

M. S. Pittler
D. M. Powers
D. L. Schnabel

System Development and Technology Aspects of the IBM 3081 Processor Complex

The IBM 3081 Processor Complex consists of a 3081 Processor Unit and supporting units for processor control, power, and cooling. The Processor Unit, completely implemented in LSI technology, has a dyadic organization of two central processors, each with a 26-ns machine cycle time, and executes System/370 instructions at approximately twice the rate of the IBM 3033. This paper presents an overview of the advances in technology and in the design process, as well as enhancements in the system design that were associated with the development of the IBM 3081. Application of LSI technology to the design of the 3081 Processor Unit, which contains almost 800 000 logic circuits, required extensions to silicon device packaging, interconnection, and cooling technologies. A key achievement in the 3081 is the incorporation of the thermal conduction module (TCM), which contains as many as 45 000 logic circuits, provides a cooling capacity of up to 300 W, and allows the elimination of one complete level of packaging—the card level. Reliability and serviceability objectives required new approaches to error detection and fault isolation. Innovations in system packaging and organization, and extensive design verification by software and hardware models, led to the realization of the increased performance characteristics of the system.

Introduction

The trend in large computer systems has been toward higher performance and improved reliability, availability, and serviceability (RAS). A significant factor influencing such progress has been the remarkable advance in electronics technology, which has evolved from discrete devices to encapsulated transistors to hybrid integrated circuits and, more recently, to large-scale integration. The trend of improvement in large-processor performance over the period 1965–1981 at IBM is illustrated in Fig. 1.

There were two basic product goals for the development of the IBM 3081: better price/performance levels and compatible growth for System/370 and the 303X product line. Specific objectives were set relative to IBM's previous most powerful processor, the IBM 3033: These included improved RAS characteristics, increased I/O channel capability, and an internal performance that would approximately double the instruction-execution rate.

Some of the key technical problems that were solved and significant achievements that were realized in the development of the 3081 Processor Unit are the following:

- Implementation in LSI technology of a total of about 800 000 circuits, using a maximum of 704 logic circuits per chip, to provide the performance and RAS characteristics required.
- Elimination of one complete level of packaging—the card level.
- Removal of heat at a rate beyond the capabilities of conventional air cooling.
- Development of analysis routines which run concurrently with customer operation and automatically isolate the failing replaceable unit for a high percentage of failures.
- Development of computer aids for extensive design verification.
- Realization of a 26-ns machine cycle.

The 3081 marks the beginning of an era in IBM's application of LSI technology to large computer systems. In this special issue we shall discuss the advances that have been made in LSI technology and the challenges of applying this technology to the development of a large-systems processor. This introductory paper will highlight some of the more

Copyright 1982 by International Business Machines Corporation. Copying is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract may be used without further permission in computer-based and other information-service systems. Permission to *republish* other excerpts should be obtained from the Editor.

significant technical challenges and design decisions and will serve as an overview of the other papers in the issue. Papers in this issue which discuss the 3081 refer to the Model D.

The remainder of this paper is organized into four sections. The first section, "3081 Processor Complex," provides a brief description of the processor complex, the characteristics of the processor and its packaging, and the processor controller. This is followed by a section on "3081 Technology" which describes the key technical decisions made in circuit selection and chip design and discusses the challenges and achievements in packaging, interconnections, and cooling. The third section, "Reliability, Availability, and Serviceability (RAS)," reviews the significant highlights of the design for improving processor error detection, fault isolation, and hardware error recovery. The last section, "Computer Support," discusses the extent to which computers were utilized and the effects on design and manufacturing, and presents a summary of the design verification methodology and procedures used.

3081 Processor Complex

The 3081 Processor Complex consists of a number of units, including a 3081 Processor Unit, a 3082 Processor Controller, a 3087 Coolant Distribution Unit, and a 3089 Power Unit (or other appropriate power source). A 3278 Display Console is used as a system console. An additional console device is required (as an operator console) for communication with the system control program.

The use of LSI technology in the design of the 3081 Processor Unit has resulted in a hardware system which requires less space as well as lower power and cooling requirements. Table 1 lists comparative physical characteristics of the 3081 processor complex and the 3033 processor complex with 16 channels. The space requirement for each system includes recommended service clearances.

• 3081 Processor Unit characteristics

Application of LSI technology requires a balance between packaging design and system design in order to achieve an optimal utilization of the package and to satisfy product and development goals. Two key decisions were made in the development of the 3081 Processor Unit which affected both packaging and system design. The first was to apply LSI technology to the complete processor; the second was to meet system performance objectives by reducing the processor cycle time and by the use of two *central processor* components which prepare and execute instructions in parallel.

The two central processors are an integral part of the 3081 Processor Unit. Each central processor has access to channels and to central storage via a single *system controller*. The central processors share main (central) storage (16, 24, or 32M bytes) and operate under a single control program.




Figure 1 IBM large-processor performance trend (relative to System/360 Model 65).

Table 1 Physical comparisons of the 3081 and the 3033 processor complexes.

	Space (m ²)	Cooling (kiloBTUs per hour)	Power (kilowatts)
3081-D16	34	60	23
3033-U16	45	199	68

Figure 2 illustrates the organization of the components of the 3081 Processor Unit. This dyadic organization provides a single-system image as well as the performance benefits of a multiprocessing organization.

Maximum use is made of microcode controls with minimum hardware assists. Part of the microcode is in a static control storage (writable for engineering changes), and the remaining microcode is obtained from a pageable control storage which is loaded from the central storage.

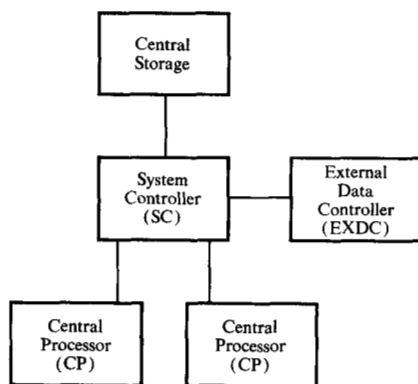


Figure 2 Organization of components of the IBM 3081 Processor Unit.

The combined effects of the organization and the 26-ns cycle result in an instruction execution rate for the 3081 Processor Unit which is generally in the range of 1.9 to 2.1 times that of the 3033 uniprocessor with identical programs and configuration, each running under the Multiple Virtual Storage/System Extensions (MVS/SE) program product. The higher performance is achieved when running batch-oriented jobs streams, both commercial and scientific. The lower performance is achieved when running an interactive workload. The difference is due to the nature of interactive transactions, which result in a reduced probability of finding data in either 32K-byte high-speed buffer (each central processor has one high-speed buffer).

The 3081 is capable of supporting up to 24 channels, which can be of either the byte-multiplexer or block-multiplexer type. A maximum of four byte-multiplexer channels can be configured within the 3081 Processor Complex. All block-multiplexer channels have a data-streaming capability which permits data rates up to 3 megabytes per second per channel. Channels are assignable to two channel sets, one set for each central processor, with a maximum of 16 channels to a set.

Channels are controlled by the *external data controller* (EXDC), which is an integrated I/O processor within the 3081 Processor Unit. The EXDC consists of two types of microcode-controlled elements; one element handles the control of I/O instructions and interrupts; the other handles the data control sequencing and provides buffering for each group of eight channels.

● *Processor packaging*

A key packaging design objective was the minimization of interconnections between packaging levels to reduce the processor cycle time and to improve the interconnection reliability. To meet this objective, one of the levels of

packaging (the card level) used in predecessor IBM large-scale processors was eliminated, leaving the, chip, module, and board levels. The packaging characteristics are described in more detail in the subsequent section, "3081 Technology."

System design objectives included the packaging of each major component of the machine within one board and the packaging of each functional element of a component within one module in order to reduce the number of interconnections between packaging levels and to reduce the cycle time.

A close working relationship between packaging and system designers and a number of iterative evaluations led to the conclusion that a field-replaceable unit containing an average of about 30 000 logic circuits would satisfy the objectives. This resulted in the design of the *thermal conduction module* (TCM), a new electronics package which can contain up to either 100 or 118 chips.

Coupled with the development of the TCM, the development of new large boards (called "TCM Boards" or "3081 Boards"), containing up to either six or nine TCMs, made it possible to package the entire processor unit on four boards in one frame. Figure 3 illustrates an assembly of a nine-module board with TCMs mounted on it.

Each of the two central processors is packaged on one nine-module board (one module position is left as a spare). One six-module board packages the system controller, and the other six-module board packages the *external data controller*, which provides for up to 24 channels.

The system design process and methodology used to organize and package the 3081 Processor Unit and to design the central processor and the external data controller are discussed more fully in the paper by Gustafson and Sparacio [1].

The system controller is more complex than the storage controllers that have been used in previous IBM large processors, mainly because data is stored directly into the high-speed buffer associated with each central processor and because of the integration of dual central processors. The functions of the system controller, which manages numerous concurrent operations, are also addressed in the paper by Gustafson and Sparacio [1].

● *Processor controller*

A significant unit in the 3081 Processor Complex is the 3082 Processor Controller, which is dedicated to monitoring, controlling, and maintaining the processor complex. Protection against power and cooling malfunctions is provided for the TCMs through the continuous monitoring of over 250

conditions which are used to control power regulators, fans, and pumps. Potential problems and malfunctions are reported to the 3081 system console and to a service support console located in the processor controller.

Diagnostics, which are used to isolate failures to the field-replaceable-unit (FRU) level in the electronics portion of the processor, are run concurrently with normal processor operation. In most cases, the failing unit is automatically identified and reported without the assistance of maintenance personnel. The processor controller is also capable of communicating with a remote service site for more detailed analysis by service specialists, who can also control and manipulate the 3081 remotely.

System operation is simplified because of the role of the processor controller in system initialization and reconfiguration. A part of the initialization process is the dynamic location of an error-free main-storage location for the assignment of the *initial program load* (IPL) of the system control program. An additional error-free main-storage area is identified for use by the processor controller for hardware support purposes. A dedicated path to main storage via the system controller allows communication (concurrent with system activity) between the processor controller and the system control program, the central processors, and the channels.

The 3081 implements two new instructions of the System/370 architecture to facilitate communication between the processor controller and the processor unit and system control program. These are MSFCALL (an extension to the DIAGNOSE instruction) and TEST BLOCK.

The processor controller is packaged in a frame separate from the processor unit and uses card-on-board logic technology similar to that used in preceding large IBM processors. The paper by Reilly *et al.* [2] provides a more detailed description of this unit of the 3081 Processor Complex.

3081 Technology

Objectives for reliability and performance of the 3081 technology led to key decisions in circuit selection and chip design and presented challenges in packaging, cooling, and interconnection technologies. These considerations are discussed in this section.

• Circuits and chips

The choice of the type of integrated logic circuit appropriate for the design of a processing system is dependent upon such considerations as inherent circuit switching speed, power requirements, circuit function, achievable manufacturing density, and process complexity. In addition, the circuit type selected should be applicable to a number of different




Figure 3 3081 nine-module board.




Figure 4 Example of TTL gate configuration.

processor product designs, each with a different circuit speed and power requirement. Commonalities in design are also sought, in order to minimize differences in manufacturing processes and to reduce design automation programming efforts.

For these reasons, the basic logic chips used in the 3081 Processor Unit contain similar but higher-speed versions of the *transistor-transistor logic* (TTL) circuits used in the IBM System/38 (announced in 1978) and in the IBM 4300 processor (announced in 1979). By lowering resistor values and optimizing the device size and critical spacings on the gate-array logic-chip layout, circuit delays were reduced to

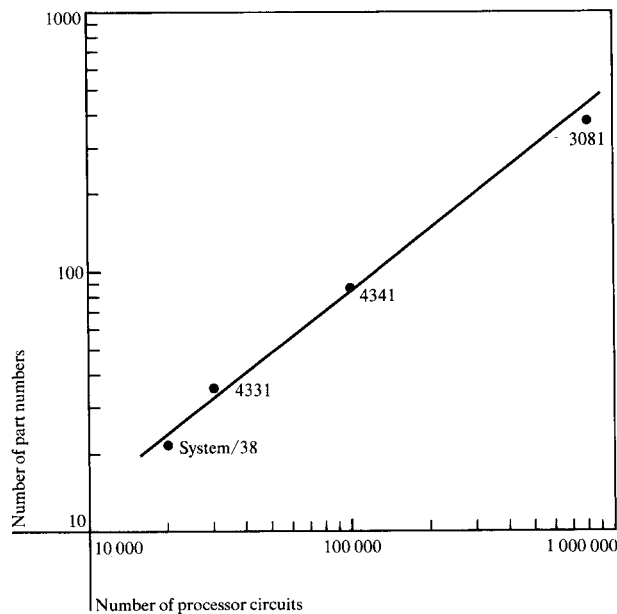


Figure 5 Logic-chip part numbers as a function of processor circuit count (704-circuit logic chip).

1.15 ns. Power dissipation averages approximately 1 mW per circuit. A -1.5 -volt supply powers the logic gates and a -4.25 -volt supply is used to bias resistors available for interchip connections on both driver and receiver circuits.

Logic circuit

Figure 4 shows an example of a TTL circuit configuration. The basic logic gate has either a three-way or a four-way input device with separate emitter and base contacts and a common collector, one output transistor, Schottky clamping diodes, and multiple resistors. This device structure provides a circuit which has the flexibility to be used as either an on-chip or an off-chip receiver. On-chip input levels are either 0.0 V or -1.2 V and off-chip input levels are either -0.8 V or -2.0 V.

Logic and termination chips

Logic-chip designs in the IBM System/38, 4300, and 3081 processors all include 704 gates per chip, but they differ in the number of off-chip drivers and in the number of solder connections. The 3081 logic chip has an 11×11 array of solder pads, of which 96 pads are reserved for logic signals. A chip used for off-chip termination contains 94 90-ohm resistors and also has the 11×11 solder-pad array used in logic chips. Packaged gate delays in the 3081 Processor Unit average approximately 2.5 ns per logic level, and power dissipation averages approximately 1.5 watts per chip.

Logic-chip part numbers

The master slice concept [3, 4] is used to permit the flexibility required in designing the logic for large computer systems while maintaining commonality in the manufacture

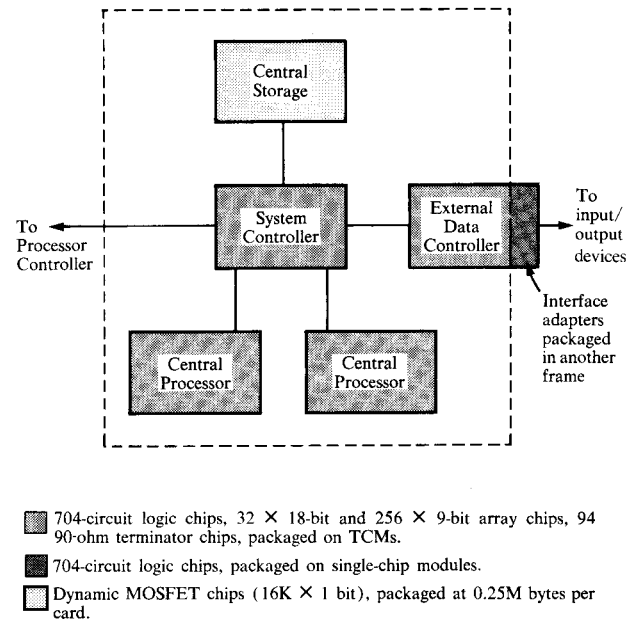


Figure 6 3081 Processor Unit organization and technology.

of silicon wafers. This reduces the turn-around time for new parts and keeps the production cost low. The master slice designs for the System/38, 4331, 4341, and 3081 processors all present the same contact image; thus, one physical design automation program can be used for wiring all these master slices.

The number of unique chips required for a design is a function of the level of circuit integration and the number of circuits in the processor. Figure 5 illustrates, for the IBM systems using the 704-logic-circuit chip, the relationship of the logic-chip part numbers to the total logic circuits in the processor. Succeeding generations of large processors will use ever-increasing numbers of circuits to provide new functions. Any future decision to limit the number of unique part numbers should be based on a tradeoff between design flexibility and cost savings resulting from multiple usage of parts.

Array chips

In addition to the 704-circuit master slice chip, custom-designed array chips are used to implement the general registers, local working registers, buffer storage, and control storage; and extensive use of arrays is made in the central processors for backup registers used for microcode tracing and for the storage of checkpoint information required for error recovery in instruction execution sequences.

Bipolar technology is used for both a 32×18 -bit array chip with 7 ns access time and for a 256×9 -bit array chip with 16 ns access time. The latter chip is used for control

storage and for buffer storage and is actually implemented as a 256×10 -bit array to provide increased yield early in the development program. Both array chips have the same 11×11 solder-pad array as the logic chip to allow flexibility of placement on module chip sites.

Figure 6 shows the types of technology used and their location within the 3081 Processor Unit.

- *Module-level packaging*

Logic chips used in the System/38 are packaged singularly on a ceramic substrate. In the 4300 processors, up to nine chips are interconnected with molybdenum-metallized alumina, within 18 or 23 layers of ceramic, on one 35-mm or 50-mm module.

The latter approach offers significant advantages over traditional single-chip packaging techniques. Chip interconnections, which previously had to be made at the card or board level, can now be made within the module itself; this reduces delays by providing shorter paths between circuits. The decrease in the number of interconnections also increases system reliability.

The thermal conduction module used in the 3081 Processor Unit represents an extension of the *multilayer ceramic* (MLC) technology used in the IBM 4300 processors; it was achieved through various optimizations among key design factors such as module size, line impedance, power distribution, and engineering-change requirements.

The 90-mm-square module accommodates either 100 or 118 chips. The 100-chip module is intended to be used primarily for logic chips, which have approximately twice as many signal I/O connections as array chips, therefore requiring larger chip sites. The 118-chip module has 48 array-chip sites and 70 logic-chip sites. Array chips can be placed on any chip site, but logic chips are usually restricted to the sites having more I/O connections.

Because of the extensive use of array chips in the processor unit, most modules contain at least a few array chips. The average mix of chips on a module in the central processing element consists of about 60 percent logic chips and 40 percent array chips.

Signal wiring, power distribution, and engineering change capability are achieved through the use of 33 layers of metallized ceramic. A module is connected to the next level of packaging through 1800 pins (1200 pins are available for signals, 500 pins are available for power, and 100 pins are spare).

As many as 45 000 logic circuits are placed on a module, and as many as 5500 signal wires are embedded in the MLC

substrate. A 3081 module has the capacity for logic circuits equivalent to that packaged on eight boards in the 3033 processor.

The paper by Blodgett and Barbour [5] in this issue more fully discusses the key considerations in the design of the thermal conduction module and describes the structural and electrical characteristics of the MLC.

- *Board-level packaging*

Although substantial improvement has been made in the circuit density of a module, the high number of total circuits in today's large processors requires a means for the interconnection of many modules. The method chosen for the 3081 processor is an advanced glass-epoxy printed-circuit board designed in two versions to accommodate either six or nine TCMs. Each module is connected to the board by means of a new connector technology which permits easy insertion of the 1800 module pins into mating board connectors and provides a reliable connection.

The printed-circuit boards measure 600×700 mm, are 4.6 mm thick, and have 20 connection planes: six for signals, twelve for power, one surface plane for engineering changes and overflow wiring, and the other surface plane for module connections. There is approximately one kilometer of signal wiring on an average nine-module board. Approximately 450 A of current are distributed in the power planes of the nine-module board, which consumes about 1400 W of power. The board has a rated capacity of 600 A.

The paper by Seraphim [6] in this issue describes the design of this board and discusses the associated manufacturing and assembly processes. Some of the key challenges in the development were: the development of a new connector technology to satisfy reliability and serviceability requirements; the control of dimensional tolerances over the full board area and in all layers of the laminate material; and the drilling of holes of 0.4 mm diameter on a board 4.6 mm thick, resulting in a high ratio of board-thickness-to-hole-diameter (referred to as aspect ratio).

- *Interconnections*

To achieve the objectives for performance and reliability made possible with LSI technology, integration must occur at all packaging levels on the system. In many respects, increased circuit density at the chip level is the easiest to achieve; therefore, small systems have been the first to benefit from LSI. The focus of attention in large computer systems is on the migration of circuit interconnections from the card and board packaging levels into chips and modules.

A measure of the improvement in LSI integration is the amount of signal wire at the chip and module levels




Figure 7 Cross section of the thermal conduction module.

Table 2 Percentage distribution of total wire count by packaging level.

Packaging level	3033 Processor	3081 Processor
Chip	63.0	94.1
Module	—	5.4
Card	33.2	—
Board	2.2	0.4
Cable	1.6	0.1
	100%	100%

Table 3 Percentage distribution of total wire length by packaging level.

Packaging level	3033 Processor	3081 Processor
Chip	1.3	9.2
Module	—	41.9
Card	38.8	—
Board	12.9	31.1
Cable	47.0	17.8
	100%	100%

compared to that previously existing at the card and board levels of packaging. Table 2 shows the percentage of total wires counted, at each level of packaging, in the IBM 3081 and 3033 processors for an equivalent amount of logic. Table 3 shows the percentage of total wire length at each packaging level.

• Module cooling

A 3081 module generates a power density (heat flux) of the order of 10^5 W/m². This is about a tenfold increase compared to the heat flux generated at the module level in the 3033 processor, which uses forced-air cooling over the modules and interboard heat exchangers to cool the air. Extensions to this cooling technology (such as increasing

forced-air flow) were not suitable for removing the heat flux of the 3081 module; a new cooling method had to be used.

Basic thermal problems associated with a high-performance package are the transfer of heat from the source (e.g., the chip) to an external package surface, the removal of heat from the external surface by a coolant, and the maintenance of coolant temperature. The above problems were addressed in the development of the *thermal conduction module* (TCM).

A cross section of the thermal conduction module is illustrated in Fig. 7. The TCM has a helium-filled module cap which provides a thermal conduction path from the back of each chip, via one contacting piston per chip, to the cap. The cap is in contact with a water-cooled cold plate (water jacket) which is removable for module service without disconnecting the supply of circulating water. A separate coolant distribution unit is used to pump water (at approximately 24°C) through the cold plate. The TCM maintains a circuit-junction temperature in the range of 40°C to 85°C and provides a cooling capability which exceeds the requirements for the 3081 processor.

The paper by Chu, Hwang, and Simons [7] in this issue describes the thermal problems, the design alternatives considered, and the theory of operation of the TCM using one-dimensional models. The paper by Oktay and Kammerer [8] in this issue discusses the thermal, mechanical, and environmental experiments, as well as the extended three-dimensional analytic modeling required to develop the TCM.

Reliability, availability, and serviceability (RAS)

The lower cost per circuit of LSI technology as well as its intrinsic reliability provides potential for improved system RAS. The design of the 3081 processor complex includes added function for error-detection and error-correction mechanisms, as well as for hardware error recovery and reconfiguration. Additionally, serviceability is improved through a new approach to diagnostics.

The intrinsic reliability of the 3081 LSI package is a major factor in the reduction of solid failures compared to those experienced using previous technologies. The intermittent type of failure, which has always been a problem, now becomes more important and requires an approach to error detection and fault isolation considerably different from the error re-creation strategy used in predecessor large-scale IBM systems.

In past systems, an error re-creation strategy was satisfactory because conventional tools and human intelligence were sufficient to isolate and repair the fault. But, within the LSI package used in this system, access to chip I/O pins is

prohibited—in fact, the field engineer does not have access to the TCM I/O pins.

Thus, LSI technology requires an approach to fault isolation and service different from that used in past systems; but LSI technology also offers the opportunity to develop a different approach with new considerations for logic design, packaging, and processor controller functions.

- **Error detection**

The design of the 3081 processor places greater emphasis on instantaneous error detection and associated analysis procedures than has been the case in previous IBM large processor systems. The strategy demands that error checkers be designed and located so that the logical state of the processor (in terms of “1” and “0” levels) is captured at the time of the error.

It takes several cycles to stop the clocks of the hardware element which has reported the error. The use of arrays with high density and speed makes it possible to capture all state data necessary to isolate a fault which has been detected in the processor. Shift latches using Level-Sensitive Scan Design (LSSD) [9] permit data to be scanned out to the processor controller for analysis. The Processor Unit contains 126 scan rings, containing a total of almost 35 000 bits of data, which are read out in less than 100 ms. In addition to permitting data to be scanned out, LSSD allows data to be scanned in to set the state of the processor to any predetermined value.

- **Hardware error recovery and reconfiguration**

After fault-related data are scanned out to the processor controller, an attempt is made to recover from a hardware error by resumption of processing of the interrupted instruction execution sequence. The instruction element is designed to establish a checkpoint at approximately every 10 to 20 instructions, to which the state of the machine may be restored in order to retry instructions. Extensive use is made of the 32×18 -bit 7-ns array chip to save the contents of altered general registers, floating-point registers, the program-status word (PSW), and buffer storage. Up to eight attempts are made to retry instructions.

In the event that there is a solid failure, the 3081 Processor Unit can be reconfigured. The dyadic organization makes it possible to logically remove a failing central processor and to continue processing with the remaining central processor until the customer chooses to schedule maintenance. Similarly, a portion of central storage, a single channel group, or a single channel interface may be logically removed from the processor organization while still permitting processing to continue, although system performance may be degraded.

- **Diagnostics and serviceability**

The diagnostics of the 3081 differ considerably from those used with previous large IBM processors. A key objective was to isolate to as few TCMs as possible per failure, ideally only the failing one. Another objective was to isolate the failing TCM (which is the field-replaceable unit), on the basis of the data captured at the time of error detection, for a very high percentage of failures. Analysis routines, which run automatically and concurrently with customer operation, identify a single failing TCM a high percentage of the time. Full isolation to a FRU is achieved for solid failures by using validation tests, automatically called when required by the processor controller, which detect and isolate the type of fault in which logic is stuck at a “1” or at a “0” level. Test data are entered into the processor via the LSSD scan-in data paths. The validation tests are run by the field engineer.

As stated previously, the RAS strategy for the 3081 Processor Unit must be, and is, significantly different than for previous large systems. The approach used for *error detection* and *fault isolation* (ED/FI) and the model used for projecting the ability of the system to detect and isolate failures are described in the paper by Bossen and Hsiao [10]. The paper by Tendolkar and Swann [11] presents the concepts of automated diagnostics implemented for the 3081 and describes the technique used to estimate the isolation effectiveness of the diagnostic package. Integral to this RAS strategy is the implementation and function of the processor controller, which was previously mentioned and which is discussed in detail in Ref. [2].

Computer support

The dependency upon computer support for physical design and test generation is well known [12]. The application of LSI technology has further extended computer support through new requirements for physical design [13], the automatic generation of tests, system design aids, and the control and tracking of complex manufacturing processes.

- **Engineering design system**

Most of the computer support is provided by an *engineering design system* (EDS), which includes programs that check for wireability and that generate tests for each packaging level, and which also generate the physical design of shapes, patterns, precise locations, and interconnections to complete the total fabrication process. These programs, and others used for design and analysis, share a substantial amount of engineering design data. The use of one system, EDS, to maintain the data in a set of interrelated data bases reduces human intervention in data transcription, increases data integrity, and avoids duplication of resources. Extensive use was made of computer support and EDS data throughout the 3081 development effort and continues to be a significant factor in the manufacturing process.

An example which illustrates the continuity of data provided by the engineering design system is the use of physical design data to test design integrity. The verification of the system design and the generation of component tests and system validation tests all require information derived from the EDS physical design data base. The fabrication of chips, module substrates, and boards, the bonding of chips onto MLC substrates, the addition of overflow wiring, and the final assembly of the thermal conduction module are all accomplished by complex tools which are numerically controlled by information provided by the engineering design system.

A high level of integrity in data released to manufacturing is brought about by design verification and the engineering design system. This factor plus the degree of automation in process control, assembly, and testing achieved for this new technology has contributed to a shorter manufacturing time compared to that of previous technologies.

The cost of computer support for the development of the 3081 was about ten times greater than that for the 3033; approximately one third of this cost was for design verification and the rest was for physical design, test generation, and logic entry. Highlights of the new tools used for design verification follow.

• *Design verification*

Even when a determined effort is made to generate accurate designs, the use of a totally new technology coupled with new considerations for machine organization and RAS will result in errors which can be very costly to the development schedule. Therefore, a goal was set to identify a substantial percentage of the logic design problems and 100% of the timing problems through design verification; the remaining design problems were to be isolated on hardware.

Functional simulation and Boolean validation

Flowcharts, manually drawn by designers to represent functional logic, are the basis for input to a simulator which models design behavior on a cycle-by-cycle basis. Corrected flowcharts provide a functional specification of the logic, against which the actual low-level logic design is validated. Hundreds of logic circuits are compared for functional equivalence via a time-efficient Boolean comparison algorithm, rather than by simulation of low-level logic.

Timing analysis

The prediction of propagation delays in logic paths, which in some cases cross chip, module, and board boundaries, has become a highly complex problem requiring new techniques. Considerable computing resources are required to identify critical paths and to calculate path-delay times and probabilities. The dependence on details of the physical design of the logic to calculate wiring and circuit delays in logic paths

means that the designer must go through an iterative process of logic design, physical design, and timing analysis to achieve machine-cycle-time objectives before the manufacture of personalized chips and modules. The use of timing analysis in the design of the 3081 Processor Unit was very effective in minimizing the timing problems to be resolved on the hardware and to maximize performance.

Descriptions of the new tools, techniques, and procedures used for verification of the design and discussions of the significance of design verification in terms of early identification of engineering changes are contained in three papers of this issue. The first paper, by Monachino [14], provides an extensive discussion on the design verification methodology used and the value of identifying problems during design prior to the manufacturing process. The next two papers concentrate on design verification tools: the paper by Hitchcock, Smith, and Cheng [15] discusses the concepts of timing analysis and the engineering considerations which were reflected in the development of that tool, and the paper by Smith, Bahnsen, and Halliwell [16] describes the functional verification methodology, the Boolean comparison process, and the algorithms used.

Summary and conclusions

The 3081 marks the beginning of IBM's application of LSI technology to large-scale processing systems. Advancements in silicon, packaging, and interconnection technologies have yielded improvements in packaged circuit performance and reliability. Card wiring has been replaced by module wiring, with more than 40% of the total wire length at the module level. This achievement, plus the placement of 94% of the total wire count at the chip level, has contributed to a processor cycle time which is less than half that of the 3033. Product objectives for improved reliability and processing capability were met while achieving significant reductions in power consumption, cooling, and floor space.

These results have been brought about largely through advancements in technology, through changes in system design and in the design process to increase integrity and reliability, and by the extensive use of computer aids for improved evaluation of design.

The future will continue the trend set by this new effort. The potential is there for improvements in performance and reliability as chip circuit densities are increased and as logic and arrays are combined in the same chip. Interconnections will be reduced and packaging density will be increased and, as a result of the cooling capability in the thermal conduction module, higher levels of circuit performance will be realized. Also, the experience gained in the application of LSI technology to system design will be used to enhance the organizational structure with more sophisticated designs.

A number of observations have been made about the interaction required among designers, particularly those engaged in system design and packaging design. It can generally be concluded that the successful application of LSI technology is the result of the symbiosis of many design disciplines; and this implies that designers will continue to require broader knowledge of the relationship of their speciality to that of others.

In conclusion, it has taken numerous technical inventions and advancements, changes in the development process, and extensive time and resources to apply LSI technology to the development of the 3081 Processor Unit. Its introduction, however, marks the beginning of a new era in large-processor design, with continued development challenges but also with the promise of greater levels of product reliability and performance at lower cost.

Acknowledgments

The adequate acknowledgment of specific individual contributions to a development effort of this size is impossible. Thanks and appreciation are expressed to all persons whose continued support and contributions have been instrumental in bringing this development effort to completion. The authors thank and acknowledge the efforts of W. J. Kleinfelder, G. F. Rollo, and C. E. Galbraith, who contributed to the technical material, the organization, and the editing of this paper. C. A. Collins provided data on interconnection comparisons; and V. L. Gani, W. S. Klara, and C. E. Vicary provided data on circuits and chips. Thanks and appreciation are also due to many others who provided detailed technical data, to those who provided insight into the impact of LSI on detailed technology and system design processes, and to those who reviewed the manuscript and suggested improvements.

References

1. R. N. Gustafson and F. J. Sparacio, "IBM 3081 Processor Unit: Design Considerations and Design Process," *IBM J. Res. Develop.* **26**, 12-21 (1982, this issue).
2. John Reilly, Arthur Sutton, Robert Nasser, and Robert Griscom, "Processor Controller for the IBM 3081," *IBM J. Res. Develop.* **26**, 22-29 (1982, this issue).

3. J. Pomeranz, R. Nijhuis, and C. Vicary, "Customized Metal Layers Vary Standard Gate-Array Chips," *Electronics*, 105-108 (March 15, 1979).
4. R. L. Golden, P. A. Latus, and P. Lowy, "Design Automation and the Programmable Logic Array Macro," *IBM J. Res. Develop.* **24**, 23-31 (1980).
5. A. J. Blodgett and D. R. Barbour, "Thermal Conduction Module: A High-Performance Multilayer Ceramic Package," *IBM J. Res. Develop.* **26**, 30-36 (1982, this issue).
6. Donald P. Seraphim, "A New Set of Printed-Circuit Technologies for the IBM 3081 Processor Unit," *IBM J. Res. Develop.* **26**, 37-44 (1982, this issue).
7. R. C. Chu, U. P. Hwang, and R. E. Simons, "Conduction Cooling for an LSI Package: A One-Dimensional Approach," *IBM J. Res. Develop.* **26**, 45-54 (1982, this issue).
8. S. Oktay and H. C. Kammerer, "A Conduction-Cooled Module for High-Performance LSI Devices," *IBM J. Res. Develop.* **26**, 55-66 (1982, this issue).
9. E. B. Eichelberger and T. W. Williams, "A Logic Design Structure for LSI Testability," *Proceedings of the 14th Design Automation Conference*, New Orleans, LA, June 1977, pp. 462-468.
10. D. C. Bossen and M. Y. Hsiao, "Model for Transient and Permanent Error-Detection and Fault-Isolation Coverage," *IBM J. Res. Develop.* **26**, 67-77 (1982, this issue).
11. Nandakumar N. Tendolkar and Robert L. Swann, "Automated Diagnostic Methodology for the IBM 3081 Processor Complex," *IBM J. Res. Develop.* **26**, 78-88 (1982, this issue).
12. P. W. Case, M. Correia, W. Gianopoulos, W. R. Heller, H. Ofek, T. C. Raymond, R. L. Simek, and C. B. Stieglitz, "Design Automation in IBM," *IBM J. Res. Develop.* **25**, 631-646 (1981).
13. W. R. Heller, W. F. Mikhail, and W. E. Donath, "Prediction of Wiring Space Requirements for LSI," *Proceedings of the 14th Annual Design Automation Conference*, New Orleans, LA, 1977, pp. 32-42.
14. Michael Monachino, "Design Verification System for Large-Scale LSI Designs," *IBM J. Res. Develop.* **26**, 89-99 (1982, this issue).
15. Robert B. Hitchcock, Sr., Gordon L. Smith, and David D. Cheng, "Timing Analysis of Computer Hardware," *IBM J. Res. Develop.* **26**, 100-105 (1982, this issue).
16. Gordon L. Smith, Ralph J. Bahnsen, and Harry Halliwell, "Boolean Comparison of Hardware and Flowcharts," *IBM J. Res. Develop.* **26**, 106-116 (1982, this issue).

Received August 4, 1980; revised July 21, 1981

The authors are located at the IBM Data Systems Division laboratory, Poughkeepsie, New York 12602.