Chapter 1

Introduction to S/36 Architecture

Computers, like skyscrapers, are built from the basement up; sometimes it takes years to reach the top floor. Most existing commercial computers have risen slowly from established architectures, and the System/36 is no exception. It stands on the foundation wrought by the IBM System/3, System/32, and System/34. But System/36 users should not fear that their trusted systems will become lost in the rubble. An investigation of the history and architecture of the System/36 reveals that System/36 reliability and cost efficiency ensure several more years of life, despite the availability of newer systems such as the AS/400.

Examining the S/36's inner parts reveals its evolution. Its unique parallel processors and modular components speak of several design phases. There is much to learn. But interesting as this knowledge might be, of what practical use is it? Why should anyone be interested in how a machine works internally? Certainly the S/36 can be programmed and operated without such detailed knowledge, just as a car can be driven without understanding what goes on under the hood. However, for the driver, a knowledge of auto mechanics comes in handy when comparison shopping for a new car or when "tuning up" the old car. Knowing something about the inner workings of your computer provides the same kind of benefit.

Understanding what happens under the computer's covers is especially helpful when planning for the inevitable upgrade to a newer platform: you can take steps today to make your applications more portable while retaining good performance. And programmers who understand the basic operation of a computer can design programs that take advantage of the computer architecture's strong points.

Genealogy

A discussion of present-day S/36 architecture must begin with a look at the architecture of its predecessors. Figure 1.1 summarizes the hardware developments contributed by each generation in the S/36 ancestry.

The S/36 architectural line began in 1969 with the S/3, one of the first computers for small businesses. The S/3 was revolutionary in one sense: it offered a fast commercial instruction set at a time when most computers relied on a scientific instruction set. A computer's instruction set consists of the lowest level machine language orders it can carry out. "Add binary," "load register," and "branch to address" exemplify machine language instructions. The scientific machine instruction set used by most computers of that era

Figure 1.1
System/36 Genealogy

Machine	Major architectural contributions
System/3	• Commercial machine language instruction set, including: multibyte, memory-to-memory instructions (up to 256 bytes per instruction); variable-length decimal; and arithmetic no hardware multiply/divide • Address Translation Register (ATR) technique for multitask memory management.
System/32	• Separate processor to handle I/O (Control Storage Processor) • Microprogram emulation of System/3 processor (Main Storage Processor) • Single program execution only • 27 MB disk drive technology (Gulliver)
System/34	• Hardwired implementation of System/3 processor • Task switching and address translation handled by CSP • Significant operating system functions implemented as Supervisor Calls • Scientific instruction set (for BASIC and FORTRAN) in microcode • Four-line communication controller (MLCA) to handle polling and teleprocessing I/O • 64 MB disk drive technology (Picollo)
System/36	• Faster main storage and control storage processors • 8 MB real main storage addressability • 192 K translated (region) main storage addressability • 128 K control storage (double that of the System/34) • Two-byte wide, bidirectional channel • Improved overlap of CSP and MSP operation • Fast task-switching hardware and multiple ATR groups • Virtual memory management for system transient routines • New MSP instructions for ease of programming • W/S controller optional to reduce cost of entry level machines • Word Processing Text Mode support for workstations • Data Storage Controller to handle disk/tape/diskette data transfers • Eight-line communication controller (ELCA) to handle polling and TP I/O • RLL/ECC disk drive technology to eliminate write verifies

performed calculations on binary numbers in registers (high-speed, scratch-pad memory). For scientific programs this type of instruction set provided much needed speed. Programs computing Laplace transforms or orbital velocities ran like the wind.

Business programs, however, have little to do with binary numbers or registers; they work instead with decimal numbers and variable length fields. A traditional scientific instruction set was ill-equipped to deal with these factors because it was constantly converting commercial formats to binary formats and vice versa. Consequently, payroll and inventory applications run on

computers with a scientific instruction set used up much processing time on data conversion and ran like glaciers.

In contrast to many of its contemporaries, the S/3 commercial instruction set could perform arithmetic operations directly on decimal numbers in memory — no intermediary registers were necessary. It also could, in a single instruction, manipulate a field of data up to 256 characters long (no more byte-by-byte translation to binary). With the S/3's multi-byte, memory-to-memory instructions, S/3 programs did not spend much time converting data, and business application performance improved.

The speed of the S/3 processor was carried over to the S/32 (sometimes called the "bionic desk" because it was an all-in-one unit), which was introduced in 1973. With the S/32, IBM heralded a new hardware technology that allowed two CPUs to function side-by-side in one machine.

One CPU, or processor, was a microprogrammed version of the S/3. This processor wasn't as fast as a S/3, but it executed the same instruction set, which meant IBM could reuse much previously developed system software. The other processor, given the job of handling all contact with peripheral devices and the outside world, was a real innovation. This second processor ran its own dedicated program using its own dedicated memory. Although in the early 1970s the S/32's dual processors represented a hardware advancement, this system was limited by its capacity to run only one program at a time.

Four years later, IBM announced the S/34, which continued the S/32 philosophy of using two processors. However, the microprogrammed processor was replaced by a much faster "hardwired" version. This processor, now faster than the fastest S/3, could run multiple programs simultaneously. The second processor, also a holdover from the S/32, took on the job of managing memory and dispatching tasks for the hardwired processor. (IBM included in this processor a scientific instruction set, emulated by microcode, which allowed faster execution of BASIC and FORTRAN programs.) Local terminal management was relegated to yet a third processor. A fourth processor was available as an option to support outside communications — the first time a processor was offered optionally. Extensive use was made of another development, the ATOM (A Tiny Optimized Microprocessor), to directly control the system printer and MICR (Magnetic Ink Character Recognition) devices.

These developments led to the S/36, which appeared in 1983. In many ways, it is a radical departure from the S/34. For example, memory addressing was re-engineered for the first time since the S/3. The two main processors also were improved and the selection of optional processors expanded.

Internally, the S/36 supports several different processors. Figure 1.2 summarizes their names, functions, and characteristics. The Main Storage Processor (MSP) is really a hardwired S/3 CPU with a few new instructions. The Control Storage Processor (CSP) controls the overall operation of the entire

Figure 1.2
System/36 Multiple Processors

Function of Processor	Internal Technology	Instruction Set	Execution Speed in MIPS	Address Space	Memory Access Time (ns)
Main Storage Processor	MSP	Enhanced System/3	0.36	8 MB	200
Control Storage Processor	CSP	Register-to-Register	1.6	128 K	200
Workstation Controller	CSP/I	Register-to-Register	1.6	128 K	200
Data Storage Controller	CSP/I	Register-to-Register	1.6	128 K	200
Eight-Line Comm Adapter	CSP/I	Register-to-Register	1.6	128 K	200
Printer Controller	ATOM	Register-to-Register	1.1	128 K	200
Magnetic Ink Character Recognition	ATOM	Register-to-Register	1.1	128 K	200
Local Area Network	PC/AT	Intel 80286	1.2	640 K	150

machine. It runs a dedicated program in its own memory (control storage) which may be either 64 K or 128 K. A third kind of processor has the same instruction set and organization as the CSP, but it is used as a dedicated controller for certain input/output (I/O) operations. IBM has designated it the CSP/I; one each is found in the Workstation Controller, Multi-Line Communications Adapter (MLCA), Eight-Line Communications Adapter (ELCA), and Data Storage Controller. If a 3262 printer or Magnetic Character Reader is attached to the system, ATOMs will control these devices. To better understand the function of each of these processors, let's first examine the two main processors (MSP and CSP) and then discuss the optional processors.

Main Storage Processor

The S/36 Main Storage Processor (MSP) runs SSP programs and user applications through the S/3-based commercial (memory-to-memory) instruction set. That is all the MSP does. It has no control over which programs are executed. It has no direct contact with the outside world. (When the MSP must perform I/O operations, it submits a request to the CSP, which handles contact with

the outside world.) And the MSP executes only .36 million instructions per second (MIPS). This rate might seem slow when compared with other computers, but because the MSP doesn't concern itself with I/O or task management, it is free to concentrate on the job at hand. This freedom makes up for the MSP's apparent lack of horsepower.

The MSP in the current S/36 can address up to 8 MB of memory, or main storage — seemingly small by today's standards, where an average PS/2 might be configured with up to 16 MB. However, the S/36's implementation of virtual memory (VM) lets you run up to 128 MB of applications simultaneously. Chapter 2 describes S/36 memory architecture, including VM, in detail.

Control Storage Processor

The Control Storage Processor (CSP) interfaces with peripheral devices, manages MSP memory and swapping, and controls the execution of the MSP. The CSP also provides special computational services to the MSP, including high-level operating system operations such as queue management and intertask communications. Through judicious task and memory management, the CSP tries to keep the MSP operating at maximum efficiency. Because the CSP is not working on business programs, it uses a more applicable register-to-register instruction set, which allows the CSP, running at speeds of 1.3 to 1.9 MIPS, to juggle many jobs at once. The services provided by the CSP simplify the programming involved in the SSP and take advantage of the four-fold speed advantage CSP has over MSP for time-critical functions.

There are three versions of the S/36 CSP. Machines shipped before October 1984 contain a Stage 1 CSP, which runs at 1.3 MIPS. Machines shipped after October 1984, including all 5362 and 5364 processors, contain a Stage 2 CSP, running at 1.6 MIPS. CPUs on the 5360 model D, and all 5363s, use a Stage 3 CSP, running at 1.9 MIPS. On small machines the performance difference between Stage 1 and Stage 2 processors is insignificant because the CSP is rarely running at anywhere near its rated capacity. However, on large 5360s running many workstations or DisplayWrite/36 jobs, the CSP may be fully utilized, and the Stage 3 processor improves performance significantly.

The S/36 CSP has a number of enhancements over the S/34 version. In addition to being faster, the S/36 CSP processes more requests in parallel with the MSP than did its predecessor. It also recognizes many new Supervisor Call (SVC) instructions, which perform operating system functions for the MSP. Included in these new SVC instructions is a "storage mapping" service, which allows SSP programs (e.g., data management) easier access to buffers in a user application.

The S/36 CSP also contains a larger control storage area than the S/34 CSP. Because the S/34 CSP contained only 64 K of control storage, control storage programs that couldn't fit in this space were read in from disk as "transients"

when required. The S/36 CSP can contain either 64 K or 128 K of control storage. (The extra storage is used to contain Workstation Controller (WSC) code if the WSC function is inboard, or to simply keep more CSP routines resident if the WSC function is outboard.) In addition to the extra control storage, the S/36 CSP offers a new Virtual Address Facility in its memory management function. This facility allows any number of MSP transients to run from the user area instead of bottle-necking in a single transient area as they did on the S/34.

An interesting and useful service provided by the CSP is the Alter/Display facility. When the MSP STOP button is pressed on the service panel, a special menu appears at the system console. This menu allows a programmer or service technician to examine and modify any location on disk, in main storage, or in control storage. This kind of tool, when it is available at all on single-processor machines, is usually implemented as a large and complex control panel. The S/36's "soft" control panel is much easier to use and provides a wider range of functions. For example, an Address Compare Stop feature can be used to stop the MSP when a certain disk or memory address is referenced or changed to a specified value. The MSP is in a suspended state while Alter/Display is being used; processing resumes at the point of interruption after exiting the Alter/Display menu. This capability is invaluable for tracking down difficult system bugs.

Optional Processors

As options, you may install other processors that take care of additional tasks. A Workstation Controller (WSC) processor deals with local workstation input/output; a Data Storage Controller (DSC) processor mediates data transfers between disk and slower devices such as diskette and tape; an MLCA (Multi-Line Communications Adapter) or ELCA (Eight-line Communications Adapter) processor handles polling and protocol for multiple communication lines; a Local Area Network processor supports IBM's Token-Ring LAN.

The Workstation Controller. An interesting difference between the S/34 and the S/36 is in the workstation controller. Every S/34 had a dedicated CSP/I with 32 K of control storage to poll workstations, process keystrokes and handle field attributes like right-adjust, zero fill, and check digits. The workstation expansion feature to support more than eight devices was simply a memory expansion of the WSC to 64 K.

Not every S/36, however, has a dedicated WSC. On all 5364 models, and on 5362 models without the workstation expansion feature, the WSC function is performed "inboard" by the CSP. Because the CSP and WSC both use identical processors, adding WSC tasks to the CSP's workload wasn't hard to do, and it allowed IBM to produce machines with full S/36 functionality at a lower price. (For a comparison of features among S/36 models, see Figure 1.3.) These models do not suffer a performance loss because the CSP has

Figure 1.3
Comparison of Features Among S/36 Models

Model	Disk Configurations	Access Time	Memory Capacity	Workstation Controller	Data Storage Controller	Comm Controller	Diskette	Tape**
5364	40,80 MB up to 2 spindles	60 ms	1 MB	N/A	N/A	PC (1 line)	5.25"	6157
5363	68 to 425 MB up to 2 spindles	35 ms	2 MB	N/A	Optional	SLCA (2 lines)	5.25"	6157
5362	30 to 520 MB up to 4 spindles	35 ms	2 MB	Optional	N/A	SLCA, MLCA (4 lines)	8"	6157
5360	30 to 1438 MB up to 4 spindles	35 ms	7 MB	Dedicated (optional 2nd avail)	Optional with tape attachment	SLCA, MLCA, ELCA (8 lines)	8"	6157 8809
9402*	160 to 640 MB up to 4 spindles	35 ms	2 MB	N/A	Optional	SLCA (2 lines)	5.25"	6157

* Also known as the AS/400 model 9402, but actually is a S/36.
IBM marketed this machine for a short time as the AS/Entry.
** The 6157 tape drive is a streaming cartridge unit with 60 MB capacity.
The 8809 tape drive is a reel-to-reel unit for mainframe data exchange.

enough additional capacity to easily take on the extra load. Larger 5362s and all 5360s have an "outboard" WSC that relieves the CSP of handling more extensive local networks.

Both the inboard and outboard WSC implementations support the new "word processing mode" for local workstations. This mode adds such functions as indentation, margin control, tab entry, and word wrap — functions used by DisplayWrite/36 to provide a user interface better adapted for word processing than the fixed-field format of data processing mode. Because these features are under the direct control of a CSP or CSP/I, they have consistently fast response time, regardless of the load on the MSP.

Remote devices, such as the 5251 Model 12 workstation and the 5294 control unit, also contain workstation controllers. In the 5251 Model 12, the WSC program is fixed in Read Only Memory (ROM) and cannot be changed. It is unable to support word processing mode. The 5294 does not contain a fixed WSC program. Instead, the host S/36 downloads WSC microcode when the 5294 goes on-line. Thus, word processing mode functions are available, unlike remote WSC that have fixed micro code.

Data Storage Controller. The S/36 supports up to two tape drives —

something unavailable on the S/34. Tape drives are attached to the system through the Data Storage Controller (DSC), which can autonomously transfer files from disk to tape without the intervention of either the MSP or CSP. In fact, the DSC also can mediate transfers between disk and diskette, diskette and tape, or disk and disk.

When a S/36 does not have a DSC, data is transferred between devices on an internal two-byte-wide path called the "channel." This same path is used for intercommunication between the MSP, CSP, and other processors. With all these devices competing for use of the channel, a sudden high-volume transfer of data can result in a logjam of information, degrading system performance significantly. The DSC operates "below" the channel, communicating directly with the devices over its own private data path. This capability reduces access contention on the main channel and eliminates the degradation that normally occurs with large file transfers. The S/34 experienced tremendous response-time degradation when transferring files between diskette and disk, or disk and disk.

To operate efficiently, the DSC contains two 16 K buffers. It initially fills both buffers; then, after one buffer is written to the output device, it starts refilling it while the second buffer is being written. This double-buffering improves the output transfer rate significantly and allows the tape drive to run in streaming (high speed, nonstop) mode. When a DSC transfer is requested, the CSP notifies the DSC where the source and destination files are and the DSC takes over, interrupting the CSP only when a diskette or tape must be changed.

Because the DSC only relieves congestion on the main channel (it doesn't actually move the data faster), no appreciable performance improvement will be noticed unless the system is heavily loaded. The DSC can only make response times more consistent. The DSC also is limited to performing one device-to-device transfer at a time. If the DSC is engaged in a transfer and the MSP requests another transfer, the second request will be queued until the DSC is free. The one exception to this rule is if the DSC is processing a tape transfer and a request for a diskette transfer is made. Because the tape transfer could require a long time (especially if the operator doesn't change reels when prompted), the diskette request is processed immediately using the main I/O channel.

MLCA and ELCA Processors: A S/36 with one or two communications lines (Single-Line Communications Adapter — SLCA) uses the CSP to poll the lines, handle bottom-layer protocol, and buffer data transfers. One line presents no problem, but two lines can put an unwieldy burden on the CSP, which is forced to drop everything it's doing to service the high-priority communication interrupts. When more than two lines are installed, the MLCA (now available only on the 5362) and ELCA processors do the dirty work. These processors are essentially identical — the ELCA is more recent and the

only product currently available on newer 5360s.

Because both communications processors are a dedicated CSP/I, they support data rates much faster than the SLCA. They also assume the responsibility for polling terminals, processing protocol messages, computing checksums, retransmitting buffers, and for the lower layers of SDLC protocol. A machine with MLCA or ELCA installed will experience much less degradation than a machine using SLCA.

Local Area Network Processor. The S/36 supports IBM's Token-Ring Local Area Network (LAN) through a specially attached PC/AT. The Token-Ring network ports appear as communications lines 9 and 10, and the Token Ring runs only at the 4 million bits per second (mbps) data rate; the 16 mbps Token Ring isn't supported.

IBM's Multiple Processor Advantage

A computer architecture, such as the S/36's, that uses multiple processors faces a significant problem: how to use the processors efficiently. The goal is to get some degree of parallel operation without unnecessarily holding up the execution of any one task. Traditional approaches to the problem treat all processors equally either by running separate application programs on each processor or by interleaving instruction execution among processors. IBM has taken a different tack with the S/36 by assigning each processor a specific, dedicated job and by designing for each a unique instruction set optimized for the tasks at hand. Figure 1.4 diagrams the major components inside the S/36 and shows how they are interconnected.

The usefulness of the multiprocessor architecture is demonstrated in the analogy that single-processor computers suffer from the same problems as single-engine airplanes: a shortage of options. When the engine quits on a single-engine airplane, there are no options from which to choose. The important decision has been made for you by the engine: the aircraft is going down. The engine in a single-processor machine can stop, too, when an invalid instruction is encountered, or when a hardware error occurs. When such an event happens, the computer often is headed in the same direction as the airplane.

The multiprocessor S/36, like a multi-engine aircraft, recovers somewhat more gracefully from serious system failures. If the MSP tries to do something crazy, the CSP gets control and executes an error-recovery procedure. Often, the error-recovery program needs only to cancel the offending task before resuming the work in progress. Sometimes even this step is not necessary because the problem can be corrected while the MSP waits. For example, if the MSP runs into a parity check (memory failure) in a main storage memory card, the task is canceled and the 2 K page of memory is taken off-line to protect other tasks from the damaged memory. Likewise, if a disk sector is found to be unusable, the CSP automatically assigns a spare from a special supply of

Figure 1.4
Overview of S/36 Internal Components

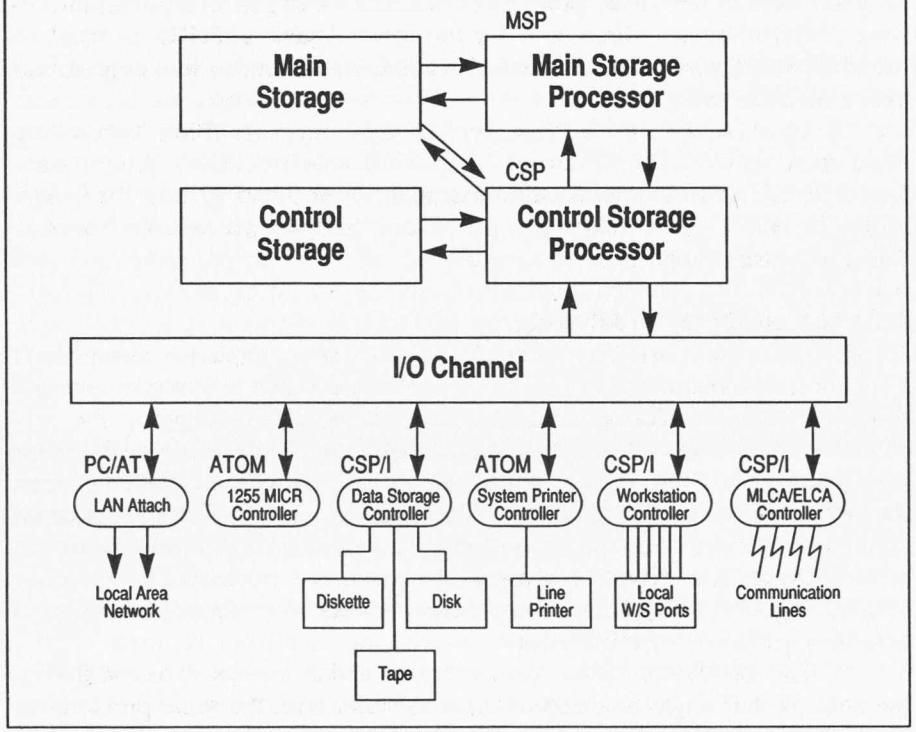

extra disk sectors, then lets the MSP proceed as if nothing happened. The CSP also keeps a detailed log — the Error Recovery Analysis Report (ERAP) — of any problems it detects for later perusal by a customer engineer.

The interdependence between the CSP and the MSP is especially important because during a typical processing day several programs compete simultaneously for use of the MSP. Competition for the MSP means the CSP must make many decisions about when to run which program. The process of allowing a program (task) to run, then stopping it and starting up another program, is called "task-switching."

A typical task-switching scenario might proceed as follows: when the MSP must perform some I/O operation, it makes a request to the CSP via a Supervisor Call instruction. The program that requested the I/O must now wait. The CSP selects another program that is ready to run and starts it, then schedules the I/O operation for the first task, thus "switching" the tasks. The steps that take place when a task switch occurs bear examining because a

major advantage of the S/36 over single-processor systems hinges on how these steps are carried out.

Many computers, including the S/36, use an "I/O-driven" mechanism for switching tasks. That is, when the execution of one task is interrupted to perform an input/output operation, the machine switches to another task. This switch makes sense because most I/O operations are quite slow when compared with the speed of the processor. For example, a disk read requires about 40 milliseconds; in the same amount of time the MSP could execute nearly 15,000 instructions. Because the task that requested the I/O operation is going to wait anyway, running another task in its stead overlaps the operation of the processor and I/O.

However, task switching is not an instantaneous event. There is a general sequence of events that must occur. First, the I/O operations that caused the interruption must be dealt with: transferring the data and controlling the device. Then the computer must determine which of several tasks should run next and maintain the various queues used to make this decision. After a new task has been selected, the environment of the old task (instruction pointer and registers) must be stored. If the new task has been swapped out to disk, it must be brought into memory. Finally, the environment for the new task must be loaded and execution started at the point of previous interruption.

Figure 1.5a shows the timeline of events for a conventional single-processor computer. Because only one processor is available, when an interrupt occurs, everything else must stop while the task switch is done. If two tasks are run together, chances are that the total amount of time to run them will be longer than if the tasks were run one after the other: more time is required to switch between tasks. When many tasks are running at the same time, task switching can account for an appreciable portion of the total execution time. In fact, as the task load increases, a point eventually will be reached where more time is used up doing task switches than running the tasks themselves. To the user, it appears as though system performance degrades rapidly, out of proportion to the number of tasks. This situation is clearly unhealthy; in most forms of accounting it is referred to as a net loss.

The same timeline for a S/36 is shown in Figure 1.5b. Here, when a task switch must be made, only the CSP is interrupted — the MSP keeps running. The CSP then sets up everything for the task switch. It takes care of I/O handling, determines which task will run next, swaps the task into storage if necessary, and then switches tasks. All time-consuming operations are performed in parallel by the CSP while the MSP continues to process user program instructions. However, the S/36 contains special "fast task-switch" hardware that allows it to save and load the MSP registers quickly, which, in turn, makes the task switch nearly instantaneous.

The MSP is not involved in the details of shifting gears and loses little

Figure 1.5a
Timeline for Single-Processor Task Switching

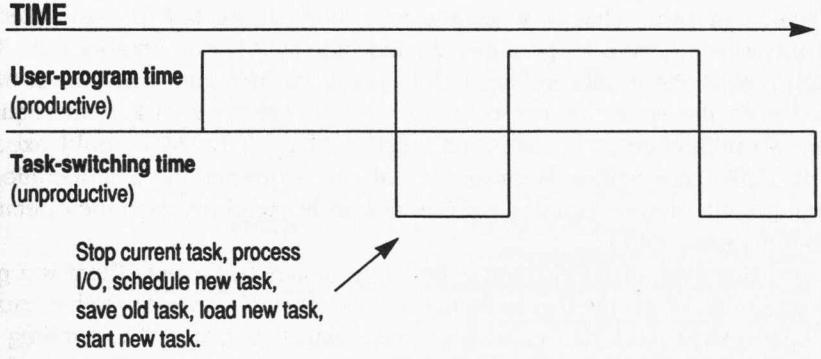

Figure 1.5b
Timeline for S/36 Dual-Processor Task Switching

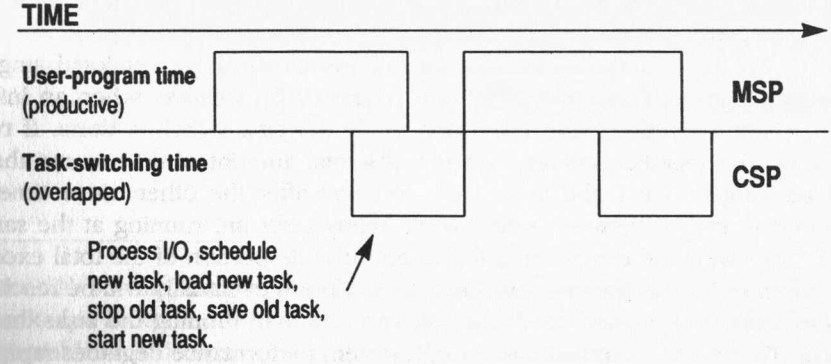

time between tasks. Within the space of a few MSP instructions, the old task is stopped and the new task is started. As new tasks are added, the MSP concentrates on running those tasks, and there are fewer abrupt changes in system performance.

The Channel

The I/O Channel is the data path used by the MSP, CSP, other processors, and peripheral devices to transfer data inside the machine. When a disk record is read it is transferred byte by byte down the channel to main storage. The MSP and CSP continue to run during this transfer, but the channel "steals" a cycle

from the MSP when it needs to access main storage. Other transfers work in a similar manner. The S/34 channel was one byte wide, meaning that, for each cycle, one byte could be transferred from one component to another inside the machine. The S/36 channel is bidirectional and two bytes wide. It can transfer two bytes at a time between components or one byte simultaneously in each direction. Because the S/36 channel can transfer twice as much data as the S/34 version, it "steals" only half as many cycles from the MSP, which improves the performance of user programs and SSP functions. Generally, any channel activity takes only half as long as it did on the S/34, greatly reducing internal traffic congestion.

The S/36 channel is actually an intelligent device, not just a data bus. It incorporates a primitive channel processor that executes a limited range of instructions specifically geared to moving data on the bus. Although these instructions are simple, it gives the channel some degree of autonomy: the CSP can issue commands to the channel and then go do other work while the commands are carried out.

Disk Drives

The hard disk technology used on the S/36 is a major advance over previous devices. For this discussion, only the 10SR 200 MB drive will be examined. But other disks (the 30/60 MB used in the 5362 and the 40/80 MB used in the 5364) are similar in operation. All S/36 disks use data encoding to increase reliability and decrease access time.

To better understand S/36 disk drives, let's again contrast the S/36 with the S/34. The S/34 Gulliver (27 MB) and Picollo (64 MB) devices required that, after every write operation, the data be re-read to make sure it was recorded correctly. While this step was handled automatically by the CSP, it was time-consuming: after writing a record, the CSP had to wait for the disk to spin around again to the starting point before the record could be re-read for verification. Thus, write operations were more than twice as long as read operations. Also, while the S/34 re-read technique provided a high level of reliability when the data was written, it provided no recourse if the data was damaged after writing (random damage). Experience with S/34 drives demonstrated that the most common random disk error was a single-bit failure within a byte. Double-bit errors within a byte also occurred but much less frequently.

The 10SR (STAR) uses a data encoding technique called Run Length Limited (RLL) encoding, which eliminates the re-read requirement and achieves reliability by detecting and correcting single- and double-bit errors at read time. The S/36 technique recognizes that bad data could be put on the disk at write time, but that most problems will be single-bit errors. On the S/36, data is not written on disk as a series of fixed-length bytes, as it is on the S/34. Instead, the bytes are encoded into variable-length bit strings containing

Figure 1.6
Run Length Limited Codes for Disk Error Correction

Input bit string	Resulting RLL code
10	0100
11	1000
000	000100
011	001000
010	100100
0010	00100100
0011	00001000

twice as many bits as originally input (Figure 1.6). Six Error Correcting Code bytes are also written for each 256-byte sector. On a S/36, then, because information is being stored redundantly, it is often possible to repair damaged bytes at read time. When a record is read, the encoded data is decoded and an error detection/correction algorithm executed on the result. The mathematics of the algorithm guarantee that any single or double-bit errors can be detected, and that single-bit errors can be corrected. The net effect is that records can be safely written to disk without re-reading for verification.

The Bottom Line

The S/36 is a prime example of building on pre-existing technology effectively. Through extensions to an established architecture, it has the ability to coexist in a distributed environment with other IBM midrange systems. The S/36's modular and general-purpose internal components let you effectively trade off performance, capacity, and cost. The architecture of the S/36 proves that IBM has acted to preserve the history of engineering, software development, manufacturing knowledge, and technical support invested in the S/36.

Where to Learn More

An excellent overview of the System/36 can be found in the IBM technical bulletin *S/36 Internals* (G361009). It outlines general concepts of both software and hardware architectures, provides a lucid explanation of memory addressing, and presents details about MSP/CSP interfaces. This volume is actually one of a series of six "Rochester Technical Bulletins" — the other five cover specific SSP topics:

- S/36 8809 Tape Support (G360-1005)
- *S/36 Performance Monitoring and Tuning* (G360-1006)
- *S/36 Query/36 Design Guide* (G360-1007)
- *S/36 Data Dictionary System Design Guide* (G360-1008)
- *S/36 Advanced Disk Data Management* (G360-1008)

Assembler language programmers will find the following two IBM volumes useful: *Programming with Assembler* (SC21-7908) and *Functions Reference Manual* (SA21-9436). The first book is provided as part of the IBM Basic Assembler Language program product and covers everything a programmer would need to know to write simple assembler programs. A more complete description of the machine, from the programmer's perspective, is found in the *Functions Reference Manual.* Machine addressing modes, instruction formats, and supervisor calls are examined in excruciating detail. The programming characteristics of every device (disk, diskette, tape, printer, display, and communications) also are set forth. Programmers who intend to write special subroutines that access input/output devices directly will be interested in this level of detail.

Technical references useful to system programmers are contained in the trilogy:

- *S/36 Program Service Information* (LY21-0590)
- *S/36 System Data Areas* (LY21-0592)
- *S/36 Program Problem Diagnosis and Diagnostic Aids* (LY21-0593)

These books are available for a charge to any licensed user of SSP. Those who plan to do serious programming in assembler language should have these manuals; they cover debugging facilities, SSP component operation, memory, and disk organization, and the formats of internal SSP data areas.

The manual, *IBM S/36 Control Storage Service Information* (LY31-0650), describes the detailed operation of control storage processor programs. It explains how the CSP communicates with and controls the MSP. The concepts are well illustrated, and an appendix contains several step-by-step examples of CSP/MSP interaction.

For hardcore hardware details, turn to the *S/36 Theory of Operation* manual, which covers detailed internal computer operations at a circuit board level. This manual is one of the large-format customer engineering books shipped with every 5360 system unit. Smaller versions of the S/36 (the 5362 and 5364) are not supplied with this manual.