKETS




Notas:




SOCKETS




Notas:




SOCKETS




Notas:




SOCKETS




Notas:



SOCKETS



Notas:



BUSES

Signal	Pin	Pin	Signal
Ground	B1	A1	-I/O CH CHK
RESET DRV	B2	A2	Data Bit 7
+5 Vdc	B3	A3	Data Bit 6
IRQ 2	B4	A4	Data Bit 5
-5 Vdc	B5	A5	Data Bit 4
DRQ 2	B6	A6	Data Bit 3
-12 Vdc	B7	A7	Data Bit 2
-CARD SLCTD	B8	A8	Data Bit 1
+12 Vdc	B9	A9	Data Bit 0
Ground	B10	A10	-I/O CH RDY
-SMEMW	B11	A11	AEN
-SMEMR	B12	A12	Address 19
-IOW	B13	A13	Address 18
-IOR	B14	A14	Address 17
-DACK 3	B15	A15	Address 16
DRQ 3	B16	A16	Address 15
-DACK 1	B17	A17	Address 14
DRQ 1	B18	A18	Address 13
-Refresh	B19	A19	Address 12
CLK(4.77MHz)	B20	A20	Address 11
IRQ 7	B21	A21	Address 10
IRQ 6	B22	A22	Address 9
IRQ 5	B23	A23	Address 8
IRQ 4	B24	A24	Address 7
IRQ 3	B25	A25	Address 6
-DACK 2	B26	A26	Address 5
T/C	B27	A27	Address 4
BALE	B28	A28	Address 3
+5 Vdc	B29	A29	Address 2
OSC(14.3MHz)	B30	A30	Address 1
Ground	B31	A31	Address 0

El bus ISA de 8 bits

Notas:



BUSES

Signal	Pin	Pin	Signal
Ground	B1	A1	-I/O CH CHK
RESET DRV	B2	A2	Data Bit 7
+5 Vdc	B3	A3	Data Bit 6
IRQ 9	B4	A4	Data Bit 5
-5 Vdc	B5	A5	Data Bit 4
DRQ 2	B6	A6	Data Bit 3
-12 Vdc	B7	A7	Data Bit 2
-0 WAIT	B8	A8	Data Bit 1
+12 Vdc	B9	A9	Data Bit 0
Ground	B10	A10	-I/O CH RDY
-SMEMW	B11	A11	AEN
-SMEMR	B12	A12	Address 19
-IOW	B13	A13	Address 18
-IOR	B14	A14	Address 17
-DACK 3	B15	A15	Address 16
DRQ 3	B16	A16	Address 15
-DACK 1	B17	A17	Address 14
DRQ 1	B18	A18	Address 13
-Refresh	B19	A19	Address 12
CLK(8.33MHz)	B20	A20	Address 11
IRQ 7	B21	A21	Address 10
IRQ 6	B22	A22	Address 9
IRQ 5	B23	A23	Address 8
IRQ 4	B24	A24	Address 7
IRQ 3	B25	A25	Address 6
-DACK 2	B26	A26	Address 5
T/C	B27	A27	Address 4
BALE	B28	A28	Address 3
+5 Vdc	B29	A29	Address 2
OSC(14.3MHz)	B30	A30	Address 1
Ground	B31	A31	Address 0

-MEM CS16	D1	C1	-SBHE
-I/O CS16	D2	C2	Latch Address 23
IRQ 10	D3	C3	Latch Address 22
IRQ 11	D4	C4	Latch Address 21
IRQ 12	D5	C5	Latch Address 20
IRQ 15	D6	C6	Latch Address 19
IRQ 14	D7	C7	Latch Address 18
-DACK 0	D8	C8	Latch Address 17
DRQ 0	D9	C9	-MEMR
-DACK 5	D10	C10	-MEMW
DRQ5	D11	C11	Data Bit 8
-DACK 6	D12	C12	Data Bit 9
DRQ 6	D13	C13	Data Bit 10
-DACK 7	D14	C14	Data Bit 11
DRQ 7	D15	C15	Data Bit 12
+5 Vdc	D16	C16	Data Bit 13
-Master	D17	C17	Data Bit 14
Ground	D18	C18	Data Bit 15

El bus ISA de 16 bits



BUSES

Lower Signal	Upper Signal	Pin	Pin	Upper Signal	Lower Signal
Ground	Ground	B1	A1	-I/O CH CHK	-CMD
+5 Vdc	RESET DRV	B2	A2	Data Bit 7	-START
+5 Vdc	+5 Vdc	B3	A3	Data Bit 6	EXADY
Reserved	IRQ 9	B4	A4	Data Bit 5	-EX32
Reserved	-5 Vdc	B5	A5	Data Bit 4	Ground
KEY	DRQ 2	B6	A6	Data Bit 3	KEY
Reserved	-12 Vdc	B7	A7	Data Bit 2	-EX16
Reserved	-0 WAIT	B8	A8	Data Bit 1	-SLBURST
+12 Vdc	+12 Vdc	B9	A9	Data Bit 0	-MSBURST
M-I/O	Ground	B10	A10	-I/O CH RDY	W-R
-LOCK	-SMEMW	B11	A11	AEN	Ground
Reserved	-SMEMR	B12	A12	Address 19	Reserved
Ground	-IOW	B13	A13	Address 18	Reserved
Reserved	-IOR	B14	A14	Address 17	Reserved
-BE 3	-DACK 3	B15	A15	Address 16	Ground
KEY	DRQ 3	B16	A16	Address 15	KEY
-BE 2	-DACK 1	B17	A17	Address 14	-BE 1
-BE 0	DRQ 1	B18	A18	Address 13	Latch Address 31
Ground	-Refresh	B19	A19	Address 12	Ground
+5 Vdc	CLK(8 33MHz)	B20	A20	Address 11	-Latch Address 30
Latch Address 29	IRQ 7	B21	A21	Address 10	-Latch Address 28
Ground	IRQ 6	B22	A22	Address 9	-Latch Address 27
Latch Address 26	IRQ 5	B23	A23	Address 8	-Latch Address 25
Latch Address 24	IRQ 4	B24	A24	Address 7	Ground
KEY	IRQ 3	B25	A25	Address 6	KEY
Latch Address 16	-DACK 2	B26	A26	Address 5	Latch Address 15
Latch Address 14	T/C	B27	A27	Address 4	Latch Address 13
+5 Vdc	BALE	B28	A28	Address 3	Latch Address 12
+5 Vdc	+5 Vdc	B29	A29	Address 2	Latch Address 11
Ground	OSC(14 3MHz)	B30	A30	Address 1	Ground
Latch Address 10	Ground	B31	A31	Address 0	Latch Address 9

Latch Address 8	-MEM CS16	D1	C1	-SBHE	Latch Address 7
Latch Address 6	-I/O CS16	D2	C2	Latch Address 23	Ground
Latch Address 5	IRQ 10	D3	C3	Latch Address 22	Latch Address 4
+5 Vdc	IRQ 11	D4	C4	Latch Address 21	Latch Address 3
Latch Address 4	IRQ 12	D5	C5	Latch Address 20	Ground
KEY	IRQ 15	D6	C6	Latch Address 19	KEY
Data Bit 16	IRQ 14	D7	C7	Latch Address 18	Data Bit 17
Data Bit 18	-DACK 0	D8	C8	Latch Address 17	Data Bit 19
Ground	DRQ 0	D9	C9	-MEMR	Data Bit 20
Data Bit 21	-DACK 5	D10	C10	-MEMW	Data Bit 22
Data Bit 23	DRQ5	D11	C11	Data Bit 8	Ground
Data Bit 24	-DACK 6	D12	C12	Data Bit 9	Data Bit 25
Ground	DRQ 6	D13	C13	Data Bit 10	Data Bit 26
Data Bit 27	-DACK 7	D14	C14	Data Bit 11	Data Bit 28
KEY	DRQ 7	D15	C15	Data Bit 12	KEY
Data Bit 29	+5 Vdc	D16	C16	Data Bit 13	Ground
+5 Vdc	-Master	D17	C17	Data Bit 14	Data Bit 30
+5 Vdc	Ground	D18	C18	Data Bit 15	Data Bit 31
-MAKx		D19	C19		-MREQx


Bus EISA

BUSES



Signal	Pm	Pm	Signal
AUDIO GND	B1	A1	-CD SETUP
AUDIO	B2	A2	MADE 24
Ground	B3	A3	Ground
OSC (14.3 MHZ)	B4	A4	Address 11
Ground	B5	A5	Address 10
Address 23	B6	A6	Address 9
Address 22	B7	A7	+5 Vdc
Address 21	B8	A8	Address 8
Ground	B9	A9	Address 7
Address 20	B10	A10	Address 6
Address 19	B11	A11	+5 Vdc
Address 18	B12	A12	Address 5
Ground	B13	A13	Address 4
Address 17	B14	A14	Address 3
Address 16	B15	A15	+5 Vdc
Address 15	B16	A16	Address 2
Ground	B17	A17	Address 1
Address 14	B18	A18	Address 0
Address 13	B19	A19	+12 Vdc
Address 12	B20	A20	-ADL
Ground	B21	A21	-PREEMPT
-IRQ 9	B22	A22	-BURST
-IRQ 3	B23	A23	-12 Vdc
-IRQ 4	B24	A24	ARB 00
Ground	B25	A25	ARB 01
-IRQ 5	B26	A26	ARB 02
-IRQ 6	B27	A27	-12 Vdc
-IRQ 7	B28	A28	ARB 03
Ground	B29	A29	ARB-GNT
Reserved	B30	A30	-TC
Reserved	B31	A31	+5 Vdc
-CHK	B32	A32	-SO
Ground	B33	A33	-S1
-CMD	B34	A34	M-IO
CHRDYRTN	B35	A35	+12 Vdc
-CD SFDBK	B36	A36	CD CHRDY
Ground	B37	A37	Data Bit 0
Data Bit 1	B38	A38	Data Bit 2
Data Bit 3	B39	A39	+5 Vdc
Data Bit 4	B40	A40	Data Bit 5
Ground	B41	A41	Data Bit 6
CHRESET	B42	A42	Data Bit 7
Reserved	B43	A43	Ground
Reserved	B44	A44	-DS 16 RTN
Ground	B45	A45	-REFRESH
KEY	B46	A46	KEY
KEY	B47	A47	KEY
Data Bit 8	B48	A48	+5 Vdc
Data Bit 9	B49	A49	Data Bit 10
Ground	B50	A50	Data Bit 11
Data Bit 12	B51	A51	Data Bit 13
Data Bit 14	B52	A52	+12 Vdc
Data Bit 15	B53	A53	Reserved
Ground	B54	A54	-SBHE
-IRQ 10	B55	A55	-CD DS 16
-IRQ 11	B56	A56	+5 Vdc
-IRQ 12	B57	A57	-IRQ 14
Ground	B58	A58	-IRQ 15
Reserved	B59	A59	Reserved
Reserved	B60	A60	Reserved

El bus MCA de 16 bits




BUSES

Signal	Pin	Pin	Signal	Signal	Pin	Pin	Signal
AUDIO GND	B1	A1	-CD SETUP	Reserved	B61	A61	Ground
AUDIO	B2	A2	MADE 24	Reserved	B62	A62	Reserved
Ground	B3	A3	Ground	Ground	B63	A63	Reserved
OSC (14.3 MHZ)	B4	A4	Address 11	Data Bit 16	B64	A64	Reserved
Ground	B5	A5	Address 10	Data Bit 17	B65	A65	+ 12 Vdc
Address 23	B6	A6	Address 9	Data Bit 18	B66	A66	Data Bit 19
Address 22	B7	A7	+5 Vdc	Ground	B67	A67	Data Bit 20
Address 21	B8	A8	Address 8	Data Bit 22	B68	A68	Data Bit 21
Ground	B9	A9	Address 7	Data Bit 23	B69	A69	+ 5 Vdc
Address 20	B10	A10	Address 6	Reserved	B70	A70	Data Bit 24
Address 19	B11	A11	+5 Vdc	Ground	B71	A71	Data Bit 25
Address 18	B12	A12	Address 5	Data Bit 27	B72	A72	Data Bit 26
Ground	B13	A13	Address 4	Data Bit 28	B73	A73	+ 5 Vdc
Address 17	B14	A14	Address 3	Data Bit 29	B74	A74	Data Bit 30
Address 16	B15	A15	+5 Vdc	Ground	B75	A75	Data Bit 31
Address 15	B16	A16	Address 2	-BE 0	B76	A76	Reserved
Ground	B17	A17	Address 1	-BE 1	B77	A77	+ 12 Vdc
Address 14	B18	A18	Address 0	-BE 2	B78	A78	-BE 3
Address 13	B19	A19	+12 Vdc	Ground	B79	A79	-DS 32 RTN
Address 12	B20	A20	-ADL	TR 32	B80	A80	-CD DS 32
Ground	B21	A21	-PREEMPT	Address 24	B81	A81	+ 5 Vdc
-IRQ 9	B22	A22	-BURST	Address 25	B82	A82	Address 26
-IRQ 3	B23	A23	-12 Vdc	Ground	B83	A83	Address 27
-IRQ 4	B24	A24	ARB 00	Address 29	B84	A84	Address 28
Ground	B25	A25	ARB 01	Address 30	B85	A85	+ 5 Vdc
-IRQ 5	B26	A26	ARB 02	Address 31	B86	A86	Reserved
-IRQ 6	B27	A27	-12 Vdc	Ground	B87	A87	Reserved
-IRQ 7	B28	A28	ARB 03	Reserved	B88	A88	Reserved
Ground	B29	A29	ARB/-GNT	Reserved	B89	A89	Ground
Reserved	B30	A30	-TC				
Reserved	B31	A31	+5Vdc				
-CHCK	B32	A32	-SO				
Ground	B33	A33	-S1				
-CMD	B34	A34	M/-IO				
CHRDYRTN	B35	A35	+12 Vdc				
-CD SFDBK	B36	A36	CD CHRDY				
Ground	B37	A37	Data Bit 0				
Data Bit 1	B38	A38	Data Bit 2				
Data Bit 3	B39	A39	+5 Vdc				
Data Bit 4	B40	A40	Data Bit 5				
Ground	B41	A41	Data Bit 6				
CHRESET	B42	A42	Data Bit 7				
Reserved	B43	A43	Ground				
Reserved	B44	A44	-DS 16 RTN				
Ground	B45	A45	-REFRESH				
KEY	B46	A46	KEY				
KEY	B47	A47	KEY				
Data Bit 8	B48	A48	+5 Vdc				
Data Bit 9	B49	A49	Data Bit 10				
Ground	B50	A50	Data Bit 11				
Data Bit 12	B51	A51	Data Bit 13				
Data Bit 14	B52	A52	+12 Vdc				
Data Bit 15	B53	A53	Reserved				
Ground	B54	A54	-SBHE				
-IRQ 10	B55	A55	-CD DS 16				
-IRQ 11	B56	A56	+5 Vdc				
-IRQ 12	B57	A57	-IRQ 14				
Ground	B58	A58	-IRQ 15				
Reserved	B59	A59	Reserved				
Reserved	B60	A60	Reserved				

El bus MCA de 32 bits

CPU Pin Assignments



Pin Name	Description	Type
ADD AD7	Address Data Bus	Bi-directional tri-state
AB A15	Address Bus	Output tri-state
A16 S3 A17 S4	Address Segment ident for	Output tri-state
A18 S5	Address interrupt enable status	Output tri-state
A19 S6	Address status	Output tri-state
S50	Status output	Output tri-state
RD	Read control	Output tri-state
RD GT0	Wait state request	Input
RD GT1	Wait for test control	Input
INTR	Interrupt request	Input
NMI	Non maskable interrupt request	Input
RESET	System Reset	Input
CLK	System Clock	Input
MN M2	- GND for a maximum system	
S0 S1 S2	Machine cycle status	Output tri-state
RD GT0 RD GT1	Local bus priority control	Bi-directional
Q50 Q51	Instruction queue status	Output
LOCK	Bus hold control	Output tri-state
MN M2	= VCC for a minimum system	
IO M	Memory output access	Output tri-state
WR	Write control	Output tri-state
ALE	Address Latch enable	Output
DT R	Data transmit/receive	Output tri-state
DT R	Data enable	Output tri-state
INTA	Interrupt acknowledge	Output tri-state
HOLD	Hold request	Input
HLDA	Hold acknowledge	Output
VCC	Power supply	



Minimum System Signals

8088 Pin and Signal Assignments



CPU 8088

2.2.1 Single Adapter Installation

Caution : Make absolutely sure that the computer is powered off before the installation procedure is started.

Avoid touching the components on the circuit board. A static electricity discharge may cause damage to the components on the board.

The SCSI adapter can be plugged into any of the 16 bit AT slots on the computer motherboard. Fasten the SCSI adapter metal bracket with a screw into the computer chassis.

If the floppy controller on the IN-2000 SCSI adapter is to be used, connect the floppy cable to the floppy connector marked 'J2' on the SCSI adapter. Enable the floppy disk controller by setting Sw1-9 in the 'on' position (see next section). Insure that there are no other floppy disk controllers enabled (including one on the motherboard).

2.2.2 Multiple Adapter Installation


It is possible to install up to three IN-2000 Host Adapters in one system. The installation procedure for the second and the third adapter is identical to the first adapter except for the default DIP switch settings.

Change the DIP switches for BIOS address on the 2nd and 3rd boards so that no conflicts exist. The IN-2000 BIOS will check for conflicts during the BOOT sequence and halt the system if any conflicts are found.

This completes the physical installation.

2.3 Switch Configuration

There is only one switch area on the SCSI adapter. This is a 10 position on/off switch named Sw1. This DIP switch controls all the different options available on the adapter.



BIOS Address

It is possible to select between 3 different BIOS segment addresses. It is also possible to disable the BIOS.

Sw1-1	Sw1-2	Segment address
off	off	c800h-cb00h *
on	off	d800h-dbffh
off	on	d000h-d3ffh
on	on	BIOS disabled

* factory default

I/O Port Address

I/O address conflicts are among the most common installation problems. The SCSI adapter allows the use of 4 different I/O port base addresses. It is NOT necessary to change BIOS prom before the I/O port address is changed.

Sw1-3	Sw1-4	I/O port address
off	off	220h-22fh *
on	off	200h-20fh
off	on	110h-11fh
on	on	100h-10fh

Interrupt Channel

Four different Interrupt channels can be selected. The

original hard disk controller uses Interrupt channel 14 as default. Do not use this channel if the original hard disk controller is installed in the system.

Sw1-5	Sw1-6	Sw1-7	Interrupt
off	off	off	15 *
off	on	off	14
off	off	on	11
off	on	on	10
on	-	-	disabled

IN PS/2 installations, use INT 10

Floppy Disk Controller

It is not possible to use the SCSI adapter floppy disk controller at the same time as the original floppy controller. Make sure that only one floppy disk controller is enabled in the system.

Sw1-9	Floppy controller
off	Disabled *
on	Enabled

* factory default

SW1-8 and SW1-10 not currently used.

2.4 SCSI cables

There are two SCSI connectors on the SCSI adapter. One internal 2x25 pin header and one external 25 pin



FACULTAD DE INGENIERIA U.N.A.M.
DIVISION DE EDUCACION CONTINUA

MANTENIMIENTO DE PC'S Y PERIFERICOS
(PARTE II)

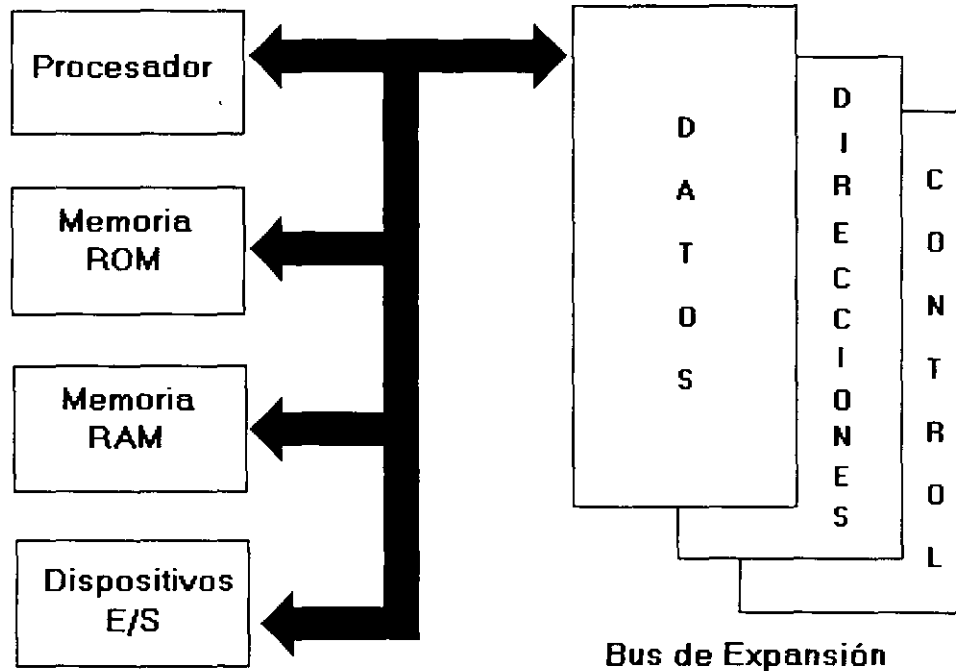
ARQUITECTURA DE UNA COMPUTADORA

JUNIO 2000

ARQUITECTURA DE LA PC

El diseño de la PC consiste de cinco áreas funcionales: El subsistema del procesador, el subsistema de memoria de solo lectura (**ROM**), el subsistema de memoria de lectura/escritura (**RAM**), los puertos de entrada/salida, y el **BUS** de expansión, llamado **PC-BUS**.

Arquitectura de una PC



El corazón del sistema es el microprocesador 8088 de Intel, este procesador opera a una velocidad de 4.77 MHz en la PC original, esta velocidad se obtiene de un cristal de 14.31818 MHz, dividido entre 3. Actualmente existe una gran cantidad de máquinas compatibles que además de operar a la frecuencia de 4.77 MHz, operan el llamado modo turbo, en frecuencias tan variadas como 7.15, 8, 9.54, 10 y 12 MHz (en máquinas XT), utilizando una pequeña variación en la arquitectura de la PC que permite el manejo de 2 o hasta 3 diferentes velocidades seleccionables.

El procesador recibe el soporte de un conjunto de circuitos que proveen: cuatro canales de acceso directo a memoria (**DMA**) de 20 bits, tres canales de 16 bits de un contador-timer, y 8 niveles de interrupciones manejables por prioridades.

Tres de los cuatro canales de **DMA** se encuentran en el **PC-BUS** y soportan transferencias a alta velocidad entre los dispositivos de e/s y la memoria sin la necesidad de la intervención del



microprocesador, el cuarto canal de **DMA** se usa para dar refresco a la memoria **RAM** dinámica. Esto se hace programando un canal del contador para que periódicamente pida una transferencia de **DMA** nula.

Los tres canales programables del contador-timer se usan en el sistema como sigue: Canal 0 se usa como un contador de propósito general, quien nos provee de una base de tiempo constante que nos permite llevar control de la fecha y de la hora; canal 1 se usa para pedir periódicamente los ciclos de refresco de memoria por **DMA**; canal 2 se usa para generar los tonos de la bocina de audio.

De los ocho niveles de interrupción disponibles, seis se encuentran en el **PC-BUS** para ser usadas por las tarjetas conectadas en él. Dos niveles se usan en la tarjeta principal, la interrupción con nivel 0, el de mayor prioridad, está conectada al canal 0 del contador y provee una interrupción periódica para el manejo del reloj/calendario. La interrupción no enmascarable (**NMI**) del 8088 se usa para reportar errores de paridad en la memoria.


Los circuitos de soporte más importantes dentro de la arquitectura de la **PC**, son los siguientes:

8259	: Controlador de Interrupciones.
8237	: Controlador de Acceso Directo a Memoria.
8284	: Generador de Reloj.
8255	: Interfaz con Periféricos Programable.
8253	: Contador Programable.
6845	: Controlador del CRT.
8272	: Controlador de Disco Flexible.

Cada uno de estos circuitos desempeña una labor específica e igual de importante dentro de la arquitectura de la **PC** y de su correcto funcionamiento depende la confiabilidad del sistema.

Una de las bondades de la **PC** es que su arquitectura abierta, permite al usuario y al diseñador de software o hardware hacer uso de cada una de sus partes funcionales, poniendo a su disposición las señales necesarias en el **PC-BUS** y los medios para software por medio del **BIOS**.







VISTA GENERAL DE UNA TARJETA PRINCIPAL

A continuación se da una descripción más detallada de los elementos principales de la arquitectura de la PC.

▣ ARQUITECTURA DEL 8086 / 88

El 8088 es un microprocesador de 16 bits que controla a todas las computadoras de las familias PC, PC/XT y compatibles,






El 8088 controla la operación básica de la **PC** enviando y recibiendo señales de control, direcciones de memoria y datos de una parte de la computadora a otra, a través de una red de conexiones o caminos a los que se les llama **BUS**. Localizados en alguna parte de este bus se encuentran los puertos de entrada/salida que conectan a todos los circuitos 000 de memoria o de soporte al bus

Dentro del 8088 existen 14 registros que nos proveen de un área de trabajo para la transferencia y el proceso de datos. Estos registros internos forman un área de 28 bytes en la cual se pueden guardar temporalmente datos, direcciones de memoria, direcciones de puertos, apuntadores e instrucciones y banderas de esta y del control. A través de estos registros el 8088 puede acceder hasta 1 Megabytes de memoria y hasta 64 kbytes de puertos de E/S.

El 8088 no es más que uno más de la familia de microprocesadores de 16 bits diseñados por Intel Corporation. El primer miembro de esta familia fue el 8086 El 8088 difiere del 8086 solo en un aspecto: Aunque el 8088 es un microprocesador de 16 bits, usa un bus de datos de 8 bits para comunicarse con el resto del sistema, a diferencia del 8086 que si tiene un bus externo de 16 bits. Prácticamente, toda la información que se refiere al 8086 se aplica al 8088 y viceversa; para propósitos de programación se consideran idénticos.





EL BUS

Como se mencionó antes, la familia de PCs se conecta internamente por medio de un bus, este bus no es más que un camino compartido al cual se conectan todas las partes funcionales de la computadora. Cuando se comunica un dato de una parte a otra de la computadora, pasa por este bus hasta llegar a su destino.

Todos los circuitos de control y toda localidad de memoria de la PC se conectan directa o indirectamente a este bus. Cuando un nuevo componente se inserta a la PC comparte este bus también, haciéndolo un usuario más de éste.



EL BUS DE DIRECCIONES

El bus de direcciones en el 8088 corresponde al de la **PC**, y utilizan 20 señales para transmitir las direcciones de memoria o de puertos de entrada/salida hacia todos los dispositivos que están conectados al bus.

Como para cada señal existen 2 diferentes valores (0 ó 1) las **PC** normales pueden direccionar hasta 2 a las 20 localidades, esto es 1 megabytes de memoria

EL BUS DE DATOS

El bus de datos trabaja en conjunto con el bus de direcciones para transferir datos a todas las partes de la computadora. El bus del 8088 es de 8 bits esto es los datos se transmiten en bytes (8 bits = 1 byte). El 8088 al ser un microprocesador de 16 bits, puede trabajar con 16 bits de datos a la vez, exactamente igual que el 8086. Pero aunque puede trabajar con datos de 16 bits, los transmite al exterior en dos partes de 8 bits. Por lo que mas de una vez se ha dicho que el 8088 no es realmente un microprocesador de 16 bits.

ORGANIZACIÓN INTERNA

El 8088 fue diseñado de forma tal que pueda procesar datos al mismo tiempo que recibe instrucciones y transmite datos a la memoria (Arquitectura de Pipeline). Para hacerlo usa registros de 16 bits. Existen 14 registros en total, cada uno con un uso especial. Cuatro registros de propósito general, en los cuales se guardan temporalmente resultados y operandos de operaciones aritméticas y lógicas. Cuatro registros de segmento, que guardan la dirección inicial de ciertos segmentos de memoria. Cinco registros índice y apuntadores que guardan la dirección relativa (offset) que se usan en conjunto con los segmentos para referirse a una localidad específica de memoria. Y por último, existe un registro de banderas, que contiene 9 banderas de 1 bit que se usan para guardar información de control y sobre el estado actual del 8088.

REGISTROS DE PROPÓSITO GENERAL

Los registros de propósito general son AX, BX, CX, y DX. Cada uno de ellos puede dividirse a su vez en dos registros de 8 bits, el de mayor valor o más significativo (AH, BH, CH, DH) y el de menor valor o menos significativo (AL, BL, CL, DL). El uso de los registros completos o de solo la mitad puede hacerse libremente, sin necesidad de seguir reglas especiales, como se necesiten.

Los registros de propósito general se usan en la mayoría de los casos como áreas temporales de trabajo, particularmente para operaciones aritméticas. Las sumas y restas se pueden hacer directamente en memoria, pero el uso de registros incrementa la velocidad de ejecución. Aunque estos registros están disponibles para cualquier tipo de trabajo, cada uno de ellos tiene algunos usos particulares como por ejemplo.

- ☐ **AX** Es el acumulador, y es el registro principal para efectuar operaciones aritméticas
- ☐ **BX** Este registro (base) se usa comúnmente como apuntador al inicio de una tabla en memoria. También puede usarse para la dirección relativa (offset) de una dirección segmentada.



- ☐ **CX** : Este registro (cuenta) se usa como un contador de repeticiones en instrucciones de ciclos o transferencias de bloques de memoria. Como por ejemplo la instrucción **LOOP** usa a **CX** para llevar el control de cuantas veces ha ejecutado el ciclo. Ninguno de los otros registros puede usarse para ello.
- ☐ **DX** : Este registro solo se usa para guardar datos de 16 bits para diferentes propósitos. Es decir, no tiene ningún uso especial.

☐ LOS REGISTROS DE SEGMENTO

Para poder conocer el uso de estos registros, es necesario conocer como maneja el 8088 las direcciones de memoria. El 8088 es un microprocesador de 16 bits, por ello no puede trabajar con números cuya representación necesite más de 16 bits, el número decimal más grande que puede manejar es 65,535 es decir 64 K. Teóricamente esto significaría que el 8088 podría acceder cuando mucho 64 K localidades de memoria. Pero como mencionamos anteriormente, el 8088 puede acceder bastante más que eso, es decir hasta 1,024 K. Esto es posible dado que se manejan 20 líneas de direcciones. Pero el 8088 está limitado por su capacidad de manejar solo 16 bits en sus registros. Para manejar 20 bits de direcciones debe usarse un método que use solo 16 bits.

El 8088 divide el espacio de memoria direccionable en un número arbitrario de **SEGMENTOS**, cada uno de los cuales no contiene más de 64 kbytes. Cada segmento empieza en una localidad que es divisible entre 16 bytes a la que se le llama la dirección del segmento o párrafo (paragraph). Para acceder localidades individuales se usa una dirección adicional llamada dirección relativa (offset) que apunta a un byte específico dentro del segmento de 64 kbytes designado por la dirección del segmento

Las direcciones se crean y se manipulan combinando un segmento de 16 bits y una dirección relativa de 16 bits. El segmento se maneja como si estuviera recorrido a la izquierda 4 bits, y después sumado a la dirección relativa de 16 bits. Lo que nos da la dirección absoluta o vector de 20 bits

	Dirección del
Segmento	1 0 1 1, 1 0 1 1, 1 0 1 0, 0 0 1 1, 0 0 0 0
	Dirección Relativa
(offset)	1 0 1 1 1 0 1 0 0 1 1 0 0 1 1 1
	Dirección Absoluta
(Vector)	1 1 0 0 0 1 1 1 0 1 0 0 1 0 0 1 0 1 1 1

Por ejemplo, para obtener la dirección absoluta de un segmento 1234 H y un offset 4321 H.

```

12340 (1234 corrido 4 bits a la izquierda)
+ 4321 (el offset)
-----
16661 La dirección de 20 bits (en hexadecimal)

```



Debido a esto, los 4 registros llamados de segmento: CS, DS, SS Y ES se usan para identificar 4 segmentos de 64 kbytes específicos. De los 4 registros de segmento, los siguientes tres se dedican a propósitos especiales:

- ☐ **CS.** Localiza el segmento de código, que contiene el programa que se está ejecutando.
- ☐ **DS.** Localiza el segmento de datos, que es el área de memoria donde se guardan los datos que están siendo usados.
- ☐ **SS.** Localiza el segmento de stack, un espacio de trabajo temporal que mantiene datos acerca de parámetros y direcciones en uso por el programa que se está ejecutando

☐ LOS REGISTROS DE OFFSET

Existen cinco registros de offset que se usan para localizar un byte o una palabra específicos dentro de un segmento de 64 kbytes.

Uno de los registros, el apuntador de instrucciones (IP), localiza la instrucción del programa que se está llevando a cabo, dentro del segmento de código; otros dos los registros de stack, están ligados íntimamente con el stack, que es el área de memoria que el 8088 utiliza para guardar datos o direcciones que necesita recordar para su uso posterior; y los dos registros restantes, los registros índice se usan para apuntar datos que se están usando en el segmento de datos.

- ☐ **IP:** El apuntador de instrucciones, nos da la dirección relativa con respecto al segmento de código. Se usa junto con el CS para mantener la dirección de la siguiente instrucción a ser ejecutada. Los programas no tienen acceso directo al IP, pero existe en algunas instrucciones como JMP o CALL, que cambian el IP directa o indirectamente.
- ☐ **SP & BP:** Los registros de stack, nos dan direcciones relativas dentro del segmento de stack. El SP (apuntador de stack) da la dirección del fin de stack y es análogo al IP. El BP (apuntador base) se usa para guardar localidades en las cuales tenemos datos específicos, que vamos a usar en un momento cercano, el BP es particularmente útil cuando se manejan parámetros entre rutinas o programas escritos en diferentes lenguajes de programación. Como por ejemplo, cuando se desea utilizar rutinas escritas ensamblador dentro de un programa escrito en PASCAL.
- ☐ **SI & DI.** Los registros índice, SI (índice de fuente) y DI (índice del destino) se usan normalmente junto con algunos de los registros de propósito general para apuntar al inicio de alguna tabla de datos almacenada en memoria, son particularmente útiles en la transferencia de datos por bloques o de cadenas de caracteres, usualmente, SI y DI incrementan sus valores automáticamente al ser usadas con ciertas instrucciones para los fines descritos anteriormente.

☐ EL REGISTRO DE BANDERAS

El último de los registros del 8088, el registro de banderas, es realmente un conjunto de banderas de un bit. Estas banderas se agrupan en un solo registro de forma que se pueden acceder como



un conjunto, de una sola vez, no es necesario el leer estas banderas de una en una, aunque el valor de estas banderas debe revisarse bit por bit.

Existen nueve banderas de un bit en este registro de 16 bits, dejando 7 bits sin usar. Estas banderas se dividen en dos grupos, las banderas de estado (6) que se usan para indicar el estado de el procesador después de ejecutar una instrucción, generalmente se usan para saber el resultado de una operación aritmética o de una comparación, y las otras 3 se usan como banderas de control, es decir en algunos casos el procesador irá a revisar el estado de estas banderas para saber que hacer en algunas instrucciones. Este registro contiene las siguientes banderas.

Tabla Banderas de Estado

Bandera	Nombre	Uso
CF	Acarreo	Indica que existió en un acarreo de un bit en una operación aritmética
OF	Sobreflujo	Indica un Sobreflujo aritmético
ZF	Cero	Indica un resultado de cero o un igual en una comparación
SF	Signo	Resultado negativo en una operación o comparación
PF	Paridad	Indica un número par de bits 1
AF	Acarreo	Indica que se necesita un ajuste en operaciones en BCD

Tabla Banderas de Control

Bandera	Nombre	Uso
DF	Dirección	Dirección hacia la izquierda o derecha al procesar cadenas de caracteres
IF	Interrupción	Habilita o Deshabilita las Interrupciones
TF	Trampa	Controla la operación instrucción por instrucción, introduciendo una trampa al finalizar cada instrucción (como lo hace DEBUG)

MEMORIA Y PERIFÉRICOS

Estas dos secciones son básicas en la arquitectura de la PC, pues en la memoria se almacena la información necesaria para el control de cualquiera de los procesos que se llevan a cabo en la PC. Y los periféricos son la forma en la que la PC se conecta con el mundo exterior, sin ellos la existencia de la PC no serviría para nada, pues por medio de ellos, se alimenta a la computadora con los datos necesarios para nuestros procesos, y por medio de ellos también es como nos enteramos de los resultados generados por estos procesos.



▣ DIRECCIONAMIENTO DE ENTRADA/SALIDA

El 8088 se comunica y controla la mayor parte de la **PC** mediante el uso de los puertos de entrada/salida. Los puertos de entrada/salida son vías por las cuales la información pasa en su camino hacia o desde un dispositivo de entrada/salida, como por ejemplo un teclado o una impresora. La mayor parte de los circuitos de soporte descritos anteriormente, usan estos puertos de e/s; de hecho, cada circuito puede usar varios de estos puertos para diferentes propósitos

Cada puerto se identifica con una dirección específica de 16 bits, esto es podemos manejar hasta 64k diferentes puertos dentro de la **PC**. El microprocesador envía datos o información de control a un puerto específico usando su dirección, y el puerto responde pasando datos o información de estado de regreso al micro.

De la misma forma que al acceder la memoria, el CPU usa los buses de datos y direcciones como conductos para la comunicación con los puertos. Para acceder un puerto, el microprocesador envía una señal por el bus de control que notifica a todos los puertos que la dirección que se envía corresponde a un puerto, y envía la dirección.

La dirección de un puerto, o el número de un puerto se maneja de forma similar a una localidad de memoria.

De hecho, algunas tarjetas de expansión hacen uso tanto de puertos como de algunas áreas de memoria, como por ejemplo las tarjetas de vídeo.

En la tabla siguiente se muestran los puertos que se encuentran asignados dentro del diseño de la **PC**:

Tabla Asignación de Puertos de E/S

Rango (Hex)	Función
000-00F	Circuito de DMA 8237A
020-021	Controlador de Interrupciones 8295A
040-043	Contador 8053
060-063	Interface Programable 8255A
080-083	Regs. de página del DMA
0Ax	Regs. de máscara de la NMI
0Cx	Reservado
0Ex	Reservado
100-1FF	No puede usarse
200-20F	Control de Juegos
201-217	Unidad de Expansión
220-24F	Reservado
278-27F	Reservado
2F0-2F7	Reservado
2F8-2FF	Puerto Serie Secundario
300-31F	Tarjeta para Prototipos
320-32F	Controlador de Discos Duros
378-37F	Puerto Paralelo (LPT1)
380-38C	Comunicaciones SDLC



380-389	Comunicaciones BCS (Secundarias)
3A0-3A9	Comunicaciones BCS (Primarias)
3B0-3BF	Pto. paralelo de la tarjeta monocromática
3C0-3CF	Reservado
3D0-3DF	Tarjeta de Color
3E0-3EF	Reservado
3F0-3F7	Controlador de Discos Flexibles
3F8-3FF	Puerto Serie (Primario, COM1)

☐ DIRECCIONAMIENTO DE MEMORIA

Como mencionamos anteriormente el 8088 tiene un esquema de manejo de memoria por segmentos, lo que permite que el espacio de memoria aprovechable sea de hasta 1 Mbyte.

La arquitectura de la PC tiene contempladas básicamente dos áreas de memoria para su funcionamiento, la memoria de solo lectura (ROM), y la memoria de lectura/escritura.

☐ MEMORIA ROM

En esta área de memoria se encuentra contenido el programa que controla el funcionamiento básico de la PC, es decir, aquí se encuentra el programa de inicialización, cuya función básica es la de Realizar diagnósticos a los circuitos principales, detectar que equipo periférico tiene conectado al PC, programar todos los circuitos de soporte como son: el detectar si existe algún otro ROM que vaya a efectuar alguna función especial, y finalmente cargar el sistema operativo de disco. Ver figura referente a la asignación del mapa de memoria ROM en ANEXO 1

La forma en que el sistema detecta si existe alguna memoria ROM o EPROM instalada en el área de expansión de 192k, es buscando cada 16k a partir de C0000 si encuentra los bytes AA 55 (hex), al encontrarlos, transfiere el control al cuarto byte del ROM (pues el tercero deberá tener de que tamaño es el ROM en bloques de 512 bytes), el programa almacenado en el ROM se deberá encargar de regresar el control al programa de inicialización usando un RET FAR. Estos 192k, se deberán acceder desde alguna tarjeta colocada en el bus, tal como lo hace la controladora de disco duro o el BIOS de las tarjetas de vídeo EGA o VGA.

El área entre las direcciones F0000 y FE000 marcada como memoria para el usuario, se encuentra ya en la tarjeta principal, en las bases para ROM que se encuentran vacías.

☐ MEMORIA RAM

La memoria RAM, es la memoria de lectura/escritura donde la máquina almacena todos los datos que usa mientras esta máquina almacena todos los datos que usa mientras está encendida y ejecutando un programa, aquí se mantienen nuestros programas, datos, parámetros del sistema operativo e incluso se guarda lo que vemos en la pantalla. Ver figura referente a la disposición del mapa de memoria RAM en el ANEXO 1

Como se aprecia en el mapa de memoria, el área aprovechable para los programas del usuario, así como para el sistema operativo, se limita hasta la dirección 9FFFF, es decir hasta 640 kbytes de memoria, pues el resto es utilizado por las tarjetas de vídeo y por el área para ROM.



Esta limitación de 640 kbytes se puede evitar usando tarjetas de expansión que se direccionan como bancos, y no como memoria continua, que es como se maneja el estándar de memoria expandida de Intel, Microsoft y Lotus (EMS).

☐ MANEJO DE INTERRUPCIONES

La forma en la que el 8088 puede comunicarse con los periféricos que componen la PC es por medio de interrupciones, en cualquier momento que un dispositivo necesita la asistencia del microprocesador, envía una señal o instrucción llamada interrupción, identificando la tarea que desea que el micro desempeñe.

Cuando el microprocesador recibe esta señal de interrupción, detiene sus actividades e inicia la ejecución de una subrutina almacenada en memoria, ya sea RAM o ROM, a la que se llama el manejador de la interrupción, este manejador corresponde a un número determinado. Después de que el manejador de la interrupción termina su tarea, las actividades del sistema continúan a partir del punto donde fueron interrumpidas.

Existen tres categorías de interrupciones: primero, las interrupciones generadas por la circuitería de la PC como respuesta a algún evento, como la presión de una tecla o la recepción de un dato por el puerto serie.

Estas interrupciones se manejan por medio del circuito controlador de interrupciones, el 8259, que les da prioridades en orden de importancia antes de enviarlas al microprocesador para su proceso. Segundo, las interrupciones generadas por errores imprevistos en las aplicaciones como por ejemplo la división entre cero.

Estas dos categorías de interrupciones son las llamadas interrupciones por hardware y se clasifican de la siguiente forma:

Asignación de las Interrupciones

Número	Causa
NMI	Error de Paridad
0	Contador
1	Teclado
2	Reservada
3	Comunicaciones por puerto serie (COM2), SDLC o BSC (Secundaria)
4	Comunicaciones por puerto serie (COM1), SDLC o BSC (Primaria)
5	Disco Duro
6	Disco Flexible
7	Puerto Paralelo

La tercera categoría de las interrupciones, son aquellas generadas deliberadamente por los programas de aplicación como una manera de llamar subrutinas lejanas y de uso constante que se encuentran en RAM o ROM. Estas rutinas son usualmente parte del BIOS o del MS-DOS, éstas pueden ser modificadas incluso crear algunas nuevas para usos particulares.




De cualquier forma en que una interrupción sea generada, el dispositivo que la genera no necesita saber donde está el manejador que le dará servicio; lo único que necesita conocer es el número de interrupción. Este número sirve como referencia a una tabla almacenada en RAM, en las localidades mas bajas, en donde se encuentra la dirección segmentada de el inicio del manejador, a esta dirección se le llama vector de interrupción, y será tratado con detalle más adelante.

Las interrupciones automáticamente guardan el segmento de código que se está usando, así como el apuntador a las instrucciones, para que al terminar el manejador el micro sepa a donde regresar.

EL PC-BUS

El bus de la PC es una extensión del bus del 8088. Se encuentra ya demultiplexado, con líneas para la fuente de poder y con el manejo de interrupciones y de acceso directo a memoria. El bus contiene un bus de datos de 8 bits, 20 líneas de direcciones, 6 niveles de interrupciones, líneas de control para memoria y los puertos de e/s, relojes, 3 canales de DMA, etc. Estas funciones se encuentran en un conector de 62 contactos.



ASPECTO TIPICO DEL BUS

Se tiene disponible una señal de "ready" para permitir la operación de tarjetas de expansión lentas, ya sea de e/s o de memoria. Si esta señal de "ready" no se activa, todos los ciclos de lectura o escritura a memoria generados por el procesador, duran 4 pulsos de reloj es decir 840 ns/byte

Los ciclos de lecturas o escrituras a los puertos de e/s generados por el procesador, requieren 5 pulsos de reloj para completar un tiempo de 105 μ s por byte. Los ciclos de refresco ocurren uno cada 72 pulsos de reloj (aproximadamente 15 μ s) y requieren de 4 ciclos de reloj, es decir ocupan aproximadamente el 7 % del ancho de banda del bus

El bus está diseñado para dar acomodo a 512 puertos de entrada/salida para el uso de las tarjetas de expansión, a pesar de que el 8088 podría manejar hasta 64 K. Existe una línea llamada "channel check" que reporta condiciones de error al procesador. Al activar esta línea se



genera una NMI (Interrupción no enmascarable) para el 8088. Las tarjetas para expansión de memoria generalmente usan esta línea para reportar errores de paridad.

El PC-Bus está diseñado para dar suficiente alimentación a todas las ranuras de expansión, para lo cual se asumen 2 cargas LS (low-power shottky) por ranura máximo. Por lo general las tarjetas de buena calidad sólo tienen una carga LS.

A continuación se da una explicación más detallada de cada una de las señales del BUS (ver figura siguiente):

EL PC - BUS

Señal	E / S	Utilización
OSC	S	Oscilador: Reloj de alta velocidad con un periodo de 70 ns. Tiene un duty del 50%
CLK	S	Reloj del sistema. Es una división entre 3 del oscilador y tiene un periodo de 210 ns. (4.77Mhz.) Tiene un duty del 33%
RESET	S	Esta línea se usa para inicializar la lógica del sistema al encender la máquina o durante una baja en el voltaje de alimentación. Está sincronizada con la bajada del reloj y es activo alto.
A0-A19	S	Bits 0 a19 de las direcciones: Estas líneas se usan para manejar la memoria y los dispositivos de E/S del sistema
D0-D7	E/S	Bits de Datos 0 de 7 : Es el bus de datos. Son activo alto
ALE	S	Habilita Direcciones. Esta línea es generada por el controlador del bus (8288) y se usa para indicar cuando existen direcciones válidas dadas por el procesador. Está disponible en el PC-BUS, y se usa en general en conjunto con AEN. Activo alto.
I/O CHK	E	Revisión del canal: Da la información sobre errores en las tarjetas del PC-BUS y se usa generalmente errores de paridad. Activo bajo.
I/O RDY	E	Canal Listo: Esta línea, activo alto, se coloca en cero por alguna tarjeta cuando ésta necesita tener un ciclo de lectura o escritura más largo, ya sea en memoria o en E/S. No puede tener abajo más de 10 ciclos de reloj.
IRQ2-IRQ7	E	Pedidos de Interrupción 2 a 7: Se usan para señalar al procesador que un dispositivo de E/S requiere atención. Están dispuestas por prioridades, teniendo IRQ2 la máxima prioridad e IRQ7 la mínima. Una interrupción se genera levantando una línea de IRQ (.de 0 a 1) y manteniéndola en alto hasta recibir el aviso de que se captado o reconocido la interrupción



IOR	S	Lectura de Dispositivo de E/S: Esta línea se usa para avisar al dispositivo, que debe poner un dato en el bus del procesador. Puede ser generada por el procesador o por el controlador de DMA. Esta señal es activo bajo.
IOW	S	Escritura a dispositivo de E/S: Esta línea se usa para avisar al dispositivo que debe leer el dato que el procesador colocó en el bus . Puede ser generada por el procesador o por el controlador DMA. Activo bajo.
MEMR	S	Lectura de Memoria: Avisa a la sección de memoria que debe poner un dato en el bus. Puede ser generada por el procesador o por el controlador de DMA. Activo bajo.
MEMW	S	Escritura a Memoria: Avisa a la sección de memoria que debe leer el dato que el procesador ha puesto en el bus. Puede ser generada por el procesador o por el controlador de DMA. Activo bajo.
DRQ1-DRQ3	E	Pedido de DMA 1 a 3: Son pedidos asíncronos usados por dispositivos, periféricos para lograr acceso directo a memoria.
DACK0 DACK3	S	Reconocimiento de DMA: Estas líneas se usan para avisar que la petición a DMA ha sido concedida (para drq1-drq3) y para el refresco de la memoria dinámica del sistema (DACK0). Son de activo bajo.
AEN	S	Habilita Dirección: Esta línea se usa para desconectar al procesador y otros dispositivos del PC-BUS para permitir el acceso directo a memoria. Cuando esta línea se coloca en activo alto, quiere decir que el controlador de DMA tiene el control del bus de datos, y de direcciones de las líneas de escritura y lectura tanto a memoria como a E/S.
T/C	S	Cuenta Terminal: Esta línea provee un pulso cuando el DMA de alguno de los canales debe terminar. Es de activo alto.

ESTÁNDARES DE BUS

Inicialmente sólo hubo un diseño de bus, el de la IBM. Posteriormente, otros fabricantes y la IBM introdujeron nuevos diseños de computadoras con diferentes configuraciones de bus. Por mucho, el más popular de éstos sigue siendo el ISA (Industry Standard Architecture), basado en el diseño de la IBM AT original. Sin embargo, hay otros diseños que también tienen seguidores.

ISA: ARQUITECTURA DE LA INDUSTRIA ESTÁNDAR

El bus estándar (ISA), algunas veces llamado Bus AT es el más popular y común en los diseños de la PC. Es un bus de datos de 16 bits basados en un diseño de conector de expansión de 98 patas. De manera similar a la mayoría de los diseños de bus, el ISA emplea conectores de doble lado. Cuando se enchufa (inserta) en una tarjeta de expansión, cada posición de conector son, de hecho, dos conectores, una del lado A de la tarjeta y otro del lado S. La figura 6 muestra el esquema de la orilla del conector de una tarjeta de expansión.



El lado de los componentes de la tarjeta de expansión lleva los conectores para las patas AI-A31 (el socket principal y parte de bus de la PC original) y CI-C18 (el socket extendido fue añadido con el modelo AT de la PC). El reverso de la tarjeta tiene las conexiones para las patas BI-B31 y DI-D18. Las A2-A9 son las primeras ocho líneas de datos y las patas C11 -C18 son el segundo juego de ocho líneas de datos. Como podrá observarse, el bus de la PC original contenía ocho líneas de datos de sus 64 conectores, las patas A2-A8. Con la PC AT, fueron añadidos ocho líneas de datos adicionales en el segundo renglón de conectores.

Si se observa algunas tarjetas de expansión, se verá que ninguna llega a usar 98 patas.

Pueden tenerse dos tipos básicos de tarjeta de expansión diseñadas para el bus ISA, tarjetas de 16 y 8 bits. Si se ve la figura 7, en ella se muestra el arreglo de doble conector de los adaptadores para el bus. La mayoría de las tarjetas de ocho bits son realmente cortas y tienen conectores solamente para el primer renglón de patas.

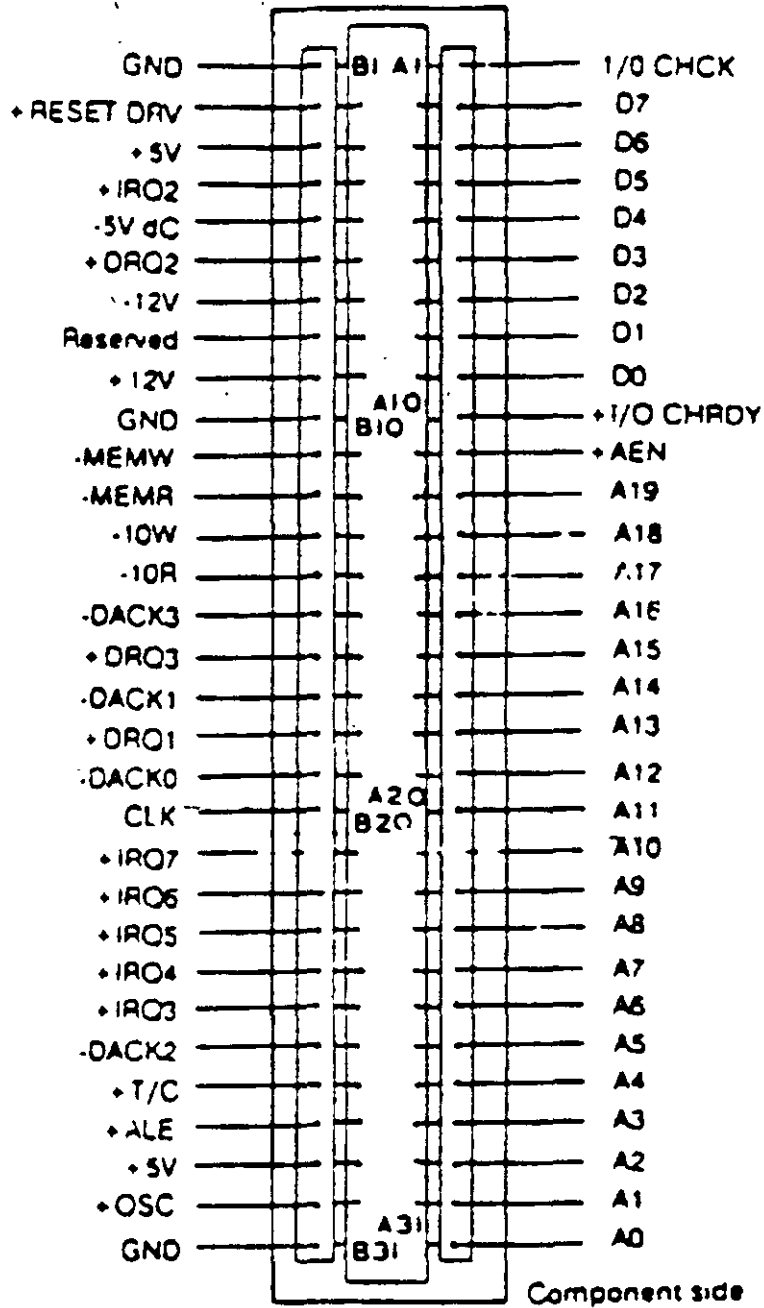
Por otro lado, una tarjeta de 16 bits tiene protuberancias y patas que se enchufan en la posición del conector del otro renglón.

Los estándares de bus ISA DE 16 Bits es adecuado para los procesadores 286 y 386SX, que pueden transferir datos de 16 bits a la vez y direccionar hasta 16 MB. Sin embargo, los procesadores 386 y posteriores pueden transferir hasta 32 bits a la vez y direccionar hasta 4 GB. Para obtener esto, se necesita un bus de 32 bits.

El bus MCA de 32 bits acepta adaptadores con 93 conectores, para dar un total de 186 líneas de señal. Por lo general, el bus MCA de 32 bits puede verse como un bus de 16 bits con líneas de señal extras. En primer lugar, hay 32 líneas de dirección (AO-A31), que proporciona hasta 232 direcciones diferentes, lo que significa que el procesador puede direccionar hasta 4 GB de memoria. También hay 32 líneas de datos (DO-D31), que permiten la transferencia simultánea de hasta 32 bits (4 bytes). Por lo tanto, el bus de 32 bits puede manejar la capacidad completa de los procesadores 386 y 486



REAR PANEL



I/O CHANNEL SLOT

DESCRIPCION DEL SLOT XT ISA

Poco después de que la IBM anunció el MCA, un grupo de compañías que fabricaban computadoras compatibles con la IBM decidieron crear una alternativa. Esta alternativa la llamaron EISA (Extended Industry Standard Architecture)



EISA: LA ARQUITECTURA INDUSTRIAL ESTANDAR EXTENDIDA

La IBM anunció el MCA junto con su línea PS/2 de PC en junio de 1987. En septiembre de 1988, un consorcio de nueve de compañías, dirigidas por Compaq, anunció que estaba desarrollando una alternativa al MCA. Este consorcio votó por mantener la venta de computadoras al estilo anterior, basadas en ISA. La alternativa a MCA fue la arquitectura extendida (EISA)

Desde el principio estuvo claro que el desarrollo de EISA estaba basado en las necesidades de ventas y no de Ingeniería. A partir del día de su publicación, se llevó cerca de dos años y medio para que los primeros adaptadores EISA llegaran al mercado. Y a la fecha existen máquinas EISA. La principal ventaja de EISA fue que permitía a los usuarios emplear sus tarjetas antiguas tipo PC/XT/AT en sus nuevas computadoras, cosa que no podían hacer con las computadoras MCA.

EL BUS LOCAL

Una solución del desplegado de video y el problema general del envío de datos y recibir datos desde el procesador hacia cualquier dispositivo de ancho de banda grande, es un diseño de bus local. Con los diseños convencionales, todo lo que viene y todo lo que va al procesador (o casi todo) se envía a través del bus principal del sistema. Como se debe mantener compatibilidad hacia atrás con los diseños anteriores, y debido a que se tiene que trabajar con un rango alto de dispositivos periféricos, este es un bus relativamente lento y con ancho de banda limitado, aún con bus de 32 bits.

Los sistemas de bus local se comunican a la velocidad del sistema, hasta 80 MHz, y siempre manejan datos en paquetes de 32 o 64 bits. Los diseños de bus local están llevando a las computadoras actuales un paso más allá del camino hacia el alto rendimiento, sin cambiar en mucho acerca de la ingeniería básica de la máquina. Las máquinas de bus local quitan del bus principal las interfaces que necesitan una respuesta rápida: memoria, video y unidades de disco. Conforme los requisitos de E/S se tomen más importantes, es probable que la conexión a la red, el audio y otras funciones sean añadidas al ambiente de bus local.

Los diseños de bus local en su mayoría han sido propiedad del hardware de un vendedor en particular. No existen muchos estándares con respecto a los buses locales pero los principales, el bus VESA VL y el bus PCI.

EL BUS LOCAL PCI

El bus local es la manera en que las computadoras actuales están aumentando el rendimiento sin ningún adelanto real de ingeniería.

El estándar PCI está diseñado y mantenido por el Peripheral Component Interconnect Special Interest Group, o PCI SIG, una asociación de representantes de la industria de minicomputadoras sin incorporar. El bus local PCI puede tener una ruta de 32 o 64 bits para transferir datos a alta velocidad. Soporta ambientes de escritorio de 5 y portátiles de 3.3 volts. Idealmente PCI puede trabajar con las PC así como otros diseños de computadoras. Como el diseño no depende de la familia ipx86 de procesadores, de acuerdo con el PCI SIG, puede trabajar con las PC actuales y con diseños futuros, sin tomar en cuenta el procesador utilizado.






FACULTAD DE INGENIERIA U.N.A.M.
DIVISION DE EDUCACION CONTINUA

MANTENIMIENTO DE PC'S Y PERIFERICOS


(PARTE II)

DISCOS DUROS



JUNIO 2000



Presenta: Maya Tecnologías de
la Información. Soluciones y
Capacitación.



Notas:





Tecnologías de Discos Duros

- Factores de Operación
 - Formato Físico
 - Formato Lógico
 - Partición
 - Instalación
 - Maestro
 - Esclavo



Notas:





Tecnologías de Discos Duros

- Formato Físico
 - Se definen
 - Cabezas
 - Cilindros
 - Sectores
 - Clusters
 - Inicialización de la superficie a 1's



Notas:

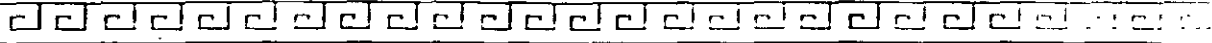



Tecnologías de Discos Duros

- Formato Lógico
- Se definen zonas de administración del S.O.
 - Boot
 - Sistema de Archivos (FAT-16, FAT-32, NTFS, INODOS)
 - DIR
 - DATOS



Notas:





Tecnologías de Discos Duros

- Métodos de Registro
 - MFM (modulación de Frecuencia modificada)
 - RLL (Longitud de Corrida limitada)
 - ZBR (Registro de bits por zona)
 - IDE
 - ESDI
 - SCSI


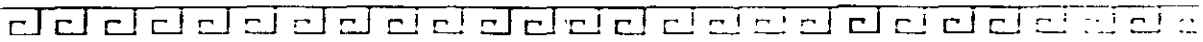


Notas:

Tecnologías de Discos Duros

Tipo	Métodos de Registro	
	Sectores por Pista	Velocidad de Mbits/s
MFM	17	5
RLL	56	7.5
ARLL	34	10
ZBR	Varia	7.5-15

Notas:





Tecnologías de Discos Duros

- Tipos de Interfaces
 - ST 506/412
 - ESDI
 - IDE
 - SCSI



Notas:




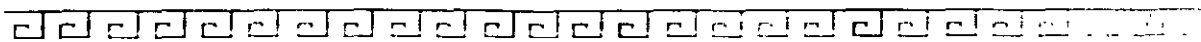
Tecnologías de Discos Duros

- Tipos de Interfaces

- ST 506/412
 - Tecnología de los 80s
 - 152 MB capacidad máxima para MFM
 - 152 MB capacidad máxima para MFM
- ESDI Enhanced Small Device Interface
 - Estandar establecido en 1983
 - Velocidad de Transferencia de 24 Mbit/sec
 - 32 sectores por track
 - Interlevo 1:1 a velocidades de 1 Mbit/sec
- Ambas Tecnologías con Tarjetas Controladoras



Notas:



Tecnologías de Discos Duros

- Tipos de Interfaces

- IDE Integrated Drive Electronics




Notas:

ST31621A
Cabo Stingray Family

2000
10
10



Modelo FE
Modelo FA
Modelo EE
Modelo EE CE
Modelo EE CE
Modelo EE CE
Modelo EE CE
Modelo EE CE



Modelo Slave	Slave	Modelo Master	Master
EE	EE	EE	EE
EE	EE	EE	EE
EE	EE	EE	EE
EE	EE	EE	EE

Disco Duro IDE

Notas:





Tecnologías de Discos Duros

- Tipos de Interfaces IDE

- Normalizada como ATA (AT Attachment)
 - La Controladora Integrada en el Dispositivo
 - Conector IDE integrado a los motherboards
 - Conector 40 pines
 - Versiones de bus
 - » AT IDE 16 Bits
 - » XT IDE 8 Bits
 - » MCA IDE 16 Bits



Notas:


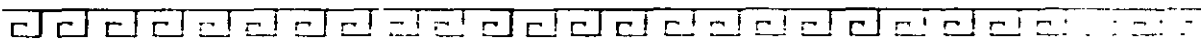


Tecnologías de Discos Duros

- ATA Comands
- Los comandos ATA identifican principalmente
 - Identificación de la unidad (Identify Drive)
 - Numero de Cilindros
 - Numero de cabezas
 - sectores por track
 - Firmware revision
 - Numero de Serie
 - Buffer type (Chache)



Notas:



Tecnologías de Discos Duros

- Parámetros actuales
- Interface IDE
 - Normalización ANSI en 1994 para ATA-1
 - Normalización ANSI en 1995 para ATA-2 (EIDE)
 - ATA - 3 en 1996
 - ATA - 4 en 1997
 - ULTRA ATA 1998




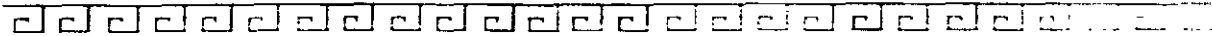
Notas:

Tecnologías de Discos Duros

- Interface IDE

Specification	ATA	ATA-2	ATA-3	ATA/ATA-16
Features	PIO-1	PIO-2 DMA-0,1,2	PIO-3 DMA-3,1,2	PIO-4 DMA-3,1,2 UDMA
Max Transfer Rate	4 Mbytes/sec	16 Mbytes/sec	33 Mbytes/sec	33 Mbytes/sec
Max Connections	2	2	2	2 per channel
Cable Length	18 inches	18 inches	18 inches	18 inches
Plug and Play	no	yes	yes	yes
Additional Functions	None	ATAPI	ATAPI Serial ATA Termination	ATAPI DMA-3 ATAPI
Year Introduced	1981	1994	1997	1997

Notas:



Tecnologías de Discos Duros

- Interface IDE
 - Parámetros actuales
 - PIO
 - DMA
 - ULTRA DMA
 - S M A R T
 - ATAPI



Notas:

Tecnologías de Discos Duros



– Interface IDE

• Parámetros actuales

– PIO mode

» pio Int	Cycle Time	M/sec
» 0 ATA	600	3.3
» 3 ATA-2, pci	180	11.1
» 4 ATA-2, pci	120	16.6

Notas:


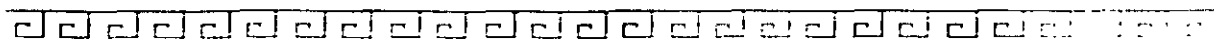


Tecnologías de Discos Duros

- Interface IDE
 - Parámetros actuales
 - DMA
 - » busmastering
 - ULTRA DMA
 - » UDMA
 - » velocidades hasta de 33 mbits/sec



Notas:




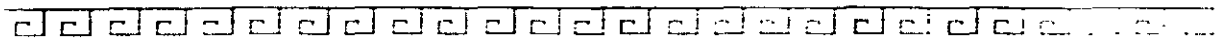
Tecnologías de Discos Duros

– Interface IDE

- Parámetros actuales
 - S M A R T
 - Self-Monitoring, Analysis and Reporting Technology
 - » Compaq
 - » Seagate Technology Inc
 - » IBM
 - » Conner Peripherals Inc , Western Digital Corporation, and Quantum Corporation



Notas:



Tecnologías de Discos Duros


- Interface IDE
 - Parámetros actuales
 - ATAPI
 - ATA Packet Interface
 - » Floppys
 - » Cd Rom
 - » DVD
 - » Soportado por sistemas operativos de red



Notas:

Tecnologías de Discos Duros

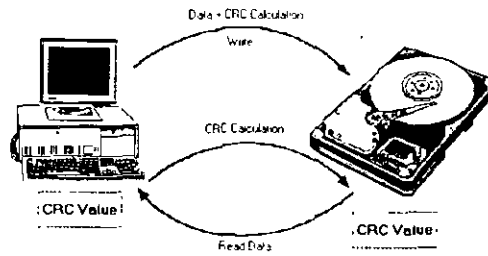
- Interface IDE Ultra ATA
- Maneja 33 mbits/sec
- Uso dual del Reloj




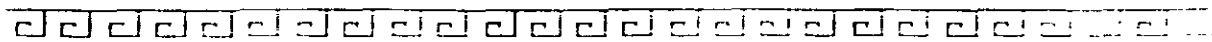
Notas:

Tecnologías de Discos Duros

- Interface IDE Ultra ATA
- Compatible con FAST ATA-2
- Data Reliability



Notas:





Tecnologías de Discos Duros

- Interface IDE Ultra ATA
 - Motherboards or Controller Cards support Ultra ATA:
 - Intel Pentium 430TX PCI set
 - Intel Pentium Pro 440LX AGP set
 - VIA/AMD VIA VPX and VP2/97 (also marketed as AMD-640)
 - SiS 5597/98, 5581/82 and 5601



Notas:



Tecnologías de Discos Duros


- Interface IDE Ultra ATA
- Software
 - Microsoft Windows 95, Windows 98 and Ultra ATA support
 - Microsoft (www.microsoft.com) for the following documents
 - Computer with Ultra DMA IDE Controller May Hang
 - Knowledge Base Article Q171353
 - <http://support.microsoft.com/support/kb/articles/q1713/53.asp>
 - Enabling IDE DMA on Windows-based Systems
 - <http://www.microsoft.com/hwdev/devdev/idedma.htm>



Notas:

Interface IDE Ultra
ATA
T31012A
edalist 4342 Family


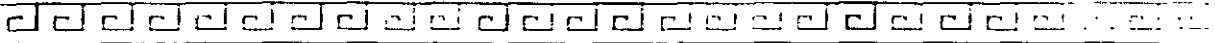
524
64
63
1.08 GE
4500 rpm
12 ms avg
310120



- 1. 1.08 GE
- 2. 4500 rpm
- 3. 12 ms avg
- 4. 310120

- 1. 1.08 GE
- 2. 4500 rpm
- 3. 12 ms avg
- 4. 310120

Notas:




Tecnologías de Discos Duros

- Tipos de Interfaces SCSI
 - Small Computer System Interface
- SCSI (ANSI X3.131-1986)
- SCSI-2 Aprobado en 1994
- SCSI-3 Se trabaja en ello
 - Ultra 2 SCSI-3
 - SCSI- FCS Fibra Chanel Solution


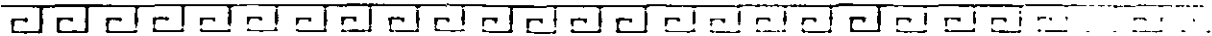


Notas:

Configuración SCSI



Notas:


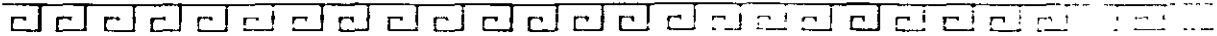


Tecnologías de Discos Duros

- Tipos de Interfaces SCSI
- Ventajas
 - Alta rendimiento
 - ZBR
 - Manejo de Dispositivos Múltiples
 - » 7 SCSI
 - » 15 SCSI-2
 - » 40 SCSI FCS (en un futuro)
 - Capacidades muy Grandes
 - Independiente del Sistema
 - Manejo Dispositivos en Cascada



Notas:


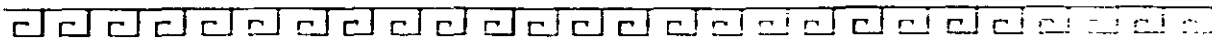


Tecnologías de Discos Duros

- Tipos de Interfaces SCSI
- Variantes SCSI-2
 - Fast SCSI 10 Mbit /sec respecto a 5Mbit/sec SCSI
 - Wide SCSI 16 bits bus respecto a 8 bit de SCSI
- Variantes SCSI-3
 - FAST and WIDE SCSI
 - Fast-20 (Ultra) SCSI
 - Ultra2 SCSI-3 WIDE



Notas:



Tecnologías de Discos Duros

- Tipos de Interfaces SCSI

- Tipos de Cable


- A

- » 50 PINS
- » 8 bits
- » 7 dispositivos

- P

- » 68 PINES
- » 16 y 32 bits
- » 6 m De longitud Maxima

- Q

- » 32 bit para Wide
- 



Notas:

Tecnologías de Discos Duros

- Tipos de Interfaces SCSI

SCSI - 1	SCSI - 2	SCSI - 3	
SCSI Standard	Revision to SCSI-1	Revision to SCSI-2	Ultra2
Asynchronous / Sync 8-bit (narrow) 10 Mbytes/sec	Differential 16-bit (wide) FAST SCSI 20 Mbytes/sec	Ultra 40 Mbytes/sec	Subset of SCSI LVD 80 Mbytes/sec
1980-1985	1986-1992	1993-1998	1995-1998

Notas:




Tecnologías de Discos Duros

SCSI - 1	SCSI - 2	SCSI - 3
...
...
...



Notas:



Chapter 15

Hard Disk Interfaces



This chapter describes the hard disk interface, from the drives to the cables and controllers that run them. You will learn about the various disk interfaces you can select, and the shortcomings and strengths of each type.

Interfaces Choices

A variety of hard disk interfaces are available today. As time has passed, the number of choices has increased, and many of the older designs no longer are viable in newer systems. You need to know about all these interfaces, from the oldest to the newest designs, because you will encounter all of them whenever upgrading or repairing systems is necessary.

The interfaces have different cabling and configuration options, and the setup and format of drives will vary as well. Special problems may arise when you are trying to install more than one drive of a particular interface type or (especially) when you are mixing drives of different interface types in one system.

This section covers the different hard disk drive interfaces, giving you all the technical information you need to deal with them in any way—troubleshooting, servicing, upgrading, and even mixing the different types.

This section also examines the standard controllers and describes how you can work with these controllers, as well as replace them with much faster units. Also discussed are the different types of drive interfaces: ST-506/412, ESDI, IDE, and SCSI. Choosing the proper interface is important, because your choice also affects your disk drive purchase and the ultimate speed of the disk subsystem.

The primary job of the hard disk controller or interface is to transmit and receive data to and from the drive. The different interface types limit how fast data can be moved from the drive to the system and offer different levels of performance. If you are putting together a system in which performance is a primary concern, you need to know how these different interfaces affect performance and what you can expect from them. Many of the statistics that appear in technical literature are not indicative of the real performance figures that you will see in practice. I will separate the myths presented by some of these overoptimistic figures from the reality of what you will actually see.

Chapter 15—Hard Disk Interfaces

With regard to disk drives, and especially hard disk drives, the specification on which people seem to focus the most is the drive's reported *average seek time*, the (average) time it takes for the heads to be positioned from one track to another. Unfortunately, the importance of this specification often is overstated, especially in relation to other specifications, such as the data-transfer rate.

The transfer rate of data between the drive and the system is more important than access time because most drives spend more time reading and writing information than they do simply moving the heads around. The speed at which a program or data file is loaded or read is affected most by the data-transfer rate. Specialized operations such as sorting large files, which involve a lot of random access to individual records of the file (and, therefore, many seek operations), are helped greatly by a faster-seeking disk drive, so seeking performance is important in these cases. Most normal file load and save operations, however, are affected most by the rate at which data can be read and written to and from the drive. The data-transfer rate depends on both the drive and the interface.

Several types of hard disk interfaces have been used in PC systems over the years:

- ST-506/412
- ESDI
- IDE
- SCSI

Of these interfaces, only ST-506/412 and ESDI are what you could call true disk-controller-to-drive interfaces. SCSI and IDE are system-level interfaces that usually incorporate a chipset-based variation of one of the other two types of disk controller interfaces internally. For example, most SCSI and IDE drives incorporate the same basic controller circuitry used in separate ESDI controllers. The SCSI interface adds another layer of interface that attaches the controller to the system bus, whereas IDE is a direct bus-attachment interface.

In data recovery, it helps to know the disk interface you are working with because many data-recovery problems involve drive setup and installation problems. Each interface requires a slightly different method of installation and drive configuration. If the installation or configuration is incorrect or accidentally altered by the system user, it may prevent access to data on a drive. Accordingly, anyone who wants to become proficient in data recovery must be an expert on installing and configuring various types of hard disks and controllers.

IBM's reliance on industry-standard interfaces such as those listed here was a boon for everybody in the IBM-compatible industry. These standards allow a great deal of cross-system and cross-manufacturer compatibility. The use of these industry-standard interfaces allows us to pick up a mail-order catalog, purchase a hard disk for the lowest possible price, and be assured that it will work with our system. This plug-and-play capability results in affordable hard disk storage and a variety of options in capacities and speed.

The ST-506/412 Interface


The ST-506/412 interface was developed by Seagate Technologies around 1980. The interface originally appeared in the Seagate ST-506 drive, which was a 5M formatted (or 6M unformatted) drive in a full-height, 5 1/4-inch form factor. By today's standards, this drive is a tank! In 1981 Seagate introduced the ST-412 drive, which added a feature called buffered seek to the interface. This drive was a 10M formatted (12M unformatted) drive that also qualifies as a tank by today's standards. Besides the Seagate ST-412, IBM also used the Miniscribe 1012 as well as the International Memories, Inc. (IMI) model 5012 drive in the XT. IMI and Miniscribe are long gone, but Seagate remains as one of the largest drive manufacturers today. Since the original XT, Seagate has supplied drives for numerous systems made by many different manufacturers.

Most drive manufacturers that made hard disks for PC systems adopted the Seagate ST-506/412 standard, a situation that helped make this interface popular. One important feature is the interface's plug-and-play design. No custom cables or special modifications are needed for the drives, which means that virtually any ST-506/412 drive will work with any ST-506/412 controller. The only real compatibility issue with this interface is the level of BIOS support provided by the system.

When introduced to the PC industry by IBM in 1983, ROM BIOS support for this hard disk interface was provided by a BIOS chip on the controller. Contrary to what most believed, the PC and XT motherboard BIOS had no inherent hard disk support. When the AT system was introduced, IBM placed the ST-506/412 interface support in the motherboard BIOS and eliminated it from the controller. Since then, any system that is compatible with the IBM AT (which includes most systems on the market today) has an enhanced version of the same support in the motherboard BIOS as well. Because this support was somewhat limited, especially in the older BIOS versions, many disk controller manufacturers also placed additional BIOS support for their controllers directly on the controllers themselves. In some cases, you would use the controller BIOS and motherboard BIOS together; in other cases, you would disable the controller or motherboard BIOS and then use one or the other. These issues will be discussed more completely later in this chapter, in the section "System Configuration."

The ST-506/412 interface does not quite make the grade in today's high-performance PC systems. This interface was designed for a 5M drive, and I have not seen any drives larger than 152M (MFM encoding) or 233M (RLL encoding) available for this type of interface. Because the capacity, performance, and expandability of ST-506/412 are so limited, this interface is obsolete and generally unavailable in new systems. However, many older systems still use drives that have this interface.

Encoding Schemes and Problems. As indicated in Chapter 14 in the section "Data Encoding Schemes," encoding schemes are used in communications for converting digital data bits to various tones for transmission over a telephone line. For disk drives, the digital bits are converted, or encoded, in a pattern of magnetic impulses, or flux transitions (also called flux reversals), which are written on the disk. These flux transitions are decoded later, when the data is read from the disk.



A device called an *Endec* (encoder/decoder) accomplishes the conversion to flux transitions for writing on the media and the subsequent reconversion back to digital data during read operations. The function of the Endec is very similar to that of a modem (modulator/demodulator) in that digital data is converted to an analog waveform, which then is converted back to digital data. Sometimes, the Endec is called a data separator because it is designed to separate data and clocking information from the flux-transition pulse stream read from the disk.

One of the biggest problems with ST-506/412 was the fact that this Endec resided on the disk controller (rather than the drive), which resulted in the possibility of corruption of the analog data signal before it reached the media. This problem became especially pronounced when the ST-506/412 controllers switched to using RLL Endecs to store 50 percent more data on the drive. With the RLL encoding scheme, the actual density of magnetic flux transitions on the disk media remains the same as with MFM encoding, but the timing between the transitions must be measured much more precisely.

In RLL encoding, the intervals between flux changes are approximately the same as with MFM, but the actual timing between them is much more critical. As a result, the transition cells in which signals must be recognized are much smaller and more precisely placed than with MFM. RLL encoding places more stringent demands on the timing of the controller and drive electronics. With RLL encoding, accurately reading the timing of the flux changes is paramount. Additionally, because RLL encodes variable-length groups of bits rather than single bits, a single error in one flux transition can corrupt two to four bits of data. For these reasons, an RLL controller usually has a more sophisticated error-detection and error-correction routine than an MFM controller.

Most of the cheaper disk drives on the market did not have data-channel circuits that were designed to be precise enough to handle RLL encoding without problems. RLL encoding also is much more susceptible to noise in the read signal, and the conventional oxide media coatings did not have a sufficient signal-to-noise ratio for reliable RLL encoding. This problem often was compounded by the fact that many drives of the time used stepper motor head positioning systems, which are notoriously inaccurate, further amplifying the signal-to-noise ratio problem.

At this time, manufacturers starting RLL-certifying drives for use with RLL Endec controllers. This stamp of approval essentially meant that the drive had passed tests and was designed to handle the precise timing requirements that RLL encoding required. In some cases, the drive electronics were upgraded between a manufacturer's MFM and RLL drive versions, but the drives are essentially the same. In fact, if any improvements were made in the so-called RLL-certified drives, the same upgrades usually were also applied to the MFM version.

The bottom line is that other than improved precision, there is no real difference between an ST-506/412 drive that is sold as an MFM model and one that is sold as an RLL model. If you want to use a drive that originally was sold as an MFM model with an RLL controller, I suggest that you do so only if the drive uses a voice coil head actuator and thin-film media. Virtually any ST-506/412 drive with these qualities is more than good enough to handle RLL encoding with no problems.


Using MFM encoding, a standard ST-506/412 format specifies that the drive will contain 17 sectors per track, with each sector containing 512 bytes of data. A controller that uses an RLL Endec raises the number of sectors per track to 25 or 26.

The real solution to reliability problems with RLL encoding was to place the Endec directly on the drive rather than on the controller. This method reduces the susceptibility to noise and interference that can plague an ST-506/412 drive system running RLL encoding. ESDI, IDE, and SCSI drives all have the Endec (and, often, the entire controller) built into the drive by default. Because the Endec is attached to the drive without cables and with an extremely short electrical distance, the propensity for timing- and noise-induced errors is greatly reduced or eliminated. This situation is analogous to a local telephone call between the Endec and the disk platters. This local communication makes the ESDI, IDE, and SCSI interfaces much more reliable than the older ST-506/412 interface; they share none of the reliability problems that once were associated with RLL encoding over the ST-506/412 interface. Virtually all ESDI, IDE, and SCSI drives use RLL encoding today with tremendously increased reliability over even MFM ST-506/412 drives.

ST-506/412 Configuration and Installation. The ST-506/412 interface is characterized by a two- or three-cable arrangement, depending on whether one or two drives are connected. One 34-connector control cable is daisy-chained between up to two drives. The daisy-chain arrangement is much like that used for floppy drives. Each drive on the daisy chain is jumped to respond to a particular Drive Select (DS) line. In the controller implementation used in all PC systems, there are two available lines, called Drive Select 1 (DS1) and Drive Select 2 (DS2). Some drives support as many as four DS lines, but only the first two are usable. Although it may appear that you could string four drives on a single daisy-chain cable, the design of the PC system and controllers uses only the first two.

The control cable usually has lines 25 through 29 twisted between the drive D and C connectors. The first drive (drive C) normally is plugged into the last control cable connector at the end of the cable opposite the controller; an optional second drive (D) can be installed in the middle control-cable connector. The twist in the lines serves to reroute the Drive Select lines so that the drive plugged into the last cable position appears to the controller to be attached to Drive Select 1, even though the jumper on the drive is set for DS2. This arrangement is very similar to the one used for floppy drives; if the cable is twisted, both drives must be set to the DS2 jumper position. If the cable does not have the twisted lines, the drive at the end of the cable (C) must be set to DS1.

Another configuration item is the terminating resistor, which must be installed on the drive at the end of the cable (C) and must be removed from the optional second drive (D) attached to the middle control-cable connector. The controller has a permanently installed terminating resistor that never has to be adjusted. Although the control cable is similar in function and appearance to the 34-pin cable used for floppy drives, the cables generally are not interchangeable because different lines are twisted. Whereas pins 25 through 29 are inverted on the hard disk control cable, pins 10 through 16 are inverted on the floppy cable, rendering them incompatible.



Chapter 15—Hard Disk Interfaces

The other two cables, called data cables, are 20-connector cables, each of which runs from the controller to a single drive because this cable is not daisy-chained. A two-drive system, therefore, has one control cable from the controller to each of two drives in a daisy chain, plus two separate data cables—one for each drive. The controller has three connectors to support the two-drive maximum limit. As its name suggests, the data cable carries data to and from the drive.

If you are using a single drive, only the data cable connector closest to the control cable connector is used; the other is left unattached. Most ST-506/412 controllers also have an on-board floppy controller, which also will have a 34-pin connector for the floppy drives. Figures 15.1 and 15.2 show the control and data cable connectors on a typical combination ST-506/412 hard disk and floppy disk controller. Notice that some of these combination controllers allow the floppy controller portion to be disabled and others do not, which may cause a conflict if you have any other floppy controller in the system.

ST-506/412 Drive Configuration. With ST-506/412 and ESDI drives, you have to configure the following items on each drive:

- Drive Select (DS) jumpers
- Terminating resistor



These configuration items usually are located near the rear of the drive on the disk drive logic board.

Drive Select Jumpers. The Drive Select jumper selects the Drive Select (DS) signal to which the drive should respond. The drive controller sends control signals on two DS lines, one for each drive. Because each drive must be set to respond to a different DS signal, you can use only two drives per controller.

The DS jumpers must be set so that each drive responds to a different DS line from the controller (DS1 or DS2). If the 34-pin control cable has a twist in lines 25 through 29, both drives should be set to DS2. If the control cable is a straight-through design (no twist), the drive at the end of the cable opposite the controller (C) should be set to DS1, and a second drive attached to the middle control-cable connector (D) should be set to DS2. Notice that some drives label the DS jumpers starting with 0, so that DS1 would be labeled DS0 and DS2 would be labeled DS1.

Terminating Resistors. An ST-506/412 drive is always shipped from the factory with a terminating resistor installed. When you install these drives, you must ensure that the drive plugged into the end of the control-cable daisy chain has this terminator installed. Additionally, this terminator must be removed (or disabled with a jumper, in some cases) from the secondary drive installed in the center control-cable connector.

The functions of the terminating resistor are the same as those discussed for floppy drives. The idea is to provide electrical-signal termination so that the control signals to and from the drive and controller do not reflect back or echo along the cable. The terminating resistor provides the proper signal-to-noise ratio and the proper electrical load for the controller. Improper drive termination results in drives that do not function (or that



Control Cable




Figure 15.1

ST-506/412 control cable connector.

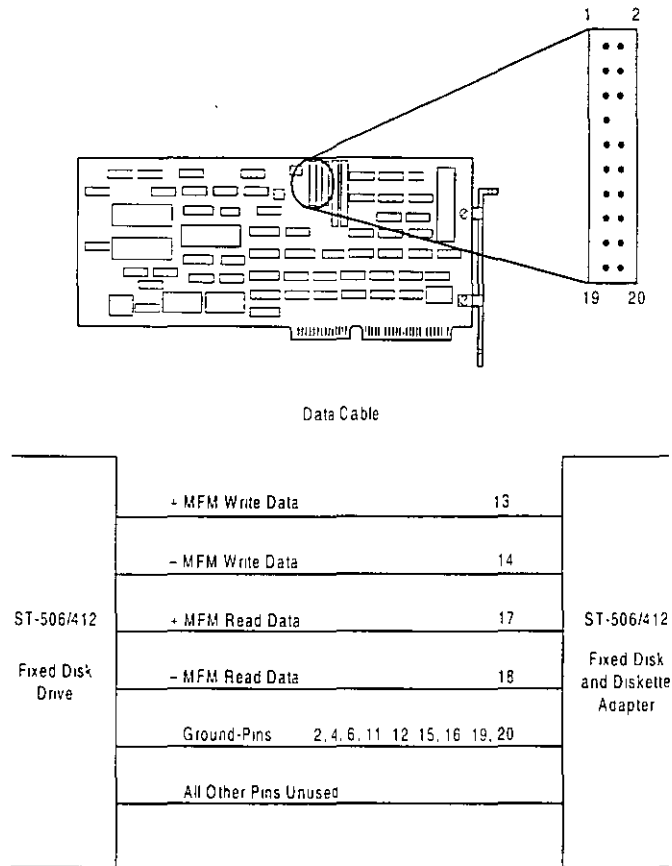


Figure 15.2
ST-506/412 data cable connectors

do so only with excessive problems). Improper termination also may damage the controller because of improper electrical loads

Control and Data Cables. The control cable connects to the controller and daisy chains to the secondary and primary drives, and separate data cables (20-pin) run from the controller to each drive. The data cable connector closest to the control-cable connector on the controller is used for the primary (C) drive.

When connecting cables, you should observe the proper pin-1 orientation from end to end. On ribbon cables, the pin-1 line usually is a different color from the other lines. (In most ribbon cables, for example, the pin-1 line is red or blue, and the rest of the cable is gray). You need to ensure that this pin-1 line is plugged into pin-1 of the controller and drive connectors. Both the controller and drive should have the pin-1 position on each connector marked. Sometimes the mark is the number 1; other times, it is a dot or some other symbol that is silkscreened on the circuit board. The cable connectors at the

controller end may be keyed, in which case pin-15 will be missing from the control-cable connector on the controller connector, and the corresponding hole will be plugged in the control cable. The data cable connectors will be missing pin-8, and the corresponding hole also will be plugged in the data cables. The edge connectors used at the drive end are normally keyed to a notch in the drive connectors. The side of the connector with the notch cut out indicates the pin-1 orientation at the drive end.

Notice that the 34-pin control cable is very similar to the 34-pin control/data cable used for floppy drives; these cables, however, usually are not interchangeable. The ST-506/412 control cable has lines 25 through 29 twisted between the secondary and primary drive connectors, whereas the 34-pin floppy cable has lines 10 through 17 twisted. As a result, the cables are incompatible and therefore noninterchangeable.

ST-506/412 Interface Connectors. The ST-506/412 Interface uses two connections, a 34-pin control connector and a 20-pin data connector. Table 15.1 and 15.2 show the pinouts for these connectors.

Power Cables. To complete the required cable connections to the hard drive, you need a spare power connector (Table 15.3 shows the pinouts for the connector). Some older power supplies have only two-drive power connectors. Several companies sell a power splitter cable, or Y cable, that can adapt one cable from the power supply so that it powers two drives. If you add a power splitter to a system, make sure that the power supply can handle the load of the additional drive or drives.

Signal Name	Pin	Pin	Signal Name
GROUND	1	2	-HD SLCT 3
GROUND	3	4	-HD SLCT 2
GROUND	5	6	-WRITE GATE
GROUND	7	8	-SEEK CMPLT
GROUND	9	10	-TRACK 0
GROUND	11	12	-WRITE FAULT
GROUND	13	14	-HD SLCT 0
KEY (no pin)	15	16	Not Connected
GROUND	17	18	-HD SLCT 1
GROUND	19	20	-INDEX
GROUND	21	22	-READY
GROUND	23	24	-STEP
GROUND	25	26	-DRV SLCT 0
GROUND	27	28	-DRV SLCT 1
GROUND	29	30	Not Connected
GROUND	31	32	Not Connected
GROUND	33	34	-DIRECTION IN

Signal Name	Pin	Pin	Signal Name
-DRV SLCTD	1	2	GROUND
Not Connected	3	4	GROUND
Not Connected	5	6	GROUND
Not Connected	7	8	KEY (no pin)
Not Connected	9	10	Not Connected
GROUND	11	12	GROUND
+MFM WRITE	13	14	-MFM WRITE
GROUND	15	16	GROUND
+MFM READ	17	18	-MFM READ
GROUND	19	20	GROUND

Pin	Wire Color	Signal
1	Yellow	+12v
2	Black	Gnd
3	Black	Gnd
4	Red	+5v

If the original power supply is not adequate, purchase an aftermarket unit that can supply adequate power. Most better aftermarket supplies have four-drive power connectors, eliminating the need for the splitter cables. Power splitter cables are available from several of the cable and accessory vendors listed in Appendix B, as well as from electronic supply stores such as Radio Shack.

Historical Notes. The following sections list some information on the original ST-506/412 controllers used in the PC environment. These were the controllers that IBM supplied in the XT and AT systems. At the time of introduction, these controllers set standards that, especially in the case of the AT controller, we still live with today. In fact, the entire IDE interface standard is based on the controller that IBM designed and used in the AT. All the conventions and standards for the hard disk interfaces that we use today started with these controllers.

Original IBM 8-Bit Controllers. The first ST-506/412 controller standard sold for PC systems was the hard disk controller used in the original 10MB IBM XT. This controller actually was made for IBM by Xebec Corporation and also was sold under the Xebec name as the Xebec 1210 controller. The Xebec 1210 is an ST-506/412 controller that uses Modified Frequency Modulation (MFM) encoding to record data on a drive. This controller's ROM, produced by IBM, contained an 8K hard disk BIOS with an internal table that had entries for four different drives. Each drive was selected by jumpers on the controller, which actually were soldered in the early IBM units. If you purchased the

controller from Xebec, you got a slightly different but completely compatible ROM, and the jumpers were not soldered, so you easily could select one of the four BIOS table entries. Xebec also allowed system integrators to copy its ROM to modify the built-in drive tables for a specific drive.

Later IBM XT systems with a 20M hard disk still used the Xebec 1210, but it had a new 4K ROM that contained different drive tables, as well as jumpers like those found on the versions also sold separately by Xebec. Xebec never sold an autoconfigure version of this controller, which would have made integrating different drives easier.

The Xebec 1210 is one of the slowest ST-506/412 controllers ever made, supporting at best a 5:1 interleave on a stock IBM PC or IBM XT system. If you use the IBM Advanced Diagnostics program for the IBM PC or IBM XT, the low-level formatter produces a standard 6:1 interleave, which results in a paltry 85K-per-second data-transfer rate. By changing the interleave to 5:1, you can wring 102K per second from this controller—still unbelievably slow by today's standards.

Xebec also made a Model 1220 that combined a hard disk and floppy disk controller, was hardware-compatible with the 1210, and worked with the IBM or standard Xebec ROM. The separate floppy controller then could be removed from the system, and you could save a slot.


I recommend replacing this controller with an autoconfigure controller whenever you get the chance. Most other controllers also are significantly faster than the Xebec.

Original IBM 16-Bit Controllers. For the AT, IBM used two controllers made by Western Digital: the WD1002-WA2 and the WD1003A-WA2. The WD1003 is an upgraded WD1002 with a much lower chip count. The WD1003 also was shorter than the WD1002 to fit into the IBM XT 286.

The WD1002 is used in the IBM AT as a combination hard disk and floppy disk controller. The WD1002 and the WD1003 are standard ST-506/412 controllers that supply MFM encoding to the drive. Neither controller contains a ROM BIOS, instead, BIOS support is built into the motherboard ROM. Both controllers support a 2:1 interleave, even on a standard 6-MHz IBM AT system. The IBM Advanced Diagnostics low-level formatter can put down a 2:1 interleave, but the default is 3:1. Most users of these controllers can realize a performance gain if they simply reformat to the lower interleave.

The ESDI Interface

ESDI, or Enhanced Small Device Interface, is a specialized hard disk interface established as a standard in 1983, primarily by Maxtor Corporation. Maxtor led a consortium of drive manufacturers to adopt its proposed interface as a high-performance standard to succeed ST-506/412. ESDI later was adopted by the ANSI (American National Standards Institute) organization and published under the ANSI X3T9.2 Committee. The latest version of the ANSI ESDI document is known as X3.170a-1991. You can obtain this document, and other ANSI-standard documents, from ANSI itself or from Global Engineering Documents. These companies are listed in Appendix B.



Chapter 15—Hard Disk Interfaces

Compared with ST-506/412, ESDI has provisions for increased reliability, such as building the Endec (encoder/decoder) into the drive. ESDI is a very-high-speed interface, capable of a maximum 24-megabits-per-second transfer rate. Most drives running ESDI, however, are limited to a maximum 10 or 15 megabits per second. Unfortunately, compatibility problems between different ESDI implementations combined with pressure from low-cost, high-performance IDE interface drives have served to make the ESDI interface obsolete. Few, if any, new systems today include ESDI drives, although ESDI became somewhat popular in high-end systems during the late '80s.

Enhanced commands enabled some ESDI controllers to read a drive's capacity parameters directly from the drive, as well as to control defect mapping, but several manufacturers had different methods for writing this information on the drive. When you install an ESDI drive, in some cases the controller automatically reads the parameter and defect information directly from the drive. In other cases, however, you still have to enter this information manually, as with ST-506/412.

The ESDI's enhanced defect-mapping commands provide a standard way for the PC system to read a defect map from a drive, which means that the manufacturer's defect list can be written to the drive as a file. The defect-list file then can be read by the controller and low-level format software, eliminating the need for the installer to type these entries from the keyboard and enabling the format program to update the defect list with new entries if it finds new defects during the low-level format or the surface analysis.

Most ESDI implementations have drives formatted to 32 sectors per track or more (80 or more sectors per track are possible)—many more sectors per track than the standard ST-506/412 implementation of 17 to 26 sectors per track. The greater density results in two or more times the data-transfer rate, with a 1:1 interleave. Almost without exception, ESDI controllers support a 1:1 interleave, which allows for a transfer rate of 1MB per second or greater.

Because ESDI is much like the ST-506/412 interface, it can replace that interface without affecting software in the system. Most ESDI controllers are register-compatible with the older ST-506/412 controllers, which enables OS/2 and other non-DOS operating systems to run with few or no problems. The ROM BIOS interface to ESDI is similar to the ST-506/412 standard, and many low-level disk utilities that run on one interface will run on the other. To take advantage of ESDI defect mapping and other special features, however, use a low-level format and surface-analysis utility designed for ESDI (such as the ones usually built into the controller ROM BIOS and called by DEBUG).

During the late 1980s, most high-end systems from major manufacturers were equipped with ESDI controllers and drives. More recently, manufacturers have been equipping high end systems with SCSI. The SCSI interface allows for much greater expandability, supports more types of devices than ESDI, and offers equal or greater performance. I no longer recommend installing ESDI drives unless you are upgrading a system that already has an ESDI controller.

ESDI Interface Connectors. ESDI (Enhanced Small Device Interface) uses two connections, a 34-pin control connector and a 20-pin data connector. Tables 15.4 and 15.5 show the pinouts for these connectors.



Signal Name	Pin	Pin	Signal Name
GROUND	1	2	-HD SLCT 3
GROUND	3	4	-HD SLCT 2
GROUND	5	6	-WRITE GATE
GROUND	7	8	-CNFG/STATUS
GROUND	9	10	-XFER ACK
GROUND	11	12	-ATTENTION
GROUND	13	14	-HD SLCT 0
KEY (no pin)	15	16	-SECTOR
GROUND	17	18	-HD SLCT 1
GROUND	19	20	-INDEX
GROUND	21	22	-READY
GROUND	23	24	-XFER REQ
GROUND	25	26	-DRV SLCT 0
GROUND	27	28	-DRV SLCT 1
GROUND	29	30	Reserved
GROUND	31	32	-READ GATE
GROUND	33	34	-CMD DATA

Signal Name	Pin	Pin	Signal Name
-DRV SLCTD	1	2	-SECTOR
-CMD COMPL	3	4	-ADDR MK EN
GROUND	5	6	GROUND
+WRITE CLK	7	8	-WRITE CLK
GROUND	9	10	+RD/REF CLK
-RD/REF CLK	11	12	GROUND
+NRZ WRITE	13	14	-NRZ WRITE
GROUND	15	16	GROUND
+NRZ READ	17	18	-NRZ READ
GROUND	19	20	-INDEX

ESDI Drive Configuration. The ESDI interface was modeled after the ST-506/412 interface and shares virtually all the same types of configuration items and procedures. The 34-pin control and 20-pin data cables are identical to those used in an ST-506/412 installation, and all the configuration procedures with regard to Drive Select jumpers, twisted cables, and terminating resistors are the same as with ST-506/412.

Follow the configuration procedures for ST-506/412 drives when configuring ESDI drives. Because ESDI was developed from the ST-506/412 interface, it shares many characteristics including the type of cabling used as well as how it is configured.

The IDE Interface

Integrated Drive Electronics (IDE) is a generic term applied to any drive with an integrated (built-in) disk controller. The IDE interface as we know it is officially called ATA (AT Attachment) and is an ANSI standard, however IDE can roughly apply to any disk drive with a built-in controller. The first drives with integrated controllers were Hardcards; today, a variety of drives with integrated controllers are available. In a drive with IDE, the disk controller is integrated into the drive, and this combination drive/controller assembly usually plugs into a bus connector on the motherboard or on a bus adapter card. Combining the drive and controller greatly simplifies installation because there are no separate power or signal cables from the controller to the drive. Also, when the controller and the drive are assembled as a unit, the number of total components is reduced, signal paths are shorter, and the electrical connections are more noise-resistant, resulting in a more reliable design than is possible when a separate controller, connected to the drive by cables, is used.

Placing the controller (including Endec) on the drive gives IDE drives an inherent reliability advantage over interfaces with separate controllers. Reliability is increased because the data encoding, from digital to analog, is performed directly on the drive in a tight noise-free environment; the timing-sensitive analog information does not have to travel along crude ribbon cables that are likely to pick up noise and to insert propagation delays into the signals. The integrated configuration allows for increases in the clock rate of the encoder, as well as the storage density of the drive.

Integrating the controller and drive also frees the controller and drive engineers from having to adhere to the strict standards imposed by the earlier interface standards. Engineers can design what essentially are custom drive and controller implementations because no other controller would ever have to be connected to the drive. The resulting drive and controller combinations can offer higher performance than earlier stand-alone controller and drive setups. IDE drives sometimes are called drives with embedded controllers.

The IDE connector on motherboards in many systems is nothing more than a stripped-down bus slot. In ATA IDE installations, these connectors normally contain a 40-pin subset of the 98 pins that would be available in a standard 16-bit ISA bus slot. The pins used are only the signal pins required by a standard-type XT or AT hard disk controller. For example, because an AT-style disk controller uses only interrupt line 14, the motherboard AT IDE connector supplies only that interrupt line; no other interrupt lines are needed. The XT IDE motherboard connector supplies interrupt line 5 because that is what an XT controller would use.

Many people who use systems with IDE connectors on the motherboard believe that a hard disk controller is built into their motherboard, but the controller really is in the drive. I do not know of any PC systems that have hard disk controllers built into the motherboard.

When IDE drives are discussed, the ATA IDE variety usually is the only kind mentioned because it is so popular. But other forms of IDE drives exist, based on other buses. For example, several PS/2 systems came with Micro Channel (MCA) IDE drives, which plug

directly into a Micro Channel Bus slot (through an angle adapter or Interposer card). An 8-bit ISA form of IDE also existed but was never very popular. Most IBM-compatible systems with the ISA or EISA Bus use AT-Bus (16-bit) IDE drives. The ATA IDE interface is by far the most popular type of drive interface available

The primary advantage of IDE drives is cost. Because the separate controller or host adapter is eliminated and the cable connections are simplified, IDE drives cost much less than a standard controller-and-drive combination. These drives also are more reliable because the controller is built into the drive. Therefore, the Endec or data separator (the converter between the digital and analog signals on the drive) stays close to the media. Because the drive has a short analog-signal path, it is less susceptible to external noise and interference.

Another advantage is performance. IDE drives are some of the highest-performance drives available—but they are also among the lowest-performance drives. This apparent contradiction is a result of the fact that all IDE drives are different. You cannot make a blanket statement about the performance of IDE drives because each drive is unique. The high-end models, however, offer performance that is equal or superior to that of any other type of drive on the market for a single-user, single-tasking operating system

IDE Origins. Technically, the first IDE drives were Hardcards. Companies such as the Plus Development division of Quantum took small 3.5-inch drives (either ST-506/412 or ESDI) and attached them directly to a standard controller. The assembly then was plugged into a bus slot as though it were a normal disk controller. Unfortunately, the mounting of a heavy, vibrating hard disk in an expansion slot with nothing but a single screw to hold it in place left a lot to be desired—not to mention the possible interference with adjacent cards due to the fact that many of these units were much thicker than a controller card alone.

Several companies got the idea that you could redesign the controller to replace the logic-board assembly on a standard hard disk and then mount it in a standard drive bay just like any other drive. Because the built-in controller in these drives still needed to plug directly into the expansion bus just like any other controller, a cable was run between the drive and one of the slots

These connection problems were solved in different ways. Compaq was the first to incorporate a special bus adapter in its system to adapt the 98-pin AT bus edge connector on the motherboard to a smaller 40-pin header style connector that the drive would plug into. The 40-pin connectors were all that were needed because it was known that a disk controller would never need more than 40 of the bus lines.

In 1987, IBM developed its own MCA IDE drives and connected them to the bus through a bus adapter device called an interposer card. These bus adapters (sometimes called paddle boards) needed only a few buffer chips and did not require any real circuitry because the drive-based controller was designed to plug directly into the bus. The paddle board nickname came from the fact that they resembled game paddle or joystick adapters, which do not have much circuitry on them. Another 8-bit variation of IDE appeared in 8-bit ISA systems such as the PS/2 Model 30. The XT IDE interface uses a 40-pin connector and cable that is similar to, but not compatible with, the 16-bit version.



IDE Bus Versions. Three main types of IDE interfaces are available, with the differences based on three different bus standards:

- AT Attachment (ATA) IDE (16-bit ISA)
- XT IDE (8-bit ISA)
- MCA IDE (16-bit Micro Channel)

The XT and ATA versions have standardized on 40-pin connectors and cables, but the connectors have slightly different pinouts, rendering them incompatible with one another. MCA IDE uses a completely different 72-pin connector and is designed for MCA bus systems only.

In most cases, you must use the type of IDE drive that matches your system bus. This situation means that XT IDE drives work only in XT-class 8-bit ISA slot systems, AT IDE drives work only in AT-class 16-bit ISA or EISA slot systems, and MCA IDE drives work only in Micro Channel systems (such as the IBM PS/2 Model 50 or higher). A company called Silicon Valley offers adapter cards for XT systems that will run ATA IDE drives. Other companies, such as Arco Electronics and Sigma Data, have IDE adapters for Micro Channel systems that allow ATA IDE drives to be used on these systems. (You can find these vendors in Appendix B.) These adapters are very useful for XT or PS/2 systems, because there is a very limited selection of XT or MCA IDE drives, whereas the selection of ATA drives is virtually unlimited.

In most modern ISA and EISA systems, you will find an ATA connector on the motherboard. If your motherboard does not have one of these connectors and you want to attach an AT IDE drive to your system, you can purchase an adapter card that changes your 98-pin slot connector to the 40-pin IDE connector. These adapter cards are nothing more than buffered cables; they are not really controllers. The controller is built into the drive. Some of the cards offer additional features, such as an on-board ROM BIOS or cache memory.

ATA IDE. CDC, Western Digital, and Compaq actually created what could be called the first ATA type IDE interface drive and were the first to establish the 40-pin IDE connector pinout. The first ATA IDE drives were 5-1/4-inch half-height CDC 40M units (I believe that they had a green activity LED) with integrated WD controllers sold in the first Compaq 386 systems way back in '86. After that, Compaq helped found a company called Conner Peripherals to supply Compaq with IDE drives. Conner originally made drives only for Compaq, but later, Compaq sold much of its ownership of Conner.

Eventually, the 40-pin IDE connector and drive interface method was placed before one of the ANSI standards committees, which, in conjunction with drive manufacturers, ironed out some deficiencies, tied up some loose ends, and published what is known as the CAM ATA (Common Access Method/AT Attachment) interface. The CAM Committee was formed in October 1988, and the first working document of the AT Attachment interface was introduced in March 1989. Before the CAM ATA standard, many companies that followed CDC, such as Conner Peripherals, made proprietary changes to what

had been done by CDC. As a result, many older ATA drives are very difficult to integrate into a dual-drive setup that has newer drives.

Some areas of the ATA standard have been left open for vendor-specific commands and functions. These vendor-specific commands and functions are the main reason why it is so difficult to low-level format IDE drives. To work properly, the formatter that you are using usually must know the specific vendor-unique commands for rewriting sector headers and remapping defects. Unfortunately, these and other specific drive commands differ from OEM to OEM, clouding the "standard" somewhat.

It is important to note that only the ATA IDE interface has been standardized by the industry. The XT IDE and MCA IDE were never adopted as industry-wide standards and never became very popular. These interfaces are no longer in production, and no new systems of which I am aware come with these nonstandard IDE interfaces.

The ATA Specification. The ATA specification was introduced in March 1989 as an ANSI standard. ATA-1 was finally approved in 1994, and ATA-2 (also called Enhanced IDE) was approved in 1995. ATA-3 is currently in the works. You can obtain the current version of these standards from Global Engineering Documents, which is listed in Appendix B. The ATA standards have gone a long way toward eliminating incompatibilities and problems with interfacing IDE drives to ISA and EISA systems. The ATA specifications define the signals on the 40-pin connector, the functions and timings of these signals, cable specifications, and so on. The following section lists some of the elements and functions defined by the ATA specification.

Dual-Drive Configurations. Dual-drive ATA installations can be problematic because each drive has its own controller and both controllers must function while being connected to the same bus. There has to be a way to ensure that only one of the two controllers will respond to a command at a time.

The ATA standard provides the option of operating on the AT Bus with two drives in a daisy-chained configuration. The primary drive (drive 0) is called the master, and the secondary drive (drive 1) is the slave. You designate a drive as being master or slave by setting a jumper or switch on the drive or by using a special line in the interface called the Cable Select (CSEL) pin.

When only one drive is installed, the controller responds to all commands from the system. When two drives (and, therefore, two controllers) are installed, all commands from the system are received by both controllers. Each controller then must be set up to respond only to commands for itself. In this situation, one controller then must be designated as the master and the other as the slave. When the system sends a command for a specific drive, the controller on the other drive must remain silent while the selected controller and drive are functioning. You handle discrimination between the two controllers by setting a special bit (the DRV bit) in the Drive/Head Register of a command block.

ATA I/O Connector. The ATA interface connector is a 40-pin header-type connector that should be keyed to prevent the possibility of installing it upside down. A key is



Chapter 15—Hard Disk Interfaces

provided by the removal of pin 20, and the corresponding pin on the cable connector should be plugged to prevent a backward installation. The use of keyed connectors and cables is highly recommended because plugging an IDE cable in backward can damage both the drive and the bus adapter circuits (although I have done it myself many times with no smoked parts yet!).


Table 15.6 shows the ATA-IDE interface connector pinout.

Signal Name	Pin	Pin	Signal Name
-RESET	1	2	GROUND
Data Bit 7	3	4	Data Bit 8
Data Bit 6	5	6	Data Bit 9
Data Bit 5	7	8	Data Bit 10
Data Bit 4	9	10	Data Bit 11
Data Bit 3	11	12	Data Bit 12
Data Bit 2	13	14	Data Bit 13
Data Bit 1	15	16	Data Bit 14
Data Bit 0	17	18	Data Bit 15
GROUND	19	20	KEY (pin missing)
DRQ 3	21	22	GROUND
-IOW	23	24	GROUND
-IOR	25	26	GROUND
I/O CH RDY	27	28	SPSYNC CSEL
-DACK 3	29	30	GROUND
IRQ 14	31	32	-IOCS16
Address Bit 1	33	34	-PDIAG
Address Bit 0	35	36	Address Bit 2
-CS1FX	37	38	-CS3FX
-DA/SP	39	40	GROUND
+5 Vdc (Logic)	41	42	-5 Vdc (Motor)
GROUND	43	44	-TYPE (0=ATA)

ATA I/O Cable. A 40-conductor ribbon cable is specified to carry signals between the bus adapter circuits and the drive (controller). To maximize signal integrity and to eliminate potential timing and noise problems, the cable should not be longer than 0.46 meters (18 inches).

ATA Signals. The ATA interface signals and connector pinout are listed in Appendix A. This section describes some of the most important signals in more detail.

Pin 20 is used as a key pin for cable orientation and is not connected through in the interface. This pin should be missing from any ATA connectors, and the cable should have the pin-20 hole in the connector plugged off to prevent the cable from being plugged in backward.



Pin 39 carries the Drive Active/Slave Present (DASP) signal, which is a dual-purpose, time-multiplexed signal. During power-on initialization, this signal indicates whether a slave drive is present on the interface. After that, each drive asserts the signal to indicate that it is active. Early drives could not multiplex these functions and required special jumper settings to work with other drives. Standardizing this function to allow for compatible dual-drive installations is one of the features of the ATA standard.

Pin 28 carries the Cable Select or Spindle Synchronization signal (CSEL or SPSYNC), which is a dual-purpose conductor; a given installation, however, may use only one of the two functions. The CSEL (Cable Select) function is the most widely used and is designed to control the designation of a drive as master (drive 0) or slave (drive 1) without requiring jumper settings on the drives. If a drive sees the CSEL as being grounded, the drive is a master; if CSEL is open, the drive is a slave.

You can install special cabling to ground CSEL selectively. This installation is normally accomplished through a Y-cable arrangement, with the IDE bus connector in the middle and each drive at opposite ends of the cable. One leg of the Y has the CSEL line connected through, indicating a master drive; the other leg has the CSEL line open (conductor interrupted or removed), making the drive at that end the slave.

ATA Commands. One of the best features of the ATA IDE interface is the enhanced command set. The ATA IDE interface was modeled after the WD1003 controller that IBM used in the original AT system. All ATA IDE drives must support the original WD command set (eight commands), with no exceptions, which is why IDE drives are so easy to install in systems today. All IBM-compatible systems have built-in ROM BIOS support for the WD1003, which means that essentially, they support ATA IDE as well.

In addition to supporting all the WD1003 commands, the ATA specification added numerous other commands to enhance performance and capabilities. These commands are an optional part of the ATA interface, but several of them are used in most drives available today and are very important to the performance and use of ATA drives in general.

Perhaps the most important is the Identify Drive command. This command causes the drive to transmit a 512-byte block of data that provides all details about the drive. Through this command, any program (including the system BIOS) can find out exactly what type of drive is connected, including the drive manufacturer, model number, operating parameters, and even the serial number of the drive. Many modern BIOSes use this information to automatically receive and enter the drive's parameters into CMOS memory, eliminating the need for the user to enter these parameters manually during system configuration. This arrangement helps prevent mistakes that later can lead to data loss when the user no longer remembers what parameters he or she used during setup.

The Identify Drive data can tell you many things about your drive, including the following:

- Number of cylinders in the recommended (default) translation mode
- Number of heads in the recommended (default) translation mode

- Number of sectors per track in the recommended (default) translation mode
- Number of cylinders in the current translation mode
- Number of heads in the current translation mode
- Number of sectors per track in the current translation mode
- Manufacturer and model number
- Firmware revision
- Serial number
- Buffer type, indicating sector buffering or caching capabilities

Several public-domain programs can execute this command to the drive and report the information on-screen. I use the IDEINFO program (which can be downloaded from the IBM Hardware Forum on CompuServe) or the IDEDIAG utility (which can be downloaded from the Western Digital BBS). Phone numbers for these information services appear in Appendix B. I find these programs especially useful when I am trying to install IDE drives and need to know the correct parameters for a user-definable BIOS type. These programs get the information directly from the drive itself.

Two other very important commands are the Read Multiple and Write Multiple commands. These commands permit multiple-sector data transfers and, when combined with block-mode Programmed I/O (PIO) capabilities in the system, can result in incredible data-transfer rates many times faster than single-sector PIO transfers.



If you want the ultimate in IDE performance and installation ease, make sure that your motherboard BIOS and IDE adapter supports ATA-2 or Enhanced IDE (EIDE). This support allows your BIOS to execute data transfers to and from the IDE drive several times faster than normal, and also makes installation and configuration easier because the BIOS will be able to detect the drive-parameter information automatically. High speed PIO (Programmed Input Output) and automatic detection of the drive type are included in the latest versions of most PC BIOSes.

There are many other enhanced commands, including room for a given drive manufacturer to implement what are called vendor-unique commands. These commands often are used by a particular vendor for features unique to that vendor. Often, features such as low-level formatting and defect management are controlled by vendor-unique commands. This is why low-level format programs can be so specific to a particular manufacturer's IDE drives and why many manufacturers make their own LLF programs available.

ATA IDE Drive Categories. ATA-IDE drives can be divided into three main categories. These categories separate the drives by function (such as translation capabilities) and design (which can affect features such as low-level formatting). The three drive categories are:

- Non-Intelligent ATA-IDE drives
- Intelligent ATA-IDE drives
- Intelligent Zoned Recording ATA-IDE drives

The following sections describe these categories.

Non-Intelligent IDE. As I stated earlier, the ATA standard requires that the built-in controller respond exactly as though it were a Western Digital WD1003 controller. This controller responds to a command set of eight commands. Early IDE drives supported these commands and had few, if any, other options. These early drives actually were more like regular ST-506/412 or ESDI controllers bolted directly to the drive than the more intelligent drives that we consider today to be IDE.

These drives were not considered to be intelligent IDE drives; an intelligent drive is supposed to have several capabilities that these early IDE drives lacked. The drives could not respond to any of the enhanced commands that were specified as (an optional) part of the ATA IDE specification, including the Identify Drive command. These drives also did not support sector translation, in which the physical parameters could be altered to appear as any set of logical cylinders, heads, and sectors. Enhanced commands and sector-translation support are what make an IDE drive an intelligent IDE drive, and these features were not available in the early IDE drives.

These drives could be low-level formatted in the same manner as any normal ST-506/412 or ESDI drive. They were universally low-level formatted at the factory, with factory-calculated optimum interleave (usually, 1:1) and head- and cylinder-skew factors. Also, factory defects were recorded in a special area on the drive; they no longer were written on a sticker pasted to the exterior. Unfortunately, this arrangement means that if you low-level format these drives in the field, you most likely will alter these settings (especially the skew factors) from what the factory set as optimum, as well as wipe out the factory-written defect table.

Some manufacturers released special low-level format routines that would reformat the drives while preserving these settings, but others did not make such programs available. Because they did not want you to overwrite the defect list or potentially slow the drive, most manufacturers stated that you should never low-level format their IDE drives.

This statement started a myth that the drives could somehow be damaged or rendered inoperable by such a format, which truly is not the case. One rumor was that the servo information could be overwritten, which would mean that you would have to send the drive back to the manufacturer for re-servoing. This also is not true; the servo information is protected and cannot be overwritten. The only consequence of an improper low-level format of these drives is the possible alteration of the skew factors and the potential loss of the factory defect maps.

The Disk Manager program by Ontrack is the best special-purpose format utility to use on these drives for formatting because it is aware of these types of drives and often can restore the skew factors and preserve the defect information. If you are working with a



drive that already has had the defect map overwritten, Disk Manager can perform a very good surface analysis that will mark off any of these areas that it finds. Disk Manager allows you to specify the skew factors and to mark defects at the sector level so that they will not cause problems later. Other general-purpose diagnostics that work especially well with IDE drives such as this include the Microscope program by Micro 2000.

Intelligent IDE. Later IDE drives became known as intelligent IDE drives. These drives support enhanced ATA commands, such as the Identify Drive command, and sector-translation capabilities.

These drives can be configured in two ways: in raw physical mode or in translation mode. To configure the drive in raw physical mode, you simply enter the CMOS drive parameters during setup so that they match the true physical parameters of the drive. For example, if the drive physically has 800 cylinders, 6 heads, and 50 sectors per track, you enter these figures during setup. To configure the drive in translation mode, you simply enter any combination of cylinders, heads, and sectors that adds up to equal or less than the true number of sectors on the drive.

In the example I just used, the drive has a total of 240,000 sectors ($800 \times 6 \times 50$). All I have to do is figure out another set of parameters that adds up to equal or less than 240,000 sectors. The simplest way to do this is to cut the number of cylinders in half and double the number of heads. Thus, the new drive parameters become 400 cylinders, 12 heads, and 50 sectors per track. This method adds up to 240,000 sectors and enables the drive to work in translation mode.

When these drives are in translation mode, a low-level format cannot alter the interleave and skew factors; neither can it overwrite the factory defect-mapping information. A low-level format program can, however, perform additional defect mapping or sector sparing while in this mode.

If the drive is in true physical mode, a low-level format rewrites the sector headers and modifies the head and cylinder skewing. If done incorrectly, the format can be repaired by a proper low-level format program that allows you to set the correct head and cylinder skew. This task can be accomplished automatically by the drive manufacturer's recommended low-level format program (if available) or by other programs, such as Disk Manager by Ontrack. When you use Disk Manager, you have to enter the skew values manually; otherwise, the program uses predetermined defaults. To get the correct skew values, it is best to contact the drive manufacturer's technical-support department. You can calculate the skew values if the manufacturer cannot provide them.

To protect the skew factors and defect information on intelligent IDE drives, all you have to do is run them in translation mode. In translation mode, this information cannot be overwritten.

Intelligent Zoned Recording IDE. The last and most sophisticated IDE drives combine intelligence with Zoned Recording. With Zoned Recording, the drive has a variable number sectors per track in several zones across the surface of the drive. Because the PC BIOS can handle only a fixed number of sectors on all tracks, these drives must always run in

translation mode. Because these drives are always in translation mode, you cannot alter the factory-set interleave and skew factors or wipe out the factory defect information.

You still can low-level format these drives, however, and use such a format to map or spare additional defective sectors that crop up during the life of the drive. To low-level format intelligent Zoned Recording drives, you need either a specific utility from the drive manufacturer or an IDE-aware program, such Disk Manager by Ontrack or Microscope by Micro 2000.

ATA-2 (Enhanced IDE). ATA-2 is an extension of the original ATA (IDE) specification. The most important additions are performance enhancing features such as fast PIO and DMA modes. ATA-2 also features improvements in the Identify Drive command allowing a drive to tell the software exactly what its characteristics are; this is essential for both Plug-n-Play and compatibility with future revisions of the standard

ATA-2 is often called Enhanced IDE (EIDE). EIDE is technically a marketing program from Western Digital. Fast-ATA and Fast-ATA-2 are similar Seagate-inspired marketing programs, which are also endorsed by Quantum. As far as the hard disk and BIOS are concerned, these are all different terms for basically the same thing.

There are four main areas where ATA-2 and EIDE have improved the original ATA/IDE interface:

- Increased Maximum Drive Capacity
- Faster Data Transfer
- Secondary Two-Device Channel
- ATAPI (ATA Program Interface)

The following section describes these improvements.

Increased Drive Capacity. ATA-2/EIDE allows for increased drive capacity over the original ATA/IDE specification. This is done through an Enhanced BIOS, which makes it possible to use hard disks exceeding the 504M (528 million bytes) barrier. The origin of this limit is the disk geometry (cylinders, heads, sectors) supported by the combination of an IDE drive and the BIOS' software interface. Both IDE and the BIOS are capable of supporting huge disks, but their combined limitations conspire to restrict the useful capacity to 504MB.

An Enhanced BIOS circumvents this by using a different geometry when talking to the drive than when talking to the software. What happens in between is called *translation*. For example, if your drive has 2,000 cylinders and 16 heads, a translating BIOS will make programs think that the drive has 1,000 cylinders and 32 heads.

You can usually tell if your BIOS is enhanced by the ability to specify more than 1,024 cylinders in the BIOS setup, although this is not conclusive. If you see drive-related settings like "LBA," "ECHS," or even "Large," these are tell-tale signs of a BIOS with translation support. Most BIOSes with a date of 1994 or later are enhanced. If your system currently does not have an Enhanced BIOS you may be able to get an upgrade.



BIOS Translation. There are roughly three ways today's BIOSes can handle translation: standard CHS addressing, Extended CHS addressing, and LBA addressing. They are summarized in the following table:

BIOS Mode	Operating System to BIOS	BIOS to Drive Ports
Standard CHS	Logical CHS Parameters	Logical CHS Parameters
Extended CHS	Translated CHS Parameters	Logical CHS Parameters
LBA	Translated CHS Parameters	LBA Parameters

Standard CHS (Cylinder Head Sector). In Standard CHS, there is only one possible translation step internal to the drive. The drive's actual, physical geometry is completely invisible from the outside with all zoned recorded ATA drives today. The Cylinders, Heads, and Sectors printed on the label for use in the BIOS setup are purely logical geometry, and do not represent the actual physical parameters. Standard CHS addressing is limited to 16 heads and 1,024 cylinders, which gives us a limit of 504 Megabytes (528 Million Bytes).

This is often called "Normal" in the BIOS setup, and causes the BIOS to behave like an old-fashioned one without translation. Use this setting if your drive has fewer than 1,024 cylinders or if you want to use the drive with a non-DOS operating system that doesn't understand translation.

Extended CHS (Cylinder Head Sector). In Extended CHS, a translated logical geometry is used to communicate between the drive and the BIOS, while a different, translated geometry is used to communicate between the BIOS and everything else. In other words, there are normally two translation steps. The drive still translates internally, but has logical parameters that exceed the 1,024 cylinder limitation of the standard BIOS. In this case, the drive's cylinder count is usually divided by two, and the head count is multiplied by two to get the translated values from those actually stored in the CMOS SETUP. This type of setting breaks the 504/528M barrier.

This is often called "Large" or "ECHS" in the BIOS setup, and tells the BIOS to use Extended CHS translation. It uses a different geometry (cylinders/heads/sectors) when talking to the drive than when talking to the BIOS. This type of translation should be used with drives that have more than 1,024 cylinders but that do not support LBA (Logical Block Addressing). Note that the geometry entered in your BIOS setup is the logical geometry, not the translated one.

LBA (Logical Block Addressing) LBA is a means of linearly addressing sector's addresses, beginning at Cylinder 0, Head 0, Sector 1 as LBA 0, and proceeding on to the last physical sector on the drive. This is new in ATA-2, but has always been the one and only addressing mode in SCSI.

With LBA, each sector on the drive is numbered starting from 0. The number is a 28-bit binary number internally, which translates to a sector number of from 0 to 268,435,456. Since each sector represents 512 bytes, this results in a maximum drive capacity of exactly 128G, or 137 billion bytes. Unfortunately, the operating system still needs to see a

translated CHS, so the BIOS determines how many sectors there are, and comes up with Translated CHS to match. The BIOS CHS limits are 1,024 cylinders, 256 Heads, and 63 Sectors per track which limits total drive capacity to just under 8GB.

In other words, this breaks the 528MB barrier in essentially the same way as Extended CHS does. Because it is somewhat simpler to use a single linear number to address a sector on the hard disk compared to a CHS type address, this is the preferred method if the drive supports LBA.

A word of warning with these BIOS translation settings. If you switch between standard CHS, Extended CHS, or LBA, the BIOS may change the (translated) geometry. The same thing may happen if you transfer a disk that has been formatted on an old, non-LBA computer to a new one that uses LBA. This will cause the logical CHS geometry seen by the operating system to change, and will cause the data to appear in the wrong locations! This can cause you to lose access to your data if you are not careful. I always recommend recording the CMOS setup screens associated with the hard disk configuration so that you can properly match the setup of a drive to the settings it was originally at.


Faster Data Transfer. ATA-2/EIDE defines several high performance modes for transferring data to and from the drive. These faster modes are the main part of the new specifications and were the main reason they were initially developed. Most of the faster drives on the market today will support either Programmed I/O transfer Mode 3 or Mode 4, which results in a very fast transfer. The following section discusses these modes.

PIO (Programmed I/O Modes). The PIO mode determines how fast data is transferred to and from the drive. In the slowest possible mode, PIO mode 0, the data cycle time cannot exceed 600 nanoseconds. In a single cycle, 16 bits are transferred in or out of the drive making the theoretical transfer rate of PIO Mode 0 (600ns cycle time) 3.3 megabytes per second. Most of the high performance ATA-2 (EIDE) drives today support PIO Mode 4, which offers a 16.6MB per second transfer rate.

The following table shows the PIO modes, with their respective transfer rates:

PIO Mode	Cycle time (ns)	Transfer rate (MB/s)	Specification
0	600	3.3	ATA
1	383	5.2	ATA
2	240	8.3	ATA
3	180	11.1	ATA-2
4	120	16.6	ATA-2

To run in Mode 3 or 4 requires that the IDE port on the system be a local bus port. This means that it must operate through either a VL-Bus or PCI bus connection. Most modern motherboards with ATA-2/EIDE support have dual IDE connectors on the motherboard, with only the primary connector running through the system's PCI local bus. The secondary connector usually runs through the ISA bus, and therefore supports up to Mode 2 operation only.



Chapter 15—Hard Disk Interfaces

When interrogated with an Identify Drive command, a hard disk returns, among other things, information about the PIO and DMA modes it is capable of using. Most enhanced BIOSes will automatically set the correct mode to match the capabilities of the drive. If you set a mode faster than the drive can handle, data corruption will result.

ATA-2 drives also perform Block Mode PIO, which means that they use the Read/Write Multiple commands that greatly reduce the number of interrupts sent to host processor. This lowers the overhead, and the resulting transfers are even faster.

DMA Transfer Modes. Although it is not used by most operating system or BIOS software, ATA-2 drives also support DMA transfers. DMA (Direct Memory Access) means that the data is transferred directly between drive and memory without using the CPU as an intermediary, as opposed to PIO.

There are two distinct types of direct memory access: DMA and busmastering DMA. Ordinary DMA relies on the DMA controller on the system's mainboard to perform the complex task of arbitration, grabbing the system bus and transferring the data. In the case of busmastering DMA, all this is done by logic on the interface card itself. Of course, this adds considerably to the complexity and the price of a busmastering interface.

Unfortunately, the DMA controller on ISA systems is ancient and slow, and out of the question for use with a modern hard disk. Today, proper software support for DMA is still rare.

ATAPI (ATA Packet Interface). ATAPI is a standard designed for devices such as CD-ROMs and tape drives that plug into an ordinary ATA (IDE) connector. The principal advantage of ATAPI hardware is that it's cheap and works on your current adapter. For CD-ROMs, it has a somewhat lower CPU usage compared to proprietary adapters but there's no performance gain otherwise. For tape drives, ATAPI has potential for superior performance and reliability compared to the popular "floppy" tape devices.

While ATAPI CD-ROMs use the hard disk interface, this does not mean that they look like an ordinary hard disk; to the contrary, from a software point of view they are a completely different kind of animal. They actually most closely resemble a SCSI device.

This means that intelligent (e.g. caching) controllers that are not ATAPI aware will NOT work with these devices. This also means that, at present, you cannot boot from an ATAPI CD-ROM and you still must load a driver to use it under DOS or Windows. Windows 95 has native ATAPI support, and the first ATAPI-aware BIOS that will even allow booting from an ATAPI CD-ROM has already been introduced.

IDE Drive Configuration. IDE (Integrated Drive Electronics) drives can be both simple and troublesome to configure. Single-drive installations usually are very simple, with few, if any, special jumper settings to worry about. Multiple-drive configurations, however, can be a problem. Jumpers have to be set on both drives; and the names, locations, and even functions of these jumpers can vary from drive to drive.

Because the CAM ATA (Common Access Method AT Attachment) IDE specification was ironed out only after many companies were already making and selling drives, many older IDE drives have problems in dual-drive installations, especially when the drives are

from different manufacturers. In some cases, two particular drives may not function together at all. Fortunately, most of the newer drives follow the CAM ATA specification, which clears up this problem. Drives that follow the specification have no problems in dual-drive installations.

Cable Configuration. The cable connection to IDE drives is usually very simple. There is a single 40-pin cable that normally has three pin-header style connectors on it. One of the connectors plugs into the IDE interface connector; the other two plug into the primary and secondary drives. The cable normally runs from the IDE connector to both drives in a daisy-chain arrangement. On one end, this cable plugs into the IDE interface connector, which is located on the motherboard in many systems but also may be located on an IDE interface adapter card. The cable then connects to the secondary (D) and primary (C) drives in succession, with the primary drive usually (but not always) being at the end of the cable opposite the IDE interface connector.

There are no terminating resistors to set with IDE drives; instead, a distributed termination circuit is built into all IDE drives. The last drive on the cable need not be the primary drive, so you actually may find the primary or secondary drive at either connector. Jumpers on the drives themselves normally control whether a drive responds as primary or secondary.

You may see a different arrangement of cable connections in some IDE installations. In some installations, the middle connector is plugged into the motherboard, and the primary and secondary drives are at opposite ends of the cable in a Y arrangement. If you see this arrangement, be careful; in some of these Y-cable installations, the cable, rather than jumpers on the drives, actually controls which drive is primary and which is secondary.

Controlling master/slave selection via the cable rather than jumpers on the drive is performed via a special signal on the IDE interface called CSEL (Cable SElect), which is on pin 28 of the interface. If the CSEL line is connected through from the drive to the IDE interface connector, the drive automatically is designated as primary. If the CSEL line is open between a drive and the IDE interface connector, that drive automatically is designated as secondary.

In the Y-cable approach, the IDE interface connector is in the middle of the cable, and a separate length of cable goes to each drive. Study this type of cable closely. If one of the ends of the Y has line 28 open (usually, a hole in the cable through that wire), only the secondary drive can be plugged into that connector. HP Vectra PC systems use exactly this type of IDE cable arrangement. This type of setup eliminates the need to set jumpers on the IDE drives to configure them for primary or secondary operation, but the setup can be troublesome if you do not know about it.

IDE Drive Jumper Settings. Configuring IDE drives can be simple, as is the case with most single-drive installations, or troublesome, especially where it comes to mixing two drives from different manufacturers on a single cable.

Most IDE drives come in three configurations:



Chapter 15—Hard Disk Interfaces

- Single-drive (master)
- Master (dual-drive)
- Slave (dual-drive)

Because each IDE drive has its own controller, you must specifically tell one drive to be the master and the other to be the slave. There's no functional difference between the two, except that the drive that's specified as the slave will assert the DASP signal after a system reset that informs the master that a slave drive is present in the system. The master drive then pays attention to the Drive Select line, which it otherwise ignores. Telling a drive that it's the slave also usually causes it to delay its spinup for several seconds to allow the master to get going and thus to lessen the load on the system's power supply.

Until the ATA IDE specification, no common implementation for drive configuration was in use. Some drive companies even used different master/slave methods for different models of drives. Because of these incompatibilities, some drives work together only in a specific master/slave or slave/master order. This situation affects mostly older IDE drives that were introduced before the ATA specification.

Most drives that fully follow the ATA specification now need only one jumper (Master/Slave) for configuration. A few also need a Slave Present jumper as well. Table 15.7 shows the jumper settings required by most ATA IDE drives.

Jumper Name	Single-Drive	Dual-Drive Master	Dual-Drive Slave
Master (M/S)	On	On	Off
Slave Present (SP)	Off	On	Off

The Master jumper indicates that the drive is a master or a slave. Some drives also require a Slave Present jumper, which is used only in a dual-drive setup and then installed only on the master drive, which is somewhat confusing. This jumper tells the master that a slave drive is attached. With many ATA IDE drives, the Master jumper is optional and may be left off. Installing this jumper doesn't hurt in these cases and may eliminate confusion, so I recommend that you install the jumpers listed here.

Conner Peripherals Drives. Because they were introduced before the ATA IDE specification was formalized, Conner Peripherals drives often are different in configuration from many other-brand drives. When you mix and match IDE hard drives from different manufacturers, the drives are not always fully compatible. Table 15.8 shows the jumper settings that are correct for most Conner IDE drive installations.

The C/D jumper is used to determine whether the drive is a master (drive C) or a slave (drive D). The drive is configured as master when this jumper is on. The DSP jumper indicates that a slave drive is present. The HSP jumper causes the drive to send the Slave Present signal to the master drive. The ACT jumper enables the master drive to signal when it is active.

Jumper Name	Single-Drive	Dual-Drive Master	Dual-Drive Slave
Master or Slave (C/D)	On	On	Off
Drive Slave Present (DSP)	Off	On	Off
Host Slave Present (HSP)	Off	Off	On
Drive Active (ACT)	On	On	Off

Some Conner drives are not set up to support the industry-standard CAM ATA (Common Access Method AT Attachment) interface by default. The problems show up when you attempt to connect another manufacturer's drive to some Conner drives in either a master or slave role. Fortunately, you can correct many of these situations by changing the configuration of the drive.

You can make this change in two ways. One way is to use a special program to semipermanently change the mode of the drive. A special file available on the Conner BBS, called FEATURE.COM, contains a program that displays the current ISA/ATACAM setting and allows the setting to be changed. The change is actually stored in a feature byte in the firmware of the drive, and after this byte is changed, most other manufacturers' drives will work with the Conner drives. The program also can be used to reset the feature byte back to its original configuration, which is best when you are connecting to other Conner drives.

The second method for changing this configuration is available on some Conner drives. These drives also have a special jumper called ATA/ISA. This jumper almost always should be installed in the ATA position to provide compatibility with the ATA standard. If you are using only Conner drives, you can leave this jumper in ISA mode if you want. Some Conner drives have a separate jumper (E1) that can delay startup of the drive to minimize the load on the power supply. This jumper should be enabled on any drive that is configured as a slave. Most other drives automatically delay startup of the slave drive for a few seconds.

Most Conner drives also have a special 12-pin connector that is used to drive an optional LED (pin 1, LED +5v; and pin 2, ground), as well as to connect to special factory equipment for low-level formatting and configuration. A company called TCE (see Appendix B) sells a device called The Conner, which connects to this port and permits full factory-level initialization, formatting, and testing of Conner drives. I consider this piece of gear to be essential to anybody who services or supports a large number of Conner Peripheral drives. Notice that Compaq uses Conner drives in most of its systems.

For more information on any specific Conner drive, you can use the company's FaxBack system (see the vendor list in Appendix B) at (800) 4CONNER. Through this system, you can get drive information and jumper settings that are specific to Conner drives.

XT-Bus (8-bit) IDE. Many systems with XT ISA bus architecture used XT IDE hard drives. The IDE interface in these systems is usually built into the motherboard. The IBM PS/2 Model 25, 25-286, 30, and 30-286 systems used an 8-bit XT IDE interface. These 8-bit XT IDE drives are difficult to find, few manufacturers other than IBM, Western

Chapter 15—Hard Disk Interfaces

Digital, and Seagate made them; and none of these drives were available in capacities beyond 40M.

Since the ATA IDE interface is a 16-bit design, it could not be used in 8-bit (XT type) systems, so some of the drive manufacturers standardized on an XT-Bus (8-bit) IDE interface for XT class systems. These drives were never very popular, and were usually only available in capacities from 20M to 40M. Table 15.9 shows the industry standard 8-bit IDE connector pinout.

Signal Name	Pin	Pin	Signal Name
-RESET	1	2	GROUND
Data Bit 7	3	4	GROUND
Data Bit 6	5	6	GROUND
Data Bit 5	7	8	GROUND
Data Bit 4	9	10	GROUND
Data Bit 3	11	12	GROUND
Data Bit 2	13	14	GROUND
Data Bit 1	15	16	GROUND
Data Bit 0	17	18	GROUND
GROUND	19	20	KEY (pin missing)
AEN	21	22	GROUND
-IOW	23	24	GROUND
-IOR	25	26	GROUND
-DACK 3	27	28	GROUND
DRQ 3	29	30	GROUND
IRQ 5	31	32	GROUND
Address Bit 1	33	34	GROUND
Address Bit 0	35	36	GROUND
-CS1FX	37	38	GROUND
-Drive Active	39	40	GROUND

Notice that IBM used a custom version of the XT-Bus IDE interface in the PS/2 Model 25 and Model 30 systems. The pinout for the custom IBM XT-Bus IDE connector is shown in table 15.10.

Signal Name	Pin	Pin	Signal Name
-RESET	1	2	-Disk Installed
Data Bit 0	3	4	GROUND
Data Bit 1	5	6	GROUND
Data Bit 2	7	8	GROUND



Signal Name	Pin	Pin	Signal Name
Data Bit 3	9	10	GROUND
Data Bit 4	11	12	GROUND
Data Bit 5	13	14	GROUND
Data Bit 6	15	16	GROUND
Data Bit 7	17	18	GROUND
-IOR	12	20	GROUND
-IOW	21	22	GROUND
-CS1FX	23	24	GROUND
Address Bit 0	25	26	GROUND
Address Bit 1	27	28	GROUND
Address Bit 2	29	30	+5 Vdc
RESERVED	31	32	+5 Vdc
-DACK 3	33	34	GROUND
DRQ 3	35	36	GROUND
IRQ 5	37	38	GROUND
I/O CH RDY	39	40	+12 Vdc
Spare	41	42	+12 Vdc
Spare	39	44	+12 Vdc

The newer PS/1, PS/Valuepoint, and PS/2 systems with 16-bit ISA architecture use ATA IDE drives. Because nearly all hard disk manufacturers make a multitude of drives with the ATA IDE interface, these systems are easy to upgrade or repair. ATA IDE drives are available in capacities up to and beyond 1G.

MCA IDE. The IBM PS/2 Models 50 and higher come with Micro Channel Architecture (MCA) bus slots. Although most of these systems now use SCSI drives, for some time IBM used a type of MCA IDE drive in these systems. MCA IDE is a form of IDE interface, but it is designed for the MCA bus and is not compatible with the more industry-standard ATA IDE interface. Few companies other than IBM and Western Digital make replacement MCA IDE drives for these systems. I recommend replacing these drives with ATA IDE drives, using adapters from Arco Electronics or Sigma Data, or switching to SCSI drives instead. The IBM MCA IDE drives are expensive for the limited capacity that they offer.

The pinout of the MCA IDE connector is shown in table 15.11.

Signal Name	Pin	Pin	Signal Name
-CD SETUP	A1	B1	Address Bit 15
Address Bit 13	A2	B2	Address Bit 14
GROUND	A3	B3	GROUND
Address Bit 11	A4	B4	OSC (14.3 MHz)

(continues)

Chapter 15—Hard Disk Interfaces

Signal Name	Pin	Pin	Signal Name
Address Bit 10	A5	B5	GROUND
Address Bit 9	A6	B6	Address Bit 12
+5 Vdc	A7	B7	-CMD
Address Bit 8	A8	B8	-CD SFDBK
Address Bit 7	A9	B9	GROUND
Address Bit 6	A10	B10	Data Bit 1
+5 Vdc	A11	B11	Data Bit 3
Address Bit 5	A12	B12	Data Bit 4
Address Bit 4	A13	B13	GROUND
Address Bit 3	A14	B14	CHRESET
+5 Vdc	A15	B15	Data Bit 8
Address Bit 2	A16	B16	Data Bit 9
Address Bit 1	A17	B17	GROUND
Address Bit 0	A18	B18	Data Bit 12
+12 Vdc	A19	B19	Data Bit 14
-ADL	A20	B20	Data Bit 15
-PREEMPT	A21	B21	GROUND
-BURST	A22	B22	Data Bit 0
+5 Vdc	A23	B23	Data Bit 2
ARB 0	A24	B24	Data Bit 5
ARB 1	A25	B25	GROUND
ARB 2	A26	B26	Data Bit 6
+12 Vdc	A27	B27	Data Bit 7
ARB 3	A28	B28	Data Bit 10
-ARB/-GRANT	A29	B29	GROUND
-TC	A30	B30	Data Bit 11
+5 Vdc	A31	B31	Data Bit 13
-S0	A32	B32	-SBHE
-S1	A33	B33	GROUND
+M/-IO	A34	B34	-CD DS 16
GROUND	A35	B35	-IRQ 14
CD CHRDY	A36	B36	GROUND

Introduction to SCSI

SCSI (pronounced "scuzzy") stands for Small Computer System Interface. This interface has its roots in SASI, the Shugart Associates System Interface. SCSI is not a disk interface, but a systems-level interface. SCSI is not a type of controller, but a bus that supports as many as eight devices. One of these devices, the host adapter, functions as the gateway between the SCSI bus and the PC system bus. The SCSI bus itself does not talk directly with devices such as hard disks; instead, it talks to the controller that is built into the drive.

A single SCSI bus can support as many as eight physical units, usually called SCSI IDs. One of these units is the adapter card in your PC; the other seven can be other peripherals. You could have hard disks, tape drives, CD-ROM drives, a graphics scanner, or other devices (up to seven total) attached to a single SCSI host adapter. Most systems support up to 4 host adapters, each with 7 devices, for a total 28 devices! Some of the newer SCSI implementations allow for 15 devices on each bus.

When you purchase a SCSI hard disk, you are usually purchasing the drive, controller, and SCSI adapter in one circuit. This type of drive is usually called an embedded SCSI drive; the SCSI interface is built into the drive. Most SCSI hard drives actually are IDE drives with SCSI bus adapter circuits added. You do not need to know what type of controller is inside the SCSI drive because your system cannot talk directly to the controller as though it were plugged into the system bus, like a standard controller. Instead, communications go through the SCSI host adapter installed in the system bus. You can access the drive only with the SCSI protocols.

Apple originally rallied around SCSI as being an inexpensive way out of the bind in which it put itself with the Macintosh. When the engineers at Apple realized the problem in making the Macintosh a closed system (with no slots), they decided that the easiest way to gain expandability was to build a SCSI port into the system, which is how external peripherals can be added to the slotless Macs. Because PC systems always have been expandable, the push toward SCSI has not been as urgent. With eight bus slots supporting different devices and controllers in IBM and IBM-compatible systems, it seemed as though SCSI was not needed.

SCSI now is becoming popular in the IBM-based computer world because of the great expandability that it offers and the number of devices that are available with built-in SCSI. One thing that stalled acceptance of SCSI in the PC marketplace was the lack of a real standard; the SCSI standard was designed primarily by a committee. No single manufacturer has led the way, at least in the IBM arena, each company has its own interpretation of how SCSI should be implemented, particularly at the host-adapter level.

SCSI is a standard, in much the same way that RS-232 is a standard. The SCSI standard (like the RS-232 standard), however, defines only the hardware connections, not the driver specifications required to communicate with the devices. Software ties the SCSI subsystem into your PC, but unfortunately, most of the driver programs work only for a specific device and a specific host adapter. For example, a graphics scanner comes with its own SCSI host adapter to connect to the system, a CD-ROM drive comes with another (different) SCSI host adapter and driver software that works only with that SCSI adapter. On a system with those two SCSI adapters, you would need a third SCSI host adapter to run SCSI hard disk drives because the host adapters supplied by the scanner and CD-ROM companies do not include a built-in, self-booting BIOS that supports hard disk drives.

SCSI has become something of a mess in the IBM world because of the lack of a host-adapter standard, a software interface standard, and standard ROM BIOS support for hard disk drives attached to the SCSI bus. Fortunately, some simple recommendations can keep you from living this compatibility nightmare!



Chapter 15—Hard Disk Interfaces

The lack of capability to run hard disks off the SCSI bus, to boot from these drives and use a variety of operating systems, is a problem that results from the lack of an interface standard. The standard IBM XT and AT ROM BIOS software was designed to talk to ST-506/412 hard disk controllers. The software easily was modified to work with ESDI because ESDI controllers are similar to ST-506/412 controllers at the register level. (This similarity at the register level enabled manufacturers to easily design self-booting, ROM-BIOS-supported ESDI drives.) The same can be said of IDE, which completely emulates the WD1003 ST-506/412 controller interface and works perfectly with the existing BIOS as well. SCSI is so different from these other standard disk interfaces that a new set of ROM BIOS routines are necessary to support the system so that it can self-boot. The newer IBM PS/2 systems that come with SCSI drives have this support built into the motherboard BIOS or as an extension BIOS on the SCSI host adapter.

Companies such as Adaptec and Future Domain have produced SCSI cards with built-in ROM BIOS support for several years, but these BIOS routines were limited to running the drives only under DOS. The BIOS would not run in the AT-protected mode, and other operating systems included drivers for only the standard ST-506/412 and ESDI controllers. Thus, running SCSI was impossible under many non-DOS operating systems. This situation has changed significantly, however; IBM now supports many third-party SCSI host adapters in OS/2, especially those from Adaptec and Future Domain. For compatibility reasons, I usually recommend using SCSI adapters from these two companies, or any other adapters that are fully hardware-compatible with the Adaptec and Future Domain adapters.

Because of the lead taken by Apple in developing systems software (operating systems and ROM) support for SCSI, peripherals connect to Apple systems in fairly standard ways. Until recently, this kind of standard-setting leadership was lacking for SCSI in the IBM world. This situation changed on March 20, 1990, when IBM introduced several "standard" SCSI adapters and peripherals for the IBM PS/2 systems, with complete ROM BIOS and full operating-system support.

IBM has standardized on SCSI for nearly all its high-end systems. In these systems, a SCSI host adapter card is in one of the slots, or the system has a SCSI host adapter built into the motherboard. This arrangement is similar in appearance to the IDE interface because a single cable runs from the motherboard to the SCSI drive, but SCSI supports as many as seven devices (some of which may not be hard disks), whereas IDE supports only two devices, which must be either a hard disk or a tape drive. PS/2 systems with SCSI drives are easy to upgrade because virtually any third-party SCSI drive will plug in and function.

The example set by IBM is causing other manufacturers to supply systems with either SCSI host adapters or SCSI interfaces integrated into the motherboards. As SCSI becomes more and more popular in the PC world, SCSI peripheral integration will be easier, due to better operating-system and device-driver support.

ANSI SCSI standards. The SCSI standard defines the physical and electrical parameters of a parallel I/O bus used to connect computers and peripheral devices in daisy-chain fashion. The standard supports devices such as disk drives, tape drives, and CD-ROM

drives. The original SCSI standard (ANSI X3 131-1986) was approved in 1986; SCSI-2 was approved in January of 1994, and a new revision called SCSI-3, is being developed.

The SCSI interface is defined as a standard by ANSI (American National Standards Institute), an organization that approves and publishes standards. The X3 Task Group operates as an ASC (Accredited Standards Committee) under ANSI to develop Information Processing System standards. X3T9 is the I/O Interfaces group, and X3T9.2 specifically is in charge of low-level interfaces such as SCSI and ATA-IDE (among others). The original SCSI-1 standard was published by the X3T9 ANSI group in 1986 and is officially published by ANSI as X3.131-1986.

One problem with the original SCSI-1 document was that many of the commands and features were optional, and there was little or no guarantee that a particular peripheral would support the expected commands. This problem caused the industry as a whole to define a set of 18 basic SCSI commands called the Common Command Set (CCS), which would become the minimum set of commands supported by all peripherals. This Common Command Set became the basis for what is now the SCSI-2 specification.

In addition to formal support for the CCS, SCSI-2 provided additional definitions for commands to access CD-ROM drives (and their sound capabilities), tape drives, removable drives, optical drives, and several other peripherals. In addition, an optional higher speed called FAST SCSI-2 and a 16-bit version called WIDE SCSI-2 were defined. Another feature of SCSI-2 is command queuing, which enables a device to accept multiple commands and execute them in the order that the device deems to be most efficient. This feature is most beneficial when you are using a multitasking operating system that could be sending several requests on the SCSI bus at the same time.

The X3T9 group approved the SCSI-2 standard as X3.131-1990 in August 1990, but the document was recalled in December 1990 for changes before final ANSI publication. Final approval for the SCSI-2 document was finally made in January of 1994(!), although it has changed little from the original 1990 release. The SCSI-2 document is now called ANSI X3 131-1994. The official document is available from Global Engineering Documents or the ANSI committee, which are listed in Appendix B. You can also download working drafts of these documents from the NCR SCSI BBS, listed in the vendor list under NCR Microelectronics.

Most companies indicate that their host adapters follow both the ANSI X3.131-1986 (SCSI-1) as well as the X3.131-1994 (SCSI-2) standards. Note that since virtually all parts of SCSI-1 are supported in SCSI-2, virtually any SCSI-1 device is also considered SCSI-2 by default. Many manufacturers advertise that their devices are SCSI-2, but this does not mean that they support any of the additional optional features that were incorporated in the SCSI-2 revision.

For example, an optional part of the SCSI-2 specification includes a fast synchronous mode that doubles the standard synchronous transfer rate from 5MB per second to 10MB per second. This Fast SCSI transfer mode can be combined with 16-bit Wide SCSI for transfer rates of up to 20MB per second. There was an optional 32-bit version defined in SCSI-2, but component manufacturers have shunned this as too expensive. In essence,



Chapter 15—Hard Disk Interfaces

32-bit SCSI was a stillborn specification. Most SCSI implementations are 8-bit standard SCSI or Fast/Wide SCSI. Even devices which support none of the fast or wide modes can still be considered SCSI-2.

The SCSI-3 standard is still being defined and is still a long way from being approved. However, portions of this specification, although not final, are being sold in products today. One of these developments is the new Fast-20 mode, which is also called Ultra-SCSI. This essentially is quad-speed SCSI, and will run 20MB per second on an 8-bit standard SCSI bus, and 40MB per second on Wide (16-bit) SCSI.

Table 15.12 shows the maximum transfer rates for the SCSI bus at various speeds and widths, as well as the cable type required for the specific transfer widths.

Bus Width	Std. SCSI	Fast SCSI	Fast-20 (Ultra) SCSI	Cable Type
8-bit	5MB/s	10MB/s	20MB/s	A (50-pin)
16-bit (Wide)	10MB/s	20MB/s	40MB/s	P (68-pin)

MB/s = MegaBytes per second.

Fast-20 SCSI is also called Ultra-SCSI



The A cable is the standard 50-pin SCSI cable, whereas the P cable is a 68-pin cable designed for 16-bit. Maximum cable length is 6 meters (about 20 feet) for standard speed SCSI, and only 3 meters (about 10 feet) for Fast or Fast-20 (Ultra) SCSI. Pinouts for these cable connections are listed in this chapter.

So-called SCSI-1 adapters have no problems with SCSI-2 peripherals. In fact, as was stated earlier, virtually any SCSI-1 device can also legitimately be called SCSI-2 (or even SCSI-3). You can't take advantage of Fast, Fast-20, or Wide transfer capabilities, but the extra commands defined in SCSI-2 can be sent by means of a SCSI-1 controller. In other words, nothing is different between SCSI-1 and SCSI-2 compliant hardware. For example, I am running a Seagate Barracuda 4G Fast SCSI-2 drive with my standard IBM SCSI-1 host adapter, and it runs fine. Most adapters are similar, in that they actually are SCSI-2 compatible, even if they advertise only SCSI-1 support. Because the SCSI-2 standard was not actually approved before January 1994, any devices that claimed to be SCSI-2 before that time were not officially in compliance with the standard. This is really not a problem, however, since the SCSI 2 document had not changed appreciably since it was nearly approved in 1990. The same thing is currently happening with advertisers listing devices as "SCSI-3." The SCSI-3 specification is not yet approved, although certain areas are being worked out.

SCSI Hard Disk Evolution and Construction. SCSI is not a disk interface, but a bus that supports SCSI bus interface adapters connected to disk and other device controllers. The

first SCSI drives for PCs simply were standard ST-506/412 or ESDI drives with a separate SCSI bus interface adapter (sometimes called a bridge controller) that converted the ST-506/412 or ESDI interfaces to SCSI. This interface originally was in the form of a secondary logic board, and the entire assembly often was mounted in an external case.

The next step was to build the SCSI bus interface "converter" board directly into the drive's own logic board. Today, we call these drives embedded SCSI drives because the SCSI interface is built in.


At that point, there was no need to conform to the absolute specifications of ST-506/412 or ESDI on the internal disk interface because the only other device that the interface would ever have to talk to was built in as well. Thus, the disk-interface and controller-chipset manufacturers began to develop more customized chipsets that were based on the ST-506/412 or ESDI chipsets already available but offered more features and higher performance. Today, if you look at a typical SCSI drive, you often can identify the chip or chipset that serves as the disk controller on the drive as being exactly the same kind that would be used on an ST-506/412 or ESDI controller or as being some evolutionary customized variation thereof.

Consider some examples. An ATA (AT Attachment interface) IDE drive must fully emulate the system-level disk-controller interface introduced with the Western Digital WD1003 controller series that IBM used in the AT. These drives must act as though they have a built-in ST-506/412 or ESDI controller; in fact, they actually do. Most of these built-in controllers have more capabilities than the original WD1003 series (usually in the form of additional commands), but they must at least respond to all the original commands that were used with the WD1003.

If you follow the hard drive market, you usually will see that drive manufacturers offer most of their newer drives in both ATA IDE and SCSI versions. In other words, if a manufacturer makes a particular 500MB IDE drive, you invariably will see that the company also make a SCSI model with the same capacity and specifications, which uses the same HDA (Head Disk Assembly) and even looks the same as the IDE version. If you study these virtually identical drives, the only major difference you will find is the additional chip on the logic board of the SCSI version, called a SCSI Bus Adapter Chip (SBIC).

Figures 15.3 and 15.4 show the logic-block diagrams of the WD-AP4200 (a 200MB ATA-IDE drive) and WD-SP4200 (a 200MB SCSI drive), respectively. These drives use the same HDA; they differ only in their logic boards, and even the logic boards are the same except for the addition of an SBIC (SCSI Bus Adapter Chip) on the SCSI drive's logic board.

Notice that even the circuit designs of these two drives are almost identical. Both drives use an LSI (Large Scale Integrated circuit) chip called the WD42C22 Disk Controller and Buffer manager chip. In the ATA drive, this chip is connected through a DMA control chip directly to the AT bus. In the SCSI version, a WD33C93 SCSI bus interface controller chip is added to interface the disk-controller logic to the SCSI bus. In fact, the logic diagrams of these two drives differ only in the fact that the SCSI version has a complete subset of the ATA drive with the SCSI bus interface controller logic added. This essentially is a very condensed version of the separate drive and bridge controller setups that were used in the early days of PC SCSI!



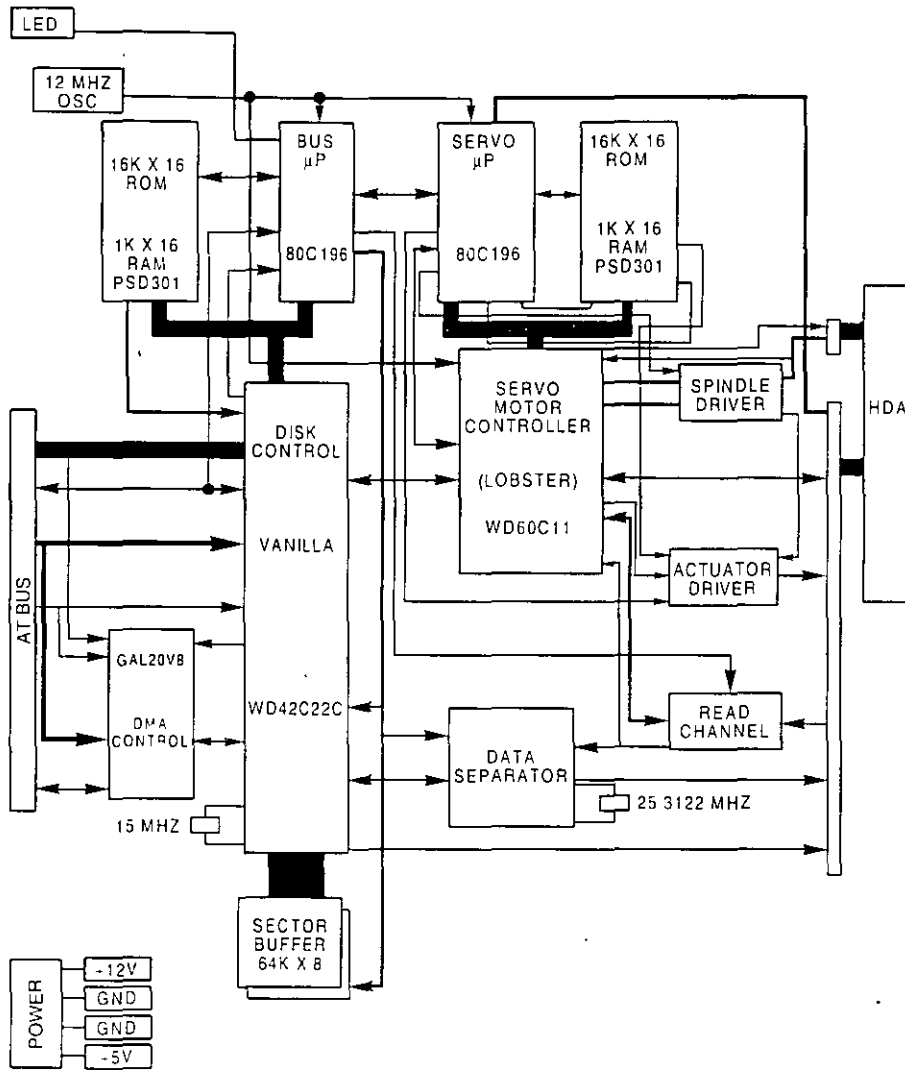


Figure 15.3

Western Digital WD-AP4200 200MB ATA-IDE drive logic-board block diagram

To top off this example, study the following logic diagram for the WD 1006V-MM1, which is an ST-506/412 controller (see fig. 15.5)

You can clearly see that the main LSI chip on board is the same WD42C22 disk controller chip used in the IDE and SCSI drives. Here is what the technical reference literature says about that chip

The WD42C22 integrates a high performance, low cost Winchester controller's architecture. The WD42C22 integrates the central elements of a Winchester

controller subsystem such as the host interface, buffer manager, disk formatter/ controller, encoder/decoder, CRC/ECC (Cyclic Redundancy Check/Error Correction Code) generator/checker, and drive interface into a single 84-pin PQFP (Plastic Quad Flat Pack) device.

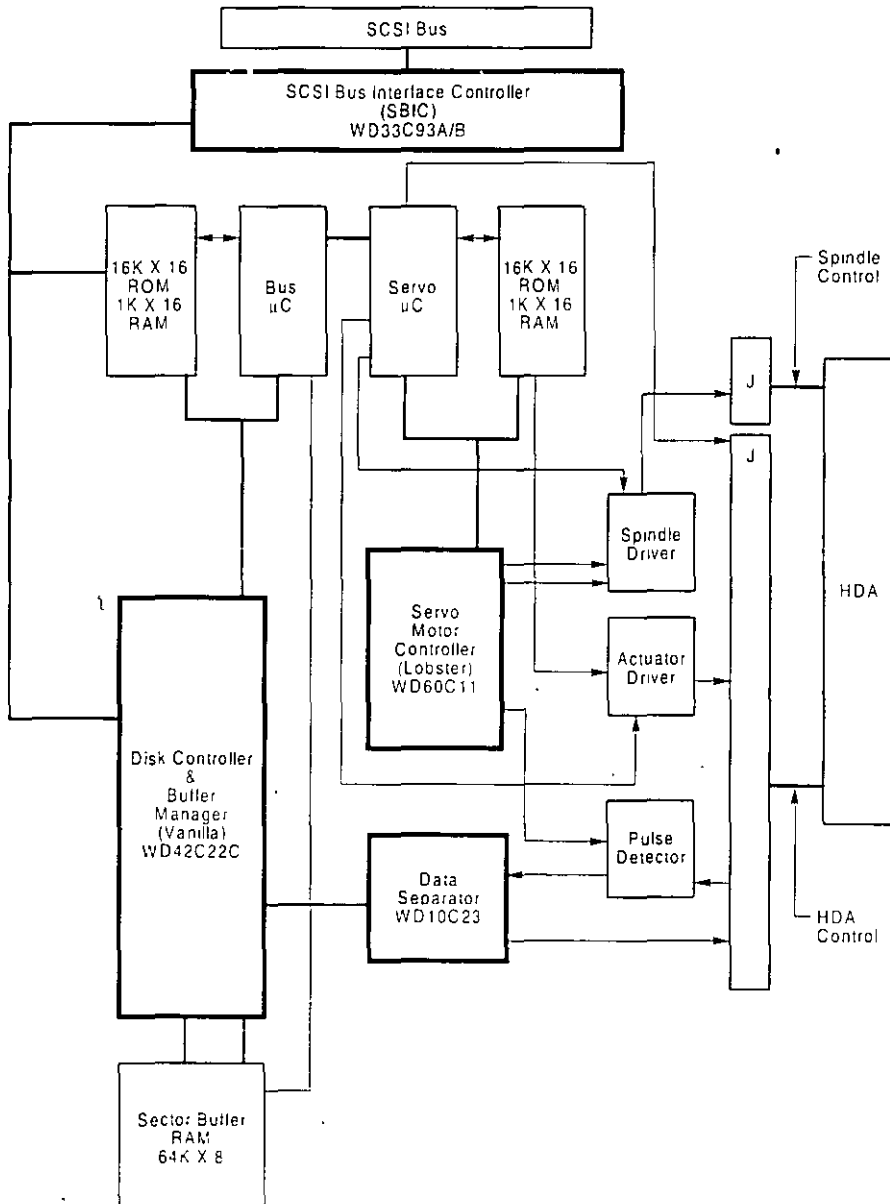


Figure 15.4

Western Digital WID-SP4200 200MB SCSI drive logic-board block diagram

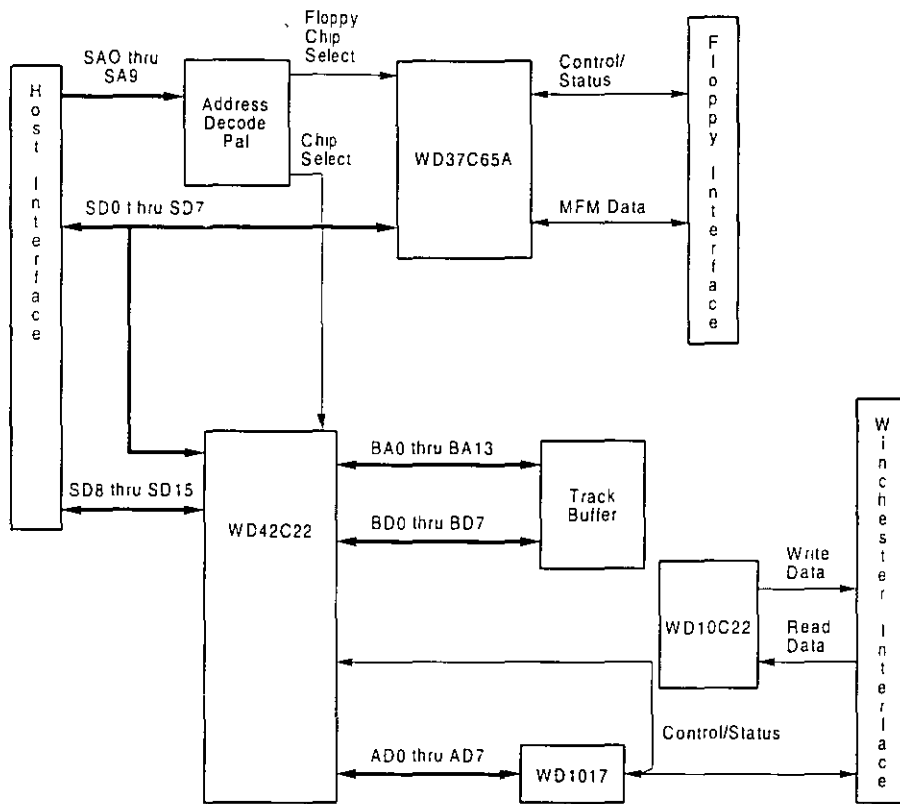


Figure 15.5

Western Digital WD1006V-MM1 S1-506/412 Disk Controller block diagram.

The virtually identical design of ATA-IDE and SCSI drives is not unique to Western Digital. Most drive manufacturers design their ATA-IDE and SCSI drives the same way, often using these very same Western Digital chips as well as disk controller and SCSI bus interface chips from other manufacturers. You now should be able to understand that most SCSI drives simply are "regular" ATA-IDE drives with SCSI bus logic added. This fact will come up again later in this chapter, in the section "SCSI versus IDE," which discusses performance and other issues differentiating these interfaces.

For another example, I have several IBM 320MB and 400MB embedded SCSI-2 hard disks; each of these drives has on-board a WD-10C00 Programmable Disk Controller in the form of a 68-pin PLCC (Plastic Leaded Chip Carrier) chip. The technical literature states, "This chip supports ST-12, ESDI, SMD and Optical interfaces. It has 27-Mbit/sec maximum transfer rate and an internal, fully programmable 48- or 32-bit ECC, 16-bit CRC-CCITT or external user defined ECC polynomial, fully programmable sector sizes and 1.25 micron low power CMOS design "

In addition, these particular embedded SCSI drives include the 33C93 SCSI Bus Interface Controller chip, which is also used in the other SCSI drive that I mentioned. Again, there is a distinctly separate disk controller, and the SCSI interface is added on.


So again, most embedded SCSI drives have a built-in disk controller (usually based on previous ST-506/412 or ESDI designs) and additional logic to interface that controller to the SCSI bus (a built-in bridge controller, if you like). Now think about this from a performance standpoint. If virtually all SCSI drives really are ATA IDE drives with a SCSI Bus Interface Controller chip added, what conclusions can you draw?

First, no drive can perform sustained data transfers faster than the data can actually be read from the disk platters. In other words, the HDA (Head Disk Assembly) limits performance to whatever it is capable of achieving. Drives can transmit data in short bursts at very high speeds because they often have built-in cache or read-ahead buffers that store data. Many of the newer high-performance SCSI and ATA-IDE drives have 1M or more of cache memory on-board! No matter how big or intelligent the cache is, however, sustained data transfer still will be limited by the HDA.

Data from the HDA must pass through the disk controller circuits which, as you have seen, are virtually identical between similar SCSI and ATA-IDE drives. In the ATA IDE drive, this data then is presented directly to the system bus. In the SCSI drive, however, the data must pass through a SCSI Bus Interface adapter on the drive, travel through the SCSI bus itself, and then pass through another SCSI Bus Interface controller in the SCSI host adapter card in your system. The longer route that a SCSI transfer must take makes this type of transfer slower than the much more direct ATA IDE transfer.

The conventional wisdom has been that SCSI is always much faster than IDE; unfortunately, this wisdom usually is wrong! This incorrect conclusion was derived by looking at the raw SCSI and ISA bus performance capabilities. An 8-bit Fast SCSI-2 bus can transfer data at 10MB (million bytes) per second, whereas the 16-bit ISA bus used directly by IDE drives can transfer data at rates ranging from 2MB to 8MB per second. Based on these raw transfer rates, SCSI seems to be faster, but the raw transfer rate of the bus is not the limiting factor. Instead, the actual HDA and disk-controller circuitry place the limits on performance. Another point to remember is that unless you are using a PCI, VL-Bus, EISA, or 32-bit MCA SCSI adapter, the SCSI data-transfer speeds will be limited by the host bus performance as well as by the drive performance.

Single-Ended or Differential SCSI. "Normal" SCSI also is called single-ended SCSI. For each signal that needs to be sent across the bus, a wire exists to carry it. With differential SCSI, for each signal that needs to be sent across the bus, a pair of wires exists to carry it. The first in this pair carries the same type of signal that the single-ended SCSI carries. The second in this pair, however, carries the logical inversion of the signal. The receiving device takes the difference of the pair (hence the name differential), which makes it less susceptible to noise and allows for greater cable length. Because of this, differential SCSI can be used with cable lengths up to 25 meters, whereas single-ended SCSI is good only for 6 meters with standard asynchronous or synchronous transfers or for only 3 meters for Fast SCSI.



You cannot mix single-ended and differential devices on a single SCSI bus; the result would be catastrophic. (That is to say that you probably will see smoke!) Notice that the cables and connectors are the same, so it's entirely possible to make this mistake. This usually is not a problem, however, because very few differential SCSI implementations exist. Especially with SCSI in the PC environment, single-ended is about all you will ever see. If, however, you do come upon a peripheral that you believe might be differential, there are a few ways to tell. One way is to look for a special symbol on the unit; the industry has adopted different universal symbols for single-ended and differential SCSI. Figure 15.6 shows these symbols.




Figure 15.6

Single-ended and differential SCSI universal symbols

If you do not see such symbols, you can tell whether you have a differential device by using an ohm meter to check the resistance between pins 21 and 22 on the device connector. On a single-ended system, the pins should be tied together and also tied to Ground. On a differential device, the pins should be open or have significant resistance between them. Again, this generally should not be a problem because virtually all devices used in the PC environment are single-ended.

SCSI-1 and SCSI-2. The SCSI-2 specification essentially is an improved version of SCSI-1 with some parts of the specification tightened and with several new features and options added. Normally, SCSI-1 and SCSI-2 devices are compatible, but SCSI-1 devices ignore the additional features in SCSI-2.

Some of the changes in SCSI-2 are very minor. For example, SCSI-1 allowed SCSI Bus parity to be optional, whereas parity must be implemented in SCSI-2. Another requirement is that initiator devices, such as host adapters, provide terminator power to the interface; most devices already did so.

SCSI-2 also has several optional features:

- Fast SCSI
- Wide SCSI
- Command queuing
- High-density cable connectors
- Improved Active (Alternative 2) termination

These features are not required; they are optional under the SCSI-2 specification. If you connect a standard SCSI host adapter to a Fast SCSI drive, for example, the interface will work, but only at standard SCSI speeds.

Fast SCSI. Fast SCSI refers to high-speed synchronous transfer capability. Fast SCSI achieves a 10MB per second transfer rate on the standard 8-bit SCSI cabling. When combined with a 16-bit Wide SCSI interface, this configuration results in data-transfer rates of 20MB per second (called Fast/Wide).

Fast-20 (Ultra) SCSI. Fast-20 or Ultra SCSI refers to high-speed synchronous transfer capability that is twice as fast as Fast-SCSI. This has been introduced in the Draft (unfinished) SCSI-3 specification and has already been adopted by the marketplace, especially for high speed hard disks. Ultra SCSI achieves a 20MB per second transfer rate on the standard 8-bit SCSI cabling. When combined with a 16-bit Wide SCSI interface, this configuration results in data-transfer rates of 40MB per second (called Ultra/Wide).

Wide SCSI. Wide SCSI allows for parallel data transfer at a bus width of 16 bits. The wider connection requires a new cable design. The standard 50-conductor 8-bit cable is called the A-cable. SCSI-2 originally defined a special 68-conductor B-cable that was supposed to be used in conjunction with the A cable for wide transfers, but the industry ignored this specification in favor of a newer 68-conductor P-cable that was introduced as part of the SCSI-3 specification. The P-cable superseded the A- and B-cable combination because the P-Cable can be used alone (without the A cable) for 16-bit Wide SCSI.

A 32-bit Wide SCSI version was originally defined on paper as a part of the SCSI-2 specification, but has not found popularity and probably never will in the PC environment. Theoretically, 32-bit SCSI implementations would require two cables, a 68-conductor P-cable and a 68-conductor Q-cable.

Termination. The single-ended SCSI bus depends on very tight termination tolerances to function reliably. Unfortunately, the original 132-ohm passive termination defined in the SCSI-1 document was not designed for use at the higher synchronous speeds now possible. These passive terminators can cause signal reflections to cause errors when transfer rates increase or when more devices are added to the bus. SCSI-2 defines an active (voltage-regulated) terminator that lowers termination impedance to 110 ohms and improves system integrity.



Command Queuing. In SCSI-1, an initiator device, such as a host adapter, was limited to sending one command per device. In SCSI-2, the host adapter can send as many as 256 commands to a single device, which will store and process those commands internally before responding on the SCSI bus. The target device even can resequence the commands to allow for the most efficient execution or performance possible. This feature is especially useful in multitasking environments, such as OS/2 and Windows NT, that can take advantage of this feature.

New Commands. SCSI-2 took the Common Command Set that was being used throughout the industry and made it an official part of the standard. The CCS was designed mainly for disk drives and did not include specific commands designed for other types of devices. In SCSI-2, many of the old commands are reworked and several new commands have been added. New command sets have been added for CD-ROMs, optical drives, scanners, communications devices, and media changers (jukeboxes).

SCSI-3. Even though the SCSI-2 specification has only recently been approved (although it has remained stable for some time), the SCSI-3 specification is already being developed. SCSI-3 will have everything that SCSI-2 has and definitely will add new commands, features, and implementations. For example, SCSI-3 will provide support for up to 32 devices on the bus instead of only 8.

One of the most exciting things about SCSI-3 is the proposed Serial SCSI, a scheme that may use only a six-conductor cable and that will be able to transfer data at up to 100MB per second! The switch to serial instead of parallel is designed to control the delay, noise, and termination problems that have plagued SCSI-2, as well as to simplify the cable connection. Serial SCSI will be capable of transferring more data over 6 wires than 32-bit Fast Wide SCSI-2 can over 128 wires! The intention is that Serial SCSI be implemented on the motherboard of future systems, giving them incredible expansion and performance capabilities.

Although Serial SCSI may not make the older host adapters and cables obsolete overnight, it does make future cabling possibilities even more of a puzzle. Serial SCSI offers the possibility of longer cable lengths; less electromagnetic interference; and easier connections on laptops, notebooks, and docking stations. Expect SCSI-3 to offer almost pain-free installations with automatic plug-and-play SCSI ID setup and termination schemes.

In any practical sense, SCSI-3 is still some ways away from being approved. Because the standard exists in draft documents before being officially approved, if the portions of the standard become stable, we may very well see products claiming SCSI-3 compatibility well before the standard truly exists. Since SCSI-3 actually incorporates all of what is in SCSI-2, technically anybody can call any SCSI-1 or SCSI-2 device a SCSI-3 device as well. Beware of product hype along these lines. Some of the new SCSI-3 features will likely be incompatible with previous SCSI implementations, and may take a while to appear on the market.

SCSI Cables and Connectors. The SCSI standards are very specific when it comes to cables and connectors. The most common connectors specified in this standard are the

50-position unshielded pin header connector for internal SCSI connections and the 50-position shielded Centronics latch-style connectors for external connections. The shielded Centronics style connector is also called Alternative 2 in the official specification. Passive or Active termination (Active is preferred) is specified for both single-ended and differential buses. The 50-conductor bus configuration is defined in the SCSI-2 standard as the A-cabled.

The SCSI-2 revision added a high-density, 50-position, D-shell connector option for the A-cable connectors. This connector is now called Alternative 1. The Alternative 2 Centronics latch-style connector remains unchanged from SCSI-1. A 68-conductor B-cable specification was added to the SCSI-2 standard to provide for 16- and 32-bit data transfers; the connector, however, had to be used in parallel with an A-cable. The industry did not widely accept the B-cable option, which has been dropped from the SCSI-3 standard.


To replace the ill-fated B-cable, a new 68-conductor P-cable was developed as part of the SCSI-3 specification. Shielded and unshielded high-density D-shell connectors are specified for both the A-cable and P-cable. The shielded high-density connectors use a squeeze-to-release latch rather than the wire latch used on the Centronics-style connectors. Active termination for single-ended buses is specified, providing a high level of signal integrity.

SCSI Cable and Connector Pinouts. The following section details the pinouts of the various SCSI cables and connectors. There are two electrically different version of SCSI, Single Ended and Differential. These two versions are electrically incompatible, and must not be interconnected or damage will result. Fortunately, there are very few Differential SCSI applications available in the PC industry, so you will rarely (if ever) encounter it. Within each electrical type (Single Ended or Differential), there are basically three SCSI cable types

- A-Cable (Standard SCSI)
- P-Cable (16- and 32-bit Wide SCSI)
- Q-Cable (32-bit Wide SCSI)

The A-Cable is used in most SCSI-1 and SCSI-2 installations, and is the most common cable you will encounter. SCSI-2 Wide (16-bit) applications use a P-Cable instead, which completely replaces the A-Cable. You can intermix standard and wide SCSI devices on a single SCSI bus by interconnecting A- and P-Cables with special adapters. 32-bit wide SCSI-3 applications use both the P- and Q-Cables in parallel to each 32-bit device. Today there are virtually no PC applications for 32-bit Wide SCSI-3, and because of the two cable requirement, it is not likely to catch on.

The A-Cables can have Pin Header (Internal) type connectors or External Shielded connectors, each with a different pinout. The P- and Q-Cables feature the same connector pinout on either Internal or External cable connections.



Chapter 15—Hard Disk Interfaces

The following tables show all the possible interface, cable, and connector pinout specifications. A hyphen preceding a signal name indicates the signal is Active Low. The RESERVED lines have continuity from one end of the SCSI bus to the other. In an A-Cable bus, the RESERVED lines should be left open in SCSI devices (but may be connected to ground), and are connected to ground in the bus terminator assemblies. In the P- and Q-Cables, the RESERVED lines are left open in SCSI devices as well as in the bus terminator assemblies.

Single-Ended SCSI Cables and Connectors. The single-ended electrical interface is the most popular type for PC systems. The following tables show all the possible single-ended cable and connector pinouts. The A-Cable is available in both internal unshielded as well as external shielded configurations. Both are shown in tables 15.13 and 15.14.

Signal Name	Pin	Pin	Signal Name
GROUND	1	2	-DB(0)
GROUND	3	4	-DB(1)
GROUND	5	6	-DB(2)
GROUND	7	8	-DB(3)
GROUND	9	10	-DB(4)
GROUND	11	12	-DB(5)
GROUND	13	14	-DB(6)
GROUND	15	16	-DB(7)
GROUND	17	18	-DB(Parity)
GROUND	19	20	GROUND
GROUND	21	22	GROUND
RESERVED	23	24	RESERVED
Open	25	26	TERMPWR
RESERVED	27	28	RESERVED
GROUND	29	30	GROUND
GROUND	31	32	-ATN
GROUND	33	34	GROUND
GROUND	35	36	-BSY
GROUND	37	38	-ACK
GROUND	39	40	-RST
GROUND	41	42	-MSG
GROUND	43	44	-SEL
GROUND	45	46	-C/D
GROUND	47	48	-REQ
GROUND	49	50	-I/O

Signal Name	Pin	Pin	Signal Name
GROUND	1	26	-DB(0)
GROUND	2	27	-DB(1)
GROUND	3	28	-DB(2)
GROUND	4	29	-DB(3)
GROUND	5	30	-DB(4)
GROUND	6	31	-DB(5)
GROUND	7	32	-DB(6)
GROUND	8	33	-DB(7)
GROUND	9	34	-DB(Parity)
GROUND	10	35	GROUND
GROUND	11	36	GROUND
RESERVED	12	37	RESERVED
Open	13	38	TERMPWR
RESERVED	14	39	RESERVED
GROUND	15	40	GROUND
GROUND	16	41	-ATN
GROUND	17	42	GROUND
GROUND	18	43	-BSY
GROUND	19	44	-ACK
GROUND	20	45	-RST
GROUND	21	46	-MSG
GROUND	22	47	-SEL
GROUND	23	48	-C/D
GROUND	24	49	-REQ
GROUND	25	50	-I/O

IBM has standardized on the SCSI interface for virtually all PS/2 systems introduced since 1990. These systems use a Micro Channel SCSI adapter or have the SCSI Host Adapter built into the motherboard. In either case, IBM's SCSI interface uses a special 60-pin mini-Centronics type external shielded connector that is unique in the industry. A special IBM cable is required to adapt this connector to the standard 50-pin Centronics style connector used on most external SCSI devices. The pinout of the IBM 60-pin mini-Centronics style External Shielded connector is shown in the following table. In the following table, notice that although the pin arrangement is unique, the pin number to signal designations correspond with the standard unshielded internal pin header type of SCSI connector.

The P-Cable (Single-Ended) and connectors are used in 16-bit wide SCSI-2 applications (see table 15.15 for the pinout).

Chapter 15—Hard Disk Interfaces

Signal Name	Pin	Pin	Signal Name
GROUND	1	60	Not Connected
-DB(0)	2	59	Not Connected
GROUND	3	58	Not Connected
-DB(1)	4	57	Not Connected
GROUND	5	56	Not Connected
-DB(2)	6	55	Not Connected
GROUND	7	54	Not Connected
-DB(3)	8	53	Not Connected
GROUND	9	52	Not Connected
-DB(4)	10	51	GROUND
GROUND	11	50	-I/O
-DB(5)	12	49	GROUND
GROUND	13	48	-REQ
-DB(6)	14	47	GROUND
GROUND	15	46	-C/D
-DB(7)	16	45	GROUND
GROUND	17	44	-SEL
-DB(Parity)	18	43	GROUND
GROUND	19	42	-MSG
GROUND	20	41	GROUND
GROUND	21	40	-RST
GROUND	22	39	GROUND
RESERVED	23	38	-ACK
RESERVED	24	37	GROUND
Open	25	36	-BSY
TERMPWR	26	35	GROUND
RESERVED	27	34	GROUND
RESERVED	28	33	GROUND
GROUND	29	32	-ATN
GROUND	30	31	GROUND

Signal Name	Pin	Pin	Signal Name
GROUND	1	35	-DB(12)
GROUND	2	36	-DB(13)
GROUND	3	37	-DB(14)
GROUND	4	38	-DB(15)
GROUND	5	39	-DB(Parity 1)
GROUND	6	40	-DB(0)



Signal Name	Pin	Pin	Signal Name
GROUND	7	41	-DB(1)
GROUND	8	42	-DB(2)
GROUND	9	43	-DB(3)
GROUND	10	44	-DB(4)
GROUND	11	45	-DB(5)
GROUND	12	46	-DB(6)
GROUND	13	47	-DB(7)
GROUND	14	48	-DB(Parity 0)
GROUND	15	49	GROUND
GROUND	16	50	GROUND
TERMPWR	17	51	TERMPWR
TERMPWR	18	52	TERMPWR
RESERVED	19	53	RESERVED
GROUND	20	54	GROUND
GROUND	21	55	-ATN
GROUND	22	56	GROUND
GROUND	23	57	-BSY
GROUND	24	58	-ACK
GROUND	25	59	-RST
GROUND	26	60	-MSC
GROUND	27	61	-SEL
GROUND	28	62	-C/D
GROUND	29	63	-REQ
GROUND	30	64	-I/O
GROUND	31	65	-DB(8)
GROUND	32	66	-DB(9)
GROUND	33	67	-DB(10)
GROUND	34	68	-DB(11)

The Q-Cable (Single-Ended) and connector is defined only for 32-bit SCSI implementations, which also require a P-Cable (see table 15.16 for the pinout). 32-bit SCSI applications are rare to virtually nonexistent.

Signal Name	Pin	Pin	Signal Name
GROUND	1	35	-DB(28)
GROUND	2	36	-DB(29)
GROUND	3	37	-DB(30)
GROUND	4	38	-DB(31)
GROUND	5	39	-DB(Parity 3)


(continues)

Chapter 15—Hard Disk Interfaces

Signal Name	Pin	Pin	Signal Name
GROUND	6	40	-DB(16)
GROUND	7	41	-DB(17)
GROUND	8	42	-DB(18)
GROUND	9	43	-DB(19)
GROUND	10	44	-DB(20)
GROUND	11	45	-DB(21)
GROUND	12	46	-DB(22)
GROUND	13	47	-DB(23)
GROUND	14	48	-DB(Parity 2)
GROUND	15	49	GROUND
GROUND	16	50	GROUND
TERMPWRQ	17	51	TERMPWRQ
TERMPWRQ	18	52	TERMPWRQ
RESERVED	19	53	RESERVED
GROUND	20	54	GROUND
GROUND	21	55	TERMINATED
GROUND	22	56	GROUND
GROUND	23	57	TERMINATED
GROUND	24	58	-ACKQ
GROUND	25	59	TERMINATED
GROUND	26	60	TERMINATED
GROUND	27	61	TERMINATED
GROUND	28	62	TERMINATED
GROUND	29	63	-REQQ
GROUND	30	64	TERMINATED
GROUND	31	65	-DB(24)
GROUND	32	66	-DB(25)
GROUND	33	67	-DB(26)
GROUND	34	68	-DB(27)

Differential SCSI Signals. Differential SCSI is not normally used in a PC environment but is very popular with minicomputer installations due to the very long bus lengths that are allowed. Although not popular in PC systems, the interface connector specifications are shown here for reference.

The A-Cable (Differential) Connector is available in both internal unshielded form as well as an external shielded form. Table 15.18 shows the pinout for the Internal cable, while table 15.19 shows the pinout for the External cable.



Signal Name	Pin	Pin	Signal Name
GROUND	1	2	GROUND
+DB(0)	3	4	-DB(0)
+DB(1)	5	6	-DB(1)
+DB(2)	7	8	-DB(2)
+DB(3)	9	10	-DB(3)
+DB(4)	11	12	-DB(4)
+DB(5)	13	14	-DB(5)
+DB(6)	15	16	-DB(6)
+DB(7)	17	18	-DB(7)
+DB(Parity)	19	20	-DP(Parity)
DIFFSENS	21	22	GROUND
RESERVED	23	24	RESERVED
TERMPWR	25	26	TERMPWR
RESERVED	27	28	RESERVED
+ATN	29	30	-ATN
GROUND	31	32	GROUND
+BSY	33	34	-BSY
+ACK	35	36	-ACK
+RST	37	38	-RST
+MSG	39	40	-MSG
+SEL	41	42	-SEL
+C/D	43	44	-C/D
+REQ	45	46	-REQ
+I/O	47	48	-I/O
GROUND	49	50	GROUND

Signal Name	Pin	Pin	Signal Name
GROUND	1	26	GROUND
-DB(0)	2	27	-DB(0)
+DB(1)	3	28	-DB(1)
-DB(2)	4	29	-DB(2)
+DB(3)	5	30	-DB(3)
+DB(4)	6	31	-DB(4)
+DB(5)	7	32	-DB(5)
+DB(6)	8	33	-DB(6)
+DB(7)	9	34	-DB(7)

(continues)

Chapter 15—Hard Disk Interfaces

Signal Name	Pin	Pin	Signal Name
+DB(Parity)	10	35	-DP(Parity)
DIFFSENS	11	36	GROUND
RESERVED	12	37	RESERVED
TERMPWR	13	38	TERMPWR
RESERVED	14	39	RESERVED
+ATN	15	40	-ATN
GROUND	16	41	GROUND
+BSY	17	42	-BSY
+ACK	18	43	-ACK
+RST	18	44	-RST
+MSG	20	45	-MSG
+SEL	21	46	-SEL
+C/D	22	47	-C/D
+REQ	23	48	-REQ
+I/O	24	49	-I/O
GROUND	25	50	GROUND

The P-Cable (Differential) and Connector is used for 16-bit wide SCSI connections. Table 15 20 has the pinouts for the P-Cable (Differential)

Signal Name	Pin	Pin	Signal Name
+DB(12)	1	35	-DB(12)
+DB(13)	2	36	-DB(13)
+DB(14)	3	37	-DB(14)
+DB(15)	4	38	-DB(15)
-DB(Parity 1)	5	39	-DB(Parity 1)
GROUND	6	40	GROUND
+DB(0)	7	41	-DB(0)
+DB(1)	8	42	-DB(1)
+DB(2)	9	43	-DB(2)
+DB(3)	10	44	-DP(3)
+DB(4)	11	45	-DB(4)
+DB(5)	12	46	-DB(5)
+DB(6)	13	47	-DB(6)
+DB(7)	14	48	-DB(7)
+DB(Parity 0)	15	49	-DB(Parity 0)
DIFFSENS	16	50	GROUND
TERMPWR	17	51	TERMPWR



Signal Name	Pin	Pin	Signal Name
TERMPWR	18	52	TERMPWR
RESERVED	19	53	RESERVED
+ATN	20	54	-ATN
GROUND	21	55	GROUND
+BSY	22	56	-BSY
+ACK	23	57	-ACK
+RST	24	58	-RST
+MSG	25	59	-MSG
+SEL	26	60	-SEL
+C/D	27	61	-C/D
+REQ	28	62	-REQ
+I/O	29	63	-I/O
GROUND	30	64	GROUND
+DB(8)	31	65	-DB(8)
+DB(9)	32	66	-DB(9)
+DB(10)	33	67	-DB(10)
+DB(11)	34	68	-DB(11)

The Q Cable (Differential) and connector is used only with the proposed 32-bit wide SCSI implementations (which have not been implemented by the marketplace as of yet), and in that case would also require a 16-bit wide P-cable. Table 15.21 shows the Q-Cable (Differential) pinout.

Signal Name	Pin	Pin	Signal Name
+DB(28)	1	35	-DB(28)
+DB(29)	2	36	-DB(29)
+DB(30)	3	37	-DB(30)
+DB(31)	4	38	-DB(31)
+DB(Parity 3)	5	39	-DB(Parity 3)
GROUND	6	40	GROUND
-DB(16)	7	41	-DB(16)
+DB(17)	8	42	-DB(17)
+DB(18)	9	43	-DB(18)
+DB(19)	10	44	-DB(19)
+DB(20)	11	45	-DB(20)
+DB(21)	12	46	-DB(21)
+DB(22)	13	47	-DB(22)
+DB(23)	14	48	-DB(23)
-DB(Parity 2)	15	49	-DB(Parity 2)

(continues)

Signal Name	Pin	Pin	Signal Name
DIFFSENS	16	50	GROUND
TERMPWRQ	17	51	TERMPWRQ
TERMPWRQ	18	52	TERMPWRQ
RESERVED	19	53	RESERVED
TERMINATED	20	54	TERMINATED
GROUND	21	55	GROUND
TERMINATED	22	56	TERMINATED
+ACKQ	23	57	-ACKQ
TERMINATED	24	58	TERMINATED
TERMINATED	25	59	TERMINATED
TERMINATED	26	60	TERMINATED
TERMINATED	27	61	TERMINATED
+REQQ	28	62	-REQQ
TERMINATED	29	63	TERMINATED
GROUND	30	64	GROUND
+DB(24)	31	65	-DB(24)
+DB(25)	32	66	-DB(25)
+DB(26)	33	67	-DB(26)
+DB(27)	34	68	-DB(27)

Termination. All busses need to be electrically terminated at each end, and the SCSI bus is no exception. Improper termination is still one of the most common problems in SCSI installations. Three types of terminators typically are available for the SCSI bus:

- Passive
- Active (also called Alternative 2)
- Forced Perfect Termination (FPT)
 - FPT-3
 - FPT-18
 - FPT-27

Typical passive terminators (a network of resistors) allow signal fluctuations in relation to the terminator power signal on the bus. Usually, passive terminating resistors suffice over short distances, such as 2 or 3 feet, but for longer distances, active termination is a real advantage. Active termination is required with Fast SCSI.

An active terminator actually has one or more voltage regulators to produce the termination voltage, rather than resistor voltage dividers. This arrangement helps ensure that the SCSI signals are always terminated to the correct voltage level. The SCSI-2 specification

recommends active termination on both ends of the bus and requires active termination whenever Fast or Wide SCSI devices are used.

A variation on active termination is available: Forced Perfect Termination. Forced Perfect Termination is an even better form of active termination, in which diode clamps are added to eliminate signal overshoot and undershoot. The trick is that instead of clamping to +5 and Ground, these terminators clamp to the output of two regulated voltages. This arrangement enables the clamping diodes to eliminate signal overshoot and undershoot, especially at higher signaling speeds and over longer distances.

FPT terminators are available in several versions. FPT-3 and FPT-18 versions are available for 8-bit standard SCSI, while the FPT-27 is available for 16-bit (Wide) SCSI. The FPT-3 version forces perfect the three most highly active SCSI signals on the 8-bit SCSI bus, while the FPT-18 forces perfect all the SCSI signals on the 8-bit bus except grounds. FPT-27 also forces perfect all of the 16-bit Wide SCSI signals except grounds.

Several companies make high-quality terminators for the SCSI bus, including Aeronics and the Data Mate division of Methode. Both of these companies make a variety of terminators, but Aeronics is well noted for some unique FPT versions that are especially suited to problem configurations that require longer cable runs or higher signal integrity. One of the best investments that you can make in any SCSI installation is in high-quality cables and terminators.

SCSI Drive Configuration. SCSI drives are not too difficult to configure, especially compared with IDE drives. The SCSI standard controls the way that the drives must be set up. You need to set two or three items when you configure an SCSI drive:

- SCSI ID setting (0 through 7)
- Terminating resistors

SCSI ID Setting. The SCSI ID setting is very simple. Up to eight SCSI devices can be used on a single SCSI bus, and each device must have a unique SCSI ID address. The host adapter takes one address, and up to seven SCSI peripherals take the others. Most SCSI host adapters are factory-set to ID 7, which is the highest-priority ID. All other devices must have unique IDs that do not conflict with one another. Some host adapters boot only from a hard disk set to a specific ID. In my system, for example, the IBM SCSI host adapter requires the boot drive to be set to ID 6. Newer IBM host adapters and systems enable you to boot from a hard disk at any SCSI ID. Older Adaptec host adapters required the boot hard disk to be ID 0; newer ones can boot from any ID.

Setting the ID usually involves changing jumpers on the drive itself. If the drive is installed in an external chassis, the chassis may have an ID selector switch that is accessible at the rear. This selector makes ID selection a simple matter of pressing a button or rotating a wheel until the desired ID number appears. If no external selector is present, you must open the external device chassis and set the ID via the jumpers on the drive.

Three jumpers are required to set the SCSI ID; the particular ID selected actually is derived from the binary representation of the jumpers themselves. For example, setting all



Chapter 15—Hard Disk Interfaces

three ID jumpers off results in a binary number of 000b, which translates to an ID of 0. A binary setting of 001b equals ID 1, 010b equals 2, 011b equals 3, and so on. Notice that as I list these values, I append a lowercase b to indicate binary numbers.

Unfortunately, the jumpers can appear either forward or backward on the drive, depending on how the manufacturer set them up. To keep things simple, I have recorded all the different ID jumper settings in the following tables. Table 15.22 shows the settings for drives that order the jumpers with the Most Significant Bit (MSB) to the left; Table 15.23 shows the settings for drives that have the jumpers ordered so that the MSB is to the right.

SCSI	ID	Jumper	Settings
0	0	0	0
1	0	0	1
2	0	1	0
3	0	1	1
4	1	0	0
5	1	0	1
6	1	1	0
7	1	1	1

1 = Jumper On, 0 = Jumper Off

SCSI	ID	Jumper	Settings
0	0	0	0
1	1	0	0
2	0	1	0
3	1	1	0
4	0	0	1
5	1	0	1
6	0	1	1
7	1	1	1

1 = Jumper On, 0 = Jumper Off

Termination. SCSI termination is very simple. Termination is required at both ends of the bus; there are no exceptions. If the host adapter is at one end of the bus, it must have termination enabled. If the host adapter is in the middle of the bus, and if both internal and external bus links are present, the host adapter must have its termination disabled, and the devices at each end of the bus must have terminators installed. Several types of terminators are available, differing both in quality and in appearance. Active terminators

are the minimum recommended, and Forced Perfect Terminators (FPT) are considered the best available. For more information on the different types, see the previous section on terminators in this chapter.

The rules are simple use the best terminators possible, and make sure that only the ends of the SCSI bus are terminated. The majority of problems that I see with SCSI installations are the result of improper termination. Some devices have built-in termination resistors that are enabled or disabled through a jumper or by being physically removed. Other devices do not have built-in terminating resistors; these devices instead rely on external terminator modules for termination.

When installing an external SCSI device, you usually will find the device in a storage enclosure with both input and output SCSI connectors so that you can use the device in a daisy chain. If the enclosure is at the end of the SCSI bus, an external terminator module most likely will have to be plugged into the second (outgoing) SCSI port to provide proper termination at that end of the bus (see fig. 15.7).




Figure 15.7

External SCSI device terminator.

External terminator modules are available in a variety of connector configurations, including pass-through designs, which are needed if only one port is available. Pass-through terminators also are commonly used in internal installations in which the device does not have built-in terminating resistors. Many hard drives use pass-through terminators for internal installations to save space on the logic-board assembly (see fig. 15.8).

The pass-through models are required when a device is at the end of the bus and only one SCSI connector is available.




Figure 15.8

Internal pin-header connector pass-through SCSI terminator

Remember to stick with high-quality active or Forced Perfect terminators at each end of the bus, and you will eliminate most common termination problems.

Other Settings. Other configuration items on a SCSI drive can be set via jumpers. Following are several of the most common additional settings that you will find:

- Start on command (delayed start)
- SCSI parity
- Terminator power
- Synchronous negotiation

These configuration items are described in the following sections.

Start On Command (Delayed Start). If you have multiple drives installed in a system, it is wise to set them up so that all the drives do not start to spin immediately when the system is powered on. A hard disk drive can consume three or four times more power during the first few seconds after power-on than during normal operation. The motor requires this additional power to get the platters spinning quickly. If several drives are drawing all this power at the same time, the power supply may be overloaded, which can cause the system to hang or to have intermittent startup problems.

Nearly all SCSI drives provide a way to delay drive spinning so that this problem does not occur. When most SCSI host adapters initialize the SCSI bus, they send out a command called Start Unit to each of the ID addresses in succession. By setting a jumper on the hard disk, you can prevent the disk from spinning until it receives the Start Unit command from the host adapter. Because the host adapter sends this command to all the ID addresses in succession, from the highest-priority address (ID 7) to the lowest (ID 0),

the higher-priority drives can be made to start first, with each lower-priority drive spinning up sequentially. Because some host adapters do not send the Start Unit command, some drives may simply delay spinup for a fixed number of seconds rather than wait for a command that never will arrive.

If drives are installed in external chassis with separate power supplies, you need not implement the delayed-start function. This function is best applied to internal drives that must be run from the same power supply that runs the system. For internal installations, I recommend setting Start on Command (delayed start) even if you have only one SCSI drive; this setting will ease the load on the power supply by spinning the drive up after the rest of the system has full power. This method is especially good for portable systems and other systems in which the power supply is limited.


SCSI Parity. SCSI parity is a limited form of error checking that helps ensure that all data transfers are reliable. Virtually all host adapters support SCSI parity checking, so this option should be enabled on every device. The only reason why it exists as an option is that some older host adapters do not work with SCSI parity, so the parity must be turned off.

Terminator Power. The terminators at each end of the SCSI bus require power from at least one device on the bus. In most cases, the host adapter supplies this terminator power; in some cases, however, it does not. For example, parallel port SCSI host adapters typically do not supply terminator power. It is not a problem if more than one device supplies terminator power because each source is diode protected. For simplicity's sake, many will configure all devices to supply terminator power. If no device supplies terminator power, the bus will not be terminated correctly and will not function properly.

SCSI Synchronous Negotiation. The SCSI bus can run in two modes: asynchronous (the default) and synchronous. The bus actually switches modes during transfers through a protocol called synchronous negotiation. Before data is transferred across the SCSI bus, the sending device (called the initiator) and the receiving device (called the target) negotiate how the transfer will take place. If both devices support synchronous transfers, they will discover this fact through the negotiation, and the transfer will take place at the faster synchronous rate.

Unfortunately, some older devices do not respond to a request for synchronous transfer and actually can be disabled when such a request is made. For this reason, both host adapters and devices that support synchronous negotiation often have a jumper that can be used to disable this negotiation so that it can work with older devices. By default, all devices today should support synchronous negotiation, and this function should be enabled.

Plug and Play (PnP) SCSI. Plug and Play SCSI was originally released in April 1994. This specification allows SCSI device manufacturers to build Plug and Play peripherals that will automatically configure when used with a Plug and Play operating system. This will allow you to easily connect or reconfigure external peripherals, such as hard disk drives, backup tapes, and CD-ROMs.



Chapter 15—Hard Disk Interfaces

To connect SCSI peripherals to the host PC, the specification requires a Plug and Play SCSI host adapter such as Plug and Play ISA or PCI. Plug and Play add-in cards enable a Plug and Play operating system automatically to configure software device drivers and system resources for the host bus interface.

The Plug and Play SCSI specification version 1.0 includes these technical highlights:

- A single cable-connector configuration
- Automatic termination of the SCSI bus
- SCAM (SCSI Configured Automatically) automatic ID assignment
- Full backward compatibility of Plug and Play SCSI devices with the installed base of SCSI systems

This should go a long way in making SCSI easier to use for the normal user.

SCSI Drivers. Each SCSI peripheral that you add to your SCSI bus (other than hard disk drives) requires an external driver to make the device work. Hard disks are the exception; driver support for them normally is provided as part of the SCSI host adapter BIOS. These external drivers are specific not only to a particular device, but also to the host adapter.

Recently, two types of standard host adapter interface drivers have become popular, greatly reducing this problem. By having a standard host adapter driver to write to, peripheral makers can more quickly create new drivers that support their devices and then talk to the universal host adapter driver. This arrangement eliminates dependence on one particular type of host adapter. These primary or universal drivers link the host adapter and the operating system.

The Advanced SCSI Programming Interface (ASPI) currently is the most popular universal driver, with most peripheral makers writing their drivers to talk to ASPI. The A in ASPI used to stand for Adaptec, the company that introduced it, but other SCSI device vendors have licensed the right to use ASPI with their products. DOS does not support ASPI directly, but it does when the ASPI driver is loaded. Windows 95, NT, OS/2 2.1, and later versions provide automatic ASPI support for several SCSI host adapters.

Future Domain and NCR have created another interface driver called the Common Access Method (CAM). CAM is an ANSI-approved protocol that enables a single driver to control several host adapters. In addition to ASPI, OS/2 2.1 and later versions currently offer support for CAM. Future Domain also provides a CAM-to-ASPI converter in the utilities that go with its host adapters.

SCSI Configuration Tips. When you are installing a chain of devices on a single SCSI bus, the installation can get complicated very quickly. Here are some tips for getting your setup to function quickly and efficiently:

- *Start by adding one device at a time.* Rather than plug numerous peripherals into a single SCSI card and then try to configure them at the same time, start by installing the host adapter and a single hard disk. Then you can continue installing devices one at a time, checking to make sure that everything works before moving on.



- *Keep good documentation.* When you add a SCSI peripheral, write down the SCSI ID address as well as any other switch and jumper settings, such as SCSI Parity, Terminator Power, and Remote Start. For the host adapter, record the BIOS addresses, Interrupt, DMA channel, and I/O Port addresses used by the adapter, as well as any other jumper or configuration settings (such as termination) that might be important to know later.
- *Use proper termination.* Each end of the bus must be terminated, preferably with active or Forced Perfect Terminators (FPT). If you are using any Fast SCSI-2 device, you must use active terminators rather than the cheaper passive types. Even with standard (slow) SCSI devices, active termination is highly recommended. If you have only internal or external devices on the bus, the host adapter and last device on the chain should be terminated. If you have external and internal devices on the chain, you generally will terminate the first and last of these devices but not the SCSI host adapter itself (which is in the middle of the bus).
- *Use high-quality shielded SCSI cables.* Make sure that your cable connectors match your devices. Use high-quality shielded cables, and observe the SCSI bus-length limitations. Use cables designed for SCSI use, and if possible, stick to the same brand of cable throughout a single SCSI bus. Different brands of cables have different impedance values; this situation sometimes causes problems, especially in long or high-speed SCSI implementations.

Following these simple tips will help minimize problems and leave you with a trouble-free SCSI installation.

IDE versus SCSI

When you compare the performance and capabilities of IDE (Integrated Drive Electronics) and SCSI (Small Computer System Interface) interfaced drives, you need to consider several factors. These two types of drives are the most popular drives used in PC systems today, and a single manufacturer may make identical drives in both interfaces. Deciding which drive type is best for your system is a difficult decision that depends on many factors.

In most cases, you will find that an IDE drive outperforms an equivalent SCSI drive at a given task or benchmark and that IDE drives usually cost less than SCSI drives, thus offering better value. In some cases, however, SCSI drives have significant performance and value advantages over IDE drives.

Performance. ATA (AT Attachment) IDE drives currently are used in most PC configurations on the market today, because the cost of an IDE-drive implementation is low and the performance capabilities are high. In comparing any given IDE and SCSI drive for performance, you have to look at the capabilities of the HDAs (Head Disk Assemblies) that are involved.

To minimize the variables in this type of comparison, it is easiest to compare IDE and SCSI drives from the same manufacturer that also use the identical HDA. You will find that in most cases, a drive manufacturer makes a given drive available in both IDE and

Chapter 15—Hard Disk Interfaces

SCSI forms. For example, Seagate makes the ST-3600A (ATA-IDE) and ST-3600N (Fast SCSI-2) drives, both of which use identical HDAs and which differ only in the logic board. The IDE version has a logic board with a built-in disk controller and a direct AT bus interface. The SCSI version has the same built-in disk controller and bus interface circuits, and also a SCSI Bus Interface Controller (SBIC) chip. The SBIC chip is a SCSI adapter that places the drive on the SCSI bus. What you will find, in essence, is that virtually all SCSI drives actually are IDE drives with the SBIC chip added.

The HDAs in these example drives are capable of transferring data at a sustained rate of 2.38M to 4M per second. Because the SCSI version always has the additional overhead of the SCSI bus to go through, in almost all cases the directly attached IDE version performs faster.

SCSI versus IDE: Advantages and Limitations. IDE drives have much less command overhead for a given sector transfer than SCSI drives do. In addition to the drive-to-controller command overhead that both IDE and SCSI must perform, a SCSI transfer involves negotiating for the SCSI bus; selecting the target drive; requesting data; terminating the transfer over the bus; and finally converting the logical data addresses to the required cylinder, head, and sector addresses.

This arrangement gives IDE an advantage in sequential transfers handled by a single-tasking operating system. In a multitasking system that can take advantage of the extra intelligence of the SCSI bus, SCSI can have the performance advantage.

SCSI drives offer significant architectural advantages over IDE and other drives. Because each SCSI drive has its own embedded disk controller that can function independently of the system CPU, the computer can issue simultaneous commands to every drive in the system. Each drive can store these commands in a queue and then perform the commands simultaneously with other drives in the system. The data could be fully buffered on the drive and transferred at high speed over the shared SCSI bus when a time slot was available.

Although IDE drives also have their own controllers, they do not operate simultaneously, and command queuing is not supported. In effect, the dual controllers in a dual-drive IDE installation work one at a time so as not to step on each other.

Although SCSI drives require an additional-cost host adapter card, more and more PCs require tape-backup, CD-ROM, or optical-drive support and thus must still be configured with a SCSI host bus adapter. This means that the incremental cost of supporting SCSI drives is virtually nil, because the SCSI host bus adapter is shared with other devices, such as tape and optical drives. In addition, all major operating systems today include software support for a wide range of SCSI devices.

What are the limitations of IDE?

- IDE does not support overlapped, multitasked I/O.
- IDE does not support command queuing
- IDE does not support bus mastering.

As you can see, SCSI has some advantages over IDE, especially where expansion is concerned, and also with regard to support for multitasking operating systems. Unfortunately it also costs much more to implement.

Recommended Aftermarket Controllers and Host Adapters

Many companies manufacture disk controllers for IBM and IBM-compatible systems. Many newer systems include IDE drives, which have built-in controllers and offer a high level of performance at a low cost. Other systems are using SCSI drives because of the inherent flexibility of the SCSI bus in supporting many drives and other peripherals.

I recommend IDE drives for most standard installations because the connections are simple and the drives are inexpensive for the power. For higher-end systems, or for systems in which upgradability and flexibility are most important, I recommend SCSI drives.

Silicon Valley makes a line of IDE adapters that are excellent for systems that do not have the special IDE connector on the motherboard. The company also makes a special adapter that enables you to put AT IDE drives in 8-bit XT systems, which few other cards do. Arco Electronics and Sigma Data make IDE adapters that enable you to install ATA IDE drives in PS/2 MCA bus systems. These companies even have unique versions for the Model 50z and 70 that replace the interposer card, therefore using no MCA slots in the system. All these companies are listed in Appendix B.

Recommended SCSI Host Adapters. For SCSI host adapters, I normally recommend Adaptec. Their adapters work well and come with the necessary formatting and operating software. Windows 95, Windows NT, and OS/2 have built-in support for Adaptec SCSI adapters. This support is a consideration in many cases because it frees you from having to deal with additional drivers.

Standard or Fast SCSI is adequately supported by the ISA bus, but if you are going to install a Fast-Wide SCSI bus, or especially an Ultra-Wide bus, then you should consider some form of local bus SCSI adapter, normally PCI. This is because the ISA bus supports a maximum transfer speed of about 8MB/sec, while a Fast-Wide SCSI bus runs up to 20MB/sec, and an Ultra-Wide SCSI bus runs up to a blazing 40MB/sec! In most cases a local bus SCSI adapter would be a PCI bus version, which is supported in most current PC systems.

One example of a popular SCSI adapter for the ISA bus is the Adaptec AHA-1540C and 1542C. The 1542 has a built-in floppy controller; the 1540 does not. These adapters are most notable for their ease of installation and use. Virtually all functions on the card can be configured and set through software. No more digging through manuals, looking for Interrupt, DMA, I/O Port, and other jumper settings—everything is controlled by software and saved in a flash memory module on the card. Following are some of the features of this card:

- Complete configuration utility built into the adapter's ROM
- Software-configurable IRQ, ROM Addresses, DMA, I/O Port Addresses, SCSI Parity, SCSI ID, and other settings
- Software-selectable termination (no resistors to pull out!)

Chapter 15—Hard Disk Interfaces

- Enhanced BIOS support for up to eight 7.88G drives
- No drivers required for more than two hard disks
- Drive spinup on a per-drive basis available
- Boots from any SCSI ID

The 1542C model includes an Intel 82077 floppy controller chip (which supports 2 88M drives), 2.88M BIOS support also must be provided by your motherboard BIOS (see fig. 15.9).

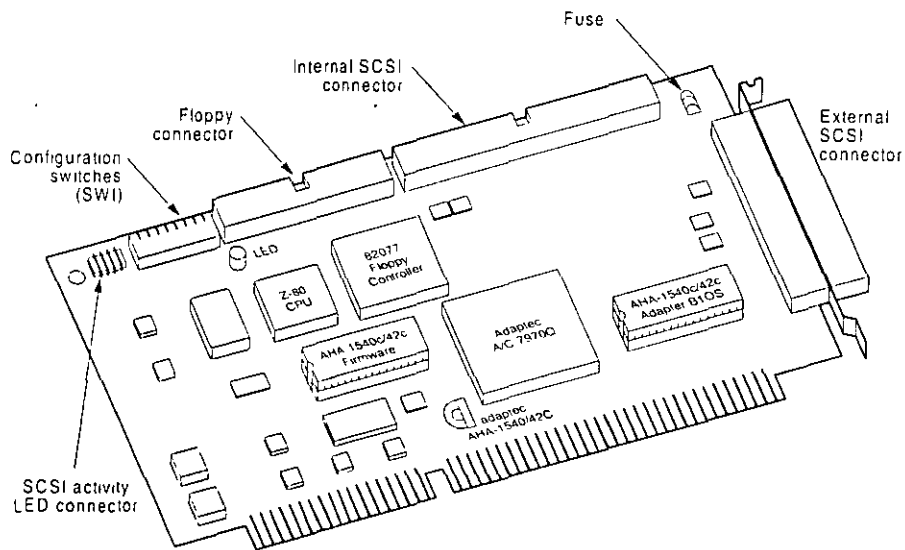


Figure 15.9

An Adaptec AHA-1542C SCSI host adapter.

More recently, Adaptec has released full Plug and Play versions of their SCSI adapters. These adapters will be automatically configured in any PC that supports the Plug and Play (PnP) specification, or they can be configured manually through supplied software in non-PnP systems. The PnP SCSI adapters are highly recommended because they can be configured without opening up the PC! All functions are set by software, and there are no jumpers or switches to attend to. Most peripheral manufacturers write drivers for Adaptec's cards first, so you will not have many compatibility or driver-support problems with any Adaptec card.

If you are using the older AHA-1542B SCSI host adapter, the current ROM revision available is version 3.20, which adds Enhanced BIOS support for multiple drives up to 7.88G with no drivers.

IBM SCSI Adapters. IBM has shipped several types of SCSI adapters since it introduced PC SCSI in 1990. Three basic types of adapters are available:

- 16-bit SCSI MCA adapter
- 32-bit SCSI MCA adapter with cache
- 16-bit Fast SCSI ISA adapter

The IBM host adapters (especially the MCA versions) are among the easiest to install and configure in the industry. Much of this ease comes from the Micro Channel architecture, which eliminates the need to worry about interrupts, DMA channels, I/O Port addresses, and even BIOS addresses. By virtue of MCA, these adapters are virtually self-configuring.

The only real problem is drivers. All these adapters have built-in support for hard disk drives, but no drivers are included for any other devices. This is important to know if you decide to attach external peripherals to these cards because you will need drivers to make the peripherals work.

IBM makes DOS drivers available only for the peripherals that it sells. Amazingly, however, IBM OS/2 includes SCSI drivers not only for numerous IBM and non-IBM peripherals, but also for non-IBM SCSI host adapters. If you cannot use OS/2, one way around the driver-support problem is to use the same peripherals that IBM sells, in which case you can get the drivers from IBM. For example, IBM sells a CD-ROM drive that actually incorporates a Toshiba mechanism. If I purchase a Toshiba CD-ROM drive, I can make it work simply by using the drivers that IBM supplies with its (Toshiba) CD-ROM drive.

IBM makes its drivers available free for the downloading on the IBM NSC BBS (the phone number is in Appendix B). I am also using a Hewlett Packard DAT (digital audiotape) drive for backup because that is the same mechanism that IBM sells in its own DAT drive. Again, I get the drivers from IBM. This system does not always work, however. For example, I want to attach a Hewlett Packard scanner to my IBM host adapter, but neither HP nor IBM (or anybody else so far) makes drivers to run the scanner with the IBM card. My only choice would be to use a different SCSI card to run the scanner—not a very good alternative, in my estimation.

The 16- and 32-bit MCA adapters have been available in two versions. The original version of the 16-bit (noncached) MCA adapter was known as FRU (Field Replacement Unit) # 15F6561 and was recalled for a defect. The replacement was labeled FRU # 85f0002. Some of the earlier cards were reworked and returned to the field without having the new FRU # applied over the old one. These repaired adapters can be identified by two yellow wires running from the upper to the lower side of the module (the module is labeled ZM10). If you have the original version of this card, getting the replacement is important (IBM should replace it free) because the original card could cause random lockups and data corruption.

The original 32-bit cached MCA adapter did not have a serious problem that required replacement, but it was updated over the years with improved functionality. The original SCSI adapter with cache does not support drive capacities greater than 1G. If you install a drive larger than 1G, the drive will work, but only the first 1G will be recognized. A new version of this card has an improved BIOS that supports drives up to 8G.

You can tell which SCSI adapter with cache you have by inspecting the adapter. The new adapter comes with an orange terminating resistor on the end of the card, similar to the noncached adapter. If you do not have both internal and external devices installed, this orange resistor should be located on the adapter. If you do have both internal and external devices, the resistor should be removed; there will be an empty socket on the adapter where the resistor originally was located. The new SCSI adapter with cache is sold under the part number 6451133 and known as FRU # 85F0063. If you have the original SCSI adapter with cache, you will not have either an orange resistor or an empty socket where the resistor is supposed to go. The original adapter does not implement termination on the card itself.

It is not widely known that you can upgrade the cache on the IBM SCSI adapter with cache from the standard 512K to 2MB simply by replacing the 256K (9-bit) SIMMs with 1MB SIMMs. The only caveat is that you must use the IBM-style 9-bit SIMMs, which have a slightly different pinout from so-called industry-standard 9-bit SIMMs. I have done this to my IBM SCSI card and have achieved perhaps a 10 percent improvement in throughput.

IBM also sells a new Fast/Wide MCA SCSI-2 adapter under part number-G451280. IBM also has a Fast SCSI-2 adapter for ISA bus systems, which actually is made by Future Domain. This adapter is sold for use in ISA bus systems such as PS/1 and PS/Valuepoints. Because of Future Domain support, many more drivers are available for this adapter than for the MCA adapters, making it very easy to add a variety of peripherals under DOS as well as OS/2. This adapter also supports Fast SCSI transfers, whereas the original MCA IBM adapters do not.

Disk Hardware and Software Limitations

By studying the capabilities of the different disk interfaces as well as the ROM BIOS and operating systems, it is possible to determine the limits on disk storage. The following section details the limits under the different interfaces and operating systems.

Disk Interface Capacity Limitations

Different disk interfaces have different limitations on the theoretical maximum drive capacities that they may support. These limitations are due to variations in the way that each interface operates at the hardware level. It is important to note that even though a particular interface may permit access to a given amount of disk real estate, the BIOS and DOS usually are much more limiting and end up being the true limits for system disk capacity.

ST-506/412, ESDI, and IDE. To determine the capacity limits for the ST-506/412, ESDI, or IDE interface, you first need to determine the limits on the maximum number of cylinders, heads, and sectors per track. To do so, look at the size of the registers that hold this data in the controller. All these interfaces have the same controller register specifications, so the capacity limits calculated here apply to all of them. As you will see, the interface capacity limits are quite high. The drive parameter limits are as follows:

Cylinders (16 bits) = 65,536

Heads (4 bits) = 16

Sectors (8 bits) = 256

This calculates to a maximum theoretical drive size of:

$65,536 \text{ Cyls} \times 16 \text{ Hds} \times 256 \text{ Secs} \times 512 \text{ Bytes} = 137,438,953,472 \text{ Bytes (128G)}$

Unfortunately, the maximum capacity—128G—is limited by the BIOS. There are two different BIOS types with regards to disk size limitations. The standard BIOS built into most systems is limited to 1,024 cylinders, 16 heads, and 63 sectors per track. If the BIOS is an enhanced version, it will be limited to 1,024 cylinders, 256 heads, and 63 sectors per track. Combining the BIOS and interface limits results in the following maximum capacities (assuming 512-byte sectors):

Limit w/ Std. BIOS: $1,024 \text{ Cyls} \times 16 \text{ Hds} \times 63 \text{ Secs} = 528,482,304 \text{ Bytes (504M)}$

Limit w/ Enh. BIOS: $1,024 \text{ Cyls} \times 256 \text{ Hds} \times 63 \text{ Secs} = 8,455,716,864 \text{ Bytes (7.88G)}$

If you do not have enhanced BIOS support on your motherboard, you could add an IDE bus adapter that has an on-board enhanced BIOS. To get around such BIOS problems, some IDE drive implementations over 528 million bytes split the drive to act as two physical units. In this case, the drive would appear on the IDE bus connector as being both master and slave, and could be used only as two 504M-maximum-size drives.

ATA-2 has defined LBA support for EIDE. This is a logical block address mode where each sector on the drive is numbered from 0 to x. The limitations are that x is a 28-bit number that has a maximum value of 268,435,456. Using 512 byte sectors, this brings the maximum drive capacity to 137,438,953,472. That is coincidentally the same as the IDE internal limit has always been. The LBA is translated by the ATA-2/EIDE enhanced BIOS to the extended Cylinder, Head, and Sector (CHS) parameters which allow a maximum of 1,024 cylinders, 256 heads, and 63 sectors respectively.

SCSI. According to the SCSI specification, drives are not addressed by cylinders, heads, and sectors but instead by what is called a Logical Block Address (LBA). This is a sector number in which all the sectors on the drive are numbered sequentially from start to finish. The LBA is specified by a 32-bit number with 512-byte sectors, results in the following limitation:

$4,294,967,296 \text{ LBAs (sectors)} \times 512 \text{ Bytes} = 2,199,023,255,552 \text{ Bytes (2,048G or 2 terabytes)}$

As you can see, SCSI drive capacity limits are extremely high. However, because the SCSI drive must appear to the BIOS as being a given number of cylinders, heads, and sectors per track, the BIOS limits SCSI capacity. Virtually all SCSI adapters have an enhanced BIOS that supports a maximum drive capacity as follows (assuming 512-byte sectors):

SCSI w/ Enh. BIOS: $1,024 \text{ Cyls} \times 256 \text{ Hds} \times 63 \text{ Secs} = 8,455,716,864 \text{ Bytes (7.88G)}$

If you do not have enhanced BIOS support in your SCSI adapter or motherboard, in some cases, you can load an external driver for your adapter to provide this support.

Most systems support up to 4 SCSI host adapters, each with up to 7 hard disk drives for a total of 28 physically installed drives.

ROM BIOS Capacity Limitations

In addition to the capacity limit of 504M, the standard ROM BIOS is limited to supporting only two hard disk drives. The enhanced BIOS is limited to 128 drives maximum. Most SCSI and IDE adapters get around the two-drive standard BIOS limits by incorporating an enhanced BIOS on board that takes over the disk interface. Some of the newer adapter on-board BIOS versions support booting from CD-ROM drives as well.

Operating-System Capacity Limitations

IBM and Microsoft officially say that DOS 5 and later versions will support up to eight physical hard disks. IBM says that OS/2 1.30.1 and later versions (including 2.x) support up to 24 physical hard disks, and because OS/2 includes DOS, that implies that DOS under OS/2 would support 24 physical drives as well. OS/2 HPFS (High Performance File System) also supports a maximum partition size of 8G and a maximum single-file size of 2G, whereas DOS and OS/2 FAT partitions have a maximum size of 2G and a maximum single-file size of 2G also. As you have seen, BIOS limitations currently limit the maximum physical hard disk size to about 7.88G (or about 8.46 million bytes).

Summary

This chapter examined hard disk interfaces and controllers. You learned about the different types of interfaces, their benefits and limitations. You also learned how to select and configure these interfaces. You can use this information to upgrade an existing system with a new type of hard disk or to troubleshoot and repair an existing system.