

N. C. Mescia
C. D. Woods

Plant Automation in a Structured Distributed System Environment

As VLSI technology evolves, miniaturization demands more sophisticated tools, instruments, and controls to manufacture the VLSI components. IBM's facility at East Fishkill, New York, has the responsibility for the development and manufacture of semiconductor products used in IBM data systems. This requires a sophisticated set of processes, inspections, and tests operating as a fully integrated system. In this paper, the design and implementation of a hierarchical distributed system for manufacturing control of integrated electronic components are described. The implementation includes distributed data bases and inter-level decoupling to ensure 24-hour manufacturing capabilities. Reasons for the choice of the processors used at various levels in the hierarchical network, and the communications required between them, are discussed.

Introduction

As Very Large Scale Integration (VLSI) technology evolves during the next and future decades, computers will continue to greatly alter the way we work and live. This alteration will stem from the further development of a wide range of computer products, from small microprocessors through large mainframes. In order to achieve this goal and others, IBM's General Technology Division (GTD) facility at East Fishkill, NY, has the responsibility for development and manufacturing of VLSI semiconductor logic devices and multilayer ceramic (MLC) substrates used in current and future IBM systems. The manufacture of these devices involves many sophisticated tools and processes. The level of automation in this environment requires that individual process and tool automation be structured to ensure that systems can easily be interconnected and that the skills developed for this purpose can be applied interchangeably across a wide base of applications. This paper describes the complex but structured hierarchical distributed manufacturing control system known as the Component Manufacturing Control System (CMCS-80), an interconnection of microcomputers, minicomputers, and mainframes used for manufacturing control at the IBM East Fishkill facility.

This facility has had previous experience in tool and instrument automation via host computers, minicomputers,

and microcomputers. As the price/performance aspects of computers have improved, it has become more economically feasible to automate both tools and processes. One of the main purposes of this paper is to illustrate what efforts must be considered for integration of these automated tools and processes into a manufacturing environment. In the 1970s, a Quick Turn-Around Time (QTAT) single-wafer semiconductor processing and production line was developed that was computer-controlled [1]. This highly automated line was developed in response to a demand for short production time of low-volume first-order parts for development of new IBM computers. This complemented the site's larger batch-processing production line. The IBM Burlington facility at Essex Junction, VT, also developed an automated manufacturing line [2] during the 1970s for processing of production VLSI memory wafers. This line, like the QTAT line at East Fishkill, employed host computers, minicomputers, and microcomputers. As previously stated, the manufacture, interconnection, and testing of these devices and substrates requires sophisticated processes, tools, and instruments. There is a great deal of sophisticated computer-controlled equipment that is being developed and offered commercially for these purposes. Of course, what is sophisticated today will be history tomorrow, and tomorrow's offerings will be even more sophisticated. No one company can keep abreast of all

© Copyright 1982 by International Business Machines Corporation. Copying in printed form for private use 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, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.




Figure 1 MLC substrate layers. This view exemplifies the various layers and the complexity of a multilayer ceramic substrate.

Table 1 Description of the various steps used in substrate manufacturing.

- **Raw material preparation:** Precise amounts of raw materials are mixed to form a thick slurry.
- **Casting:** The slurry is cast in a thin strip of "green" (unfired) ceramic material.
- **Blanking:** Green sheets, 185 mm square, are cut (blanked) from the material.
- **Punching:** As many as 35 000 "via" holes—each barely the diameter of a human hair—may be punched in a single sheet.
- **Screening:** A pattern of conductive paste is screened onto the green sheets, and into the vias, through a metal mask. Conductive paste in the via holes provides contact between layers of the substrate.
- **Stacking:** Screened sheets with varying patterns are stacked automatically in the order required for the particular substrate being built, then laminated under heat and pressure.
- **Sizing:** Excess ceramic material is trimmed from the laminate.
- **Sintering:** The laminate is then fired in a kiln at high temperature.
- **Plating:** Metals are placed over the screened patterns on the top and bottom surfaces of the substrate, preparing it for the subsequent brazing, chip placement, and engineering-change operations.
- **Testing:** A battery of electrical tests ensures that the substrate is free of defects and that the electrical connections have been properly made.
- **Brazing:** Finally, pins and a flange are brazed to the substrate in a single operation. The flange will later seal the substrate into the bottom of the frame. The completed substrate is then shipped to the appropriate assembly location.

tool/instrument offerings. Therefore, in addition to custom-built equipment, equipment must be purchased with various embedded vendor processors, for automation of state-

of-the-art manufacturing processes. As the sophistication of individual tools increases, there is a corresponding need to develop capabilities to coordinate control of work in process, to correlate test results back to previous process measurements, and to track the reliability and performance of all the process equipment.

The nature and scope of the capabilities required are beyond the traditional centralized computer support function. East Fishkill manufacturing processes are segmented into more than fourteen separate but interrelated manufacturing lines. Thus, a key feature of the system developed to serve this environment is to distribute function in a structured hierarchy, where the information base can be shared across all lines, while the manufacturing support is segmented into dedicated computers that can have some customization for each area. Another significant aspect of this structured distribution of computers is that it has sufficient autonomy, or "decoupling," at each node, to ensure minimal impact of single-system failures.

This paper describes the reasons for the selection of the CMCS-80 distributed processing hierarchy, the rationale for the hardware selection, the communications discipline chosen, the distributed data bases used, and the development of automated equipment applications.

Distributed system

The advent of automated tools and processes further created more decentralized applications. From an applications point of view, this was desirable because of previous problems associated with the availability and responsiveness of host computer control. However, centralized control was deemed to be extremely important along with the need for decentralized processors to access centralized host data bases for logistics control. A distributed system was determined to be the best approach to allow communications between these processors and the host. In selecting the proper distributed system for an application, an investigation must first determine the services required for each of the site's users. The services required at East Fishkill range from batch operations at the host to real-time applications at the tool controller. Each of the applications serves a different need, but all these needs must be blended into a customized distributed system.

In order to manufacture VLSI devices, it is necessary to process the product in a series of steps. These steps are described in Tables 1 and 2 [3]. They are grouped into different manufacturing areas, each of which is managed separately. The product may be recycled through these areas as it is processed through its repetitive layers. Figures 1 and 2 illustrate the various layers of an MLC substrate and wafer, respectively. The product is stocked at control centers

Table 2 Description of the various steps used in VLSI chip manufacturing.

- **Crystal growing:** High-purity silicon is melted and impurities are added to provide desired electrical characteristics. A "seed" crystal is brought into contact with the molten silicon. The seed is raised; a single rod-shaped crystal is drawn from the crucible.
- **Wafer slicing:** The rod is trimmed and ground to a uniform diameter. Razor-thin wafers are then sliced from it, smoothed, and polished.
- **Epitaxial growth:** The protective oxide layer is removed and an additional layer of silicon is grown to provide a working surface. A new oxide layer is then formed for protection.
- **Photoresist:** The entire wafer is coated with light-sensitive photoresist and exposed to ultraviolet light through a mask. An etching process removes exposed photoresist, leaving microscopic "windows" on each chip site on the surface of the wafer.
- **Diffusion:** Wafers, and a source material, are sealed in a quartz tube and placed in a high-temperature furnace. As the material evaporates into a gas, its atoms diffuse into the silicon through the windows created earlier. Areas altered by the diffused atoms become the transistors, resistors, and diodes required for the electronic circuitry of a chip.
- **Metal evaporation:** A protective layer of glass is formed over the wafer and etched to create more windows. Aluminum is then evaporated over the wafer. Electrical connections between the electronic components of each chip are made through its windows.
- **Electron-beam:** "Wires" are formed by coating the aluminum layer with photoresist and either "writing" a wiring pattern with an electron-beam tool or forming it with a series of masks.
- **Test:** Each chip on the wafer undergoes a battery of exhaustive electrical tests which verify that its circuitry is functioning properly.
- **Dicing:** A diamond-impregnated saw is used to cut the individual chips from the wafer.
- **Shipment:** Finished chips are packed in a special carrier and shipped to the appropriate IBM location for assembly.

between these manufacturing areas. Some of these stocking points differentiate between the manufacturing lines within East Fishkill. The semiconductor wafers, for example, are produced up to the stage called a "masterslice." These masterslices have all the basic circuits and are built to stock until the orders for custom part numbers are received from other IBM locations. These masterslices will then be personalized by further processing to interconnect the circuits as specified by the part number. In a parallel mode, MLC substrates are manufactured in various sizes and layers [Thermal Conduction Module (TCM), Multiple-Chip Module (MCM), and Single-Chip Module (SCM)] [4] in their manufacturing area. These products also have intermediate stock points within their manufacturing line (blank sheets, punched sheets, finished substrates). The chips and substrates can then be shipped separately, or after chip joining,




Figure 2 Customized bipolar wafer layers. Illustrates the interconnection of various layers of metallization and silicon dioxide deposited on the masterslice.




Figure 3 (a) Chip-joining process; (b) cross section of a chip-joined MLC substrate illustrating the various possible via paths from the chip sites to the I/O pins. (Note: C-4 = Controlled Collapse Chip Connection.)

to other IBM locations for IBM products. Figure 3 denotes the chip-joining process and a chip-joined MLC substrate.

Producing VLSI devices and multilayer ceramic substrates with automated tools requires that numeric control data define the process steps and that adjustments be sent forward to other tools based on in-line product measurements. The feed-forward of these instructions must be accomplished by the system that supervises the tools.




Figure 4 Distributed architecture. There are five levels in the CMCS-80 system hierarchy; each level can operate independently in the event of a possible outage.

Figure 4 illustrates the CMCS-80 distributed architecture and the decoupling capability of the tool controllers, area controllers, and line controllers.

It is important that the *tool controller* act autonomously to complete the processing of the job without any interaction with other controllers. The operator must be able to complete the process with the equipment and computer control at the work station. Tool controller involvement may range from monitoring to limited control and operator guidance to process automation.

Within an area it is necessary to have some supervision and coordination of equipment controllers. Therefore, *area control* systems were derived. These area control systems must be able to supervise up to 20 to 60 tools or process steps and also support 10 to 20 terminals for the display of status information, and they must be extremely reliable and responsive. The objective initially established for the area control system was 99.5% or better availability. A problem anticipated with this objective was that there was no appreciation of the numeric significance of any difference at this level of availability. We decided to transform that into something that was meaningful and that seemed significant in direct user terms. The best measure was judged to be minutes of outage. Thus the 99.5% objective was changed to fifty minutes or less per week of *unavailability* in order that we might have a user-oriented target to drive to a zero level.

Two design points of the area control system are (1) that it cannot require an operator, and (2) that it cannot require extensive system tuning. We anticipated a large number of systems across the site and did not want to allocate a

significant amount of resource for optimizing the hardware utilization. We set a planning limit of 20% throughput utilization in order to ensure that the system has the reserve capacity to handle unplanned workload and short-term bursts of transactions.

The area control system is intended to assist the line management in the day-to-day operation of the manufacturing line. With a dedicated *area controller*, it is now possible to have the computer resource and reliability necessary to implement detailed product routing and tracking that optimize the product throughput of an area. This also controls and minimizes the *work in process* (WIP), since the area controller has the necessary information, at an individual operation level, to make immediate decisions and is independent of the *line controller* for up to eight hours of operation.

Area controllers are grouped by line since the East Fishkill site is divided into multiple manufacturing lines. The differing demands, priorities, and management approaches of these lines require an approach that will handle these differences and also be able to expand a line controller in capacity over time. The choice was to segment the line controller function from the normal *site host* computer. It was anticipated to have multiple area controllers reporting into a line controller; in fact, a single line controller would initially supervise multiple manufacturing lines. Within these guidelines, it is expected that five line controllers will supervise the 14 manufacturing lines. The reliability objective for this level was originally set at 99%; again we translated this into a user-oriented term of no more than 100 minutes per week outage.

A line controller is responsible for supervising the general flow of product within a manufacturing line. It is more oriented toward control-center operations which supervise the flow of product from area to area than toward the operation within an area. The line controller will be the primary source for generating any daily reports. The data summarization for a manufacturing line will occur at this level as well, with long-term data archiving performed at the site level. In order to be effective, this system must operate with minimum operator intervention. To minimize tuning activity and to provide for response to short-term bursts of transactions, we anticipate running this system at approximately 40% utilization. This system will be independent of the site-host-level system, with product orders and process information, for up to twenty-four hours.

The *site-host* level is responsible for the site records that describe the basic characteristics of the product, the process flow, and the numeric control data necessary to build and test the product. Job release occurs at this level based on the orders received from other IBM locations. The processing

instructions and any necessary numeric control data are distributed out to the line controllers as those orders are released. When the product processing is complete, summarized data and detailed data are received by the site level. The site system retains the data necessary to do product analysis and also has perpetual inventory responsibility.

Hardware choices

At the *tool level* we had determined that we would use the Series/1 [5] with the EDX language and operating system [6] in order to provide higher-level functions for the programmers and some measure of transportability to new programming environments that may develop in the future. For example, the EDX language and operating system had already been implemented on the System/7 [7], and, in the 1977 time frame, was made available on the Series/1. In general, we do not use minimal computer hardware configurations; we use a standard configuration that satisfies a wide range of tool environments and has all the facilities possible to maintain high programmer productivity. Although the Series/1 was the initial choice for most tool-controller applications, there existed a need for a microcomputer for smaller applications. This need was satisfied with an internal offering of a *distributed system controller* (DSC) [8]. Again, the EDX language and operating system was used to preserve previously developed skills.

At the *area level*, we again chose the Series/1. We would need between 40 and 60 of these processors across the site, and there is an objective that these be installed physically within each manufacturing area. The Series/1, in its 19-inch rack, only requires 32 to 80 square feet of floor space. It also has an automatic IPL feature and satisfies the desire for a low-cost processor. Again EDX was chosen as the language and operating system at this level to preserve the skills we had developed in tool-automation programming. In the 1977 time frame, it was necessary to also develop a terminal manager, a data base manager, a transaction-scheduling supervisor, and a communication program that today are offered as part of standard packages such as the Communication Facility II (CF) [9]. The area controller Series/1 typically has a larger configuration than the tool controller Series/1.

The *line control* system seemed best suited to the System/370 environment being used at the site host system. This allows migration of programs already developed in the MVS/IMS [10] environment. We wanted to protect that programming investment as well as the skills developed by the programmers. The 43XX systems [11] satisfy the need to have physically small machines that could be afforded on an individual-manufacturing-line basis. They did not satisfy the objective of operating without operator intervention; however, we expected to install them initially in groups near the




Figure 5 Data communications hierarchy. The CMCS-80 system uses many processors, via various communications interfaces, to ensure flexibility. (Note: TCU = Transmission Control Unit, SBCA = Sensor-Based Communications Adapter, DSC = Distributed System Controller, DIS = Distributed Interface System, MHS = Magnetic Hand Scanner, and RS-232 and IEEE-488 are standard communications interfaces.)

data center. Since we would be using the 4341 processor, we expected to have four or five of these initially to cover the lines within the East Fishkill site.

The hardware choice for the *site-host-level* system was already determined when we first developed the distributed architecture. This was a network of System/370 computers with the MVS operating system using IMS and PL/I application programs.

Communications

For communications between systems, we chose a discipline that allowed only one transaction to occur at one time between any two machines. This had the effect of balancing the burst of transactions that might come up from any one controller to the next level. It also provided a simple test and debug environment. Only one transaction would be in process at a time, and the time relationship of all the events could be determined precisely within each machine. This still permits effective multiprogramming at each of the levels, since our hierarchical structure has the area, line, and site systems effectively acting as concentrators for transactions from the lower levels. An area controller with 20 tool controllers, for example, could have up to 20 transactions in process simulta-

neously—one from each tool controller. Figure 5 illustrates the communication between the various levels. The Series/1 feature #1400 (Local Communication Controller) [12], supported under EDX Communications Facility II, is under investigation as a replacement for RS-232 Communications [13] for faster transaction times between Series/1 processors.

System capacity planning

An important part of the success of the QTAT support was the ability to predict and analyze performance problems involving the entire network. Two models were needed to address that problem. The first model addressed the performance and response of the line controller under varying loads. This primarily focused on the communication load from the automated tools controlled by the satellite systems; the workload generated by people at terminals was relatively small. The response to the terminal users was presumed adequate as long as the total system utilization was below 60%.

The first model identified problems with the process measurement plans; this by itself was not that surprising. The initial set of an eventual 28 satellites had essentially turned out the "WAIT" light on the host system when they were started up at very high data collection rates. It is known that, with our high-speed communications capability and a large number of tool controllers, we would have to limit the total data collection across the system. What was lacking, and what the model provided, was a way of studying the interaction of various data collection strategies with changing product volumes.

This led to the second model, which addressed the prediction of the transaction load due to product volumes. This host system not only supported the QTAT facility but also handled the product tracking for the rest of the East Fishkill site. This model was used to evaluate the impact of current and future product volumes in conjunction with batch size and yield estimates. This provided the predicted transaction load due to total site activity. The relatively small portion of product that would be handled on a single-wafer basis in QTAT would generate a significantly higher transaction load. These wafers are tracked individually through each operation; in regular manufacturing, wafers are grouped into batches and only tracked by major process area or sector.

The predictions of both of these models taken together provided a basis for management of the total manufacturing control system. For the distributed version of CMCS-80, it is necessary for a hierarchy of models to be built. These will be designed in the same fashion as the existing models, but structured to allow the output of the *area* level models to drive the *line* and *site* models.

In the real system, the wafer and batch release process is initiated at the site level and drives transactions for routing and numerical control data distribution down to the line and area levels. For simplification of the model, this workload will be estimated by the area model and passed on to the line and site models. The prediction of data storage requirements at each level is an additional feature that will be added to these models. This is influenced by both product volumes and retention strategies; it will provide a set of tools to predict and manage the workload at each level as well as across the entire network. Each management team will be able to adjust their data collection strategies within the limitations of their area controllers. The quantified workload estimates for the line and site controllers then allow negotiation and planning for system capacity across the site.

Distributed data base

Within this multilevel hierarchy, it is necessary to partition or distribute data bases across the various levels. The product-routing data base originates from the site system, which serves as a central point of control. From there, an approved routing is distributed out to the line controller, which then distributes it to the area controllers when the product is released. Security for terminal users, on the other hand, is partitioned into local files; the management responsible for each machine ensures the integrity of their own security files.

The *numeric control* data are handled both with central distribution from the site host and with local files at the individual areas. Product descriptions come in from other IBM locations that describe the interconnection of the circuits within semiconductor chips or multilayer ceramic substrates. These data are processed at the site level to produce the control data to create and test the product. The numeric control data are distributed from the site level down through the hierarchy. Other data, which typically address parametric evaluation of the process, are generated on a local basis within a manufacturing area. These data are shared locally within the area controller and tool controllers of that manufacturing area.

Product and process measurement data may be kept at any one of the levels. Data of a very detailed nature are kept only at the tools. Other data are sent in either detailed or summarized form to the area controller, where they may be kept for up to a week. Data from the area controller are summarized at the line level for inclusion in our process data base at the site level. These data may or may not be sent on to the site-host level, depending upon the data collection strategies in place at the time as determined by the process engineers. Summary and detailed data that are to be retained for longer than the processing time of the product itself must be archived at the site-host level.

Decoupling

In constructing such a large and expanding network of systems, it is important that each system tolerate interruptions of communications. The levels above or below a particular point in the hierarchy may not be available at some time for communication. These outages, either planned or unplanned, should not disrupt the operation of any of the systems. At the tool level, each tool controller must be able to complete its work without any additional communication with the area controller. Therefore, all information necessary to process the job must be loaded from the area controller when the job is started, and the tool controller must be able to buffer any critical transactions to the area during processing. The tool controller must handle all emergencies without external intervention from either an operator or the area controller.

The area controller must be able to supervise the work on all the product that has been received into the area. In order to ensure a smooth flow of product through the area independent of any line controller outages, product information is preloaded down to each area from the line controller for eight hours' worth of new jobs. The area controller stores this information for the entire time that this product is in the area. It also has the storage capacity to buffer transactions to be sent to the line controller for the same decoupling limit of eight hours. The line controller has similar decoupling from the site host, but for a twenty-four-hour period.

CMCS-80 architecture

The following is a listing of the five functional levels of the CMCS-80 structured distributed system illustrated in Fig. 4, and a summary of the characteristics of each level.

CMCS-80 Equipment Control

Functional level: fifth—tool/instrument/process control

Characteristics:

- Real-time
- Sensor-based equipment status
- Equipment sequence control
- Process monitoring and control
- Sensor-based data collection
- Error detection and handling
- Limited operator interface

CMCS-80 Area Controller

Functional level: fourth—area control

Characteristics:

- Interactive
- Area/department/sector status reporting
- Critical data interface
- Data collection buffer to line/host
- Feed-forward of product measurements between tools
- Feed-back of process measurements between tools

- Operation-level tracking and routing verification
- 8-hour line/host independence (decoupling)

CMCS-80 Line Controller

Functional level: third—line control

Characteristics:

- On-line to major manufacturing lines
- Line control status reporting
- Job-claiming interface for area-level tracking
- Feed-forward between areas
- Numerical control data source
- Product/process data summarization
- 24-hour host independence (decoupling)

CMCS-80 Site Host

Functional levels:

first—preparation for manufacturing

second—material and information logistics

Characteristics:

- Major data base support
- Centralized history maintenance
- Primary report generation
- User-structured retrievals/access
- Interface to other locations

Information Systems responsibility for equipment automation

East Fishkill is the largest MLC and VLSI logic manufacturing facility within the IBM Corporation. East Fishkill Information Systems (I.S.) has the overall responsibility to ensure effective use of computer systems throughout the site.

Besides the applications and system programming, I.S. must provide software tools, vendor interfacing specifications, and application programming education. The I.S. mini/micro lab provides an off-line programming facility which allows off-line hardware and software development and debugging without manufacturing interruption. Once the hardware and software have been established, integration into the CMCS-80 architecture is performed. However, CMCS-80 is a site responsibility and requires the interaction of everyone. Figure 6 illustrates that interaction at the tool-application level.

In order to ensure a successful implementation of tool-application programs, cooperation of various engineering groups and I.S. is required. During the development cycle, process and equipment engineering, responsible respectively for manufacturing processes and tools, formalize the user requirements (manufacturing) collectively with I.S. These requirements are further defined in functional specifications, whereby all responsible groups sign off on this formal agreement. Also, to ensure continued successful implementation, all responsible groups are invited to I.S. formal design




Figure 6 Plant interaction. CMCS-80 is everyone's responsibility and everyone's inputs must be considered to ensure a successful automated tool/process application.




Figure 7 Course hierarchy. The courses start with introductions to various microprocessors and continue with more advanced areas, developing skills towards application programming.

reviews. This relationship has proven to be quite effective, with all responsible groups keeping abreast of all stages of implementation.

Education

The I.S. organization, in cooperation with IBM Mid-Hudson Education and Process Control Systems Engineering, has

established a Mini-Micro course hierarchy for internal education. The objective is to educate tool/process engineers, designers, scientists, and other nonprogrammers into either application programmers or more knowledgeable users, to ensure successful applications. The hierarchy starts with the "Introduction to Microprocessors" courses and finishes with advanced tool design and application programming courses. Figure 7 illustrates that hierarchy [14]. Throughout this curriculum, emphasis is placed on the total system environment within which these tools operate. This ensures that complete applications are developed that serve all needs.

Current status and future plans

Over sixty tool controllers have been tied into ten area controllers as of January 1982. The line controller applications and data bases have been developed and installed, but they coincide with the site-host applications in the same single physical machine. The site-host and line-controller applications are today linked by internal secondary transactions. These transactions will continue to be used between the site-host and the line-controller machines when they are physically separated in 1983. Our experience thus far has emphasized the importance of advanced planning and development of common applications support that can be used across all areas.

The completion of installation of the forty to sixty area controllers will probably extend into 1984. The installation of tool controllers will be a continuous process as new applications are automated, old applications are replaced, and new processes are developed. One can foresee the increased use of mechanization and automation, including the use of robots, to increase manufacturing productivity. The ongoing use of intelligent tool controllers and our evolving modular distributed system for plant supervision and control are expected to maintain the IBM East Fishkill facility as a leading manufacturer of VLSI components.

Acknowledgments

The authors would like to acknowledge the contribution of East Fishkill Information Systems Development personnel, particularly those of the Logistics Systems, Manufacturing Control Systems, Sensor-Based Control Systems, and System Architecture functions. We would also like to acknowledge for their individual contributions P. K. Bhalla, J. K. Cruger, J. E. Fischette, A. J. Fleming, T. E. Francis, E. J. Holden, J. C. Pace, F. V. Roviello, B. W. Tague, J. W. Tyger, and P. J. Winiarski, as well as V. R. Warren for his descriptive illustrations.

References and notes

1. R. H. Brunner, E. J. Holden, J. C. Luber, D. T. Mozer, and N. G. Wu, "Automated Semiconductor Line Speeds Custom Chip Production," *Electronics* **54**, 121-127 (1981).

2. Z. Apgar, "Control Systems for Memory Manufacturing," *Technical Bulletin TB-19002*, IBM General Technology Division, Essex Junction, VT, 1980.
3. *East Fishkill Circuit*, Employee Communications Dept., IBM General Technology Division, East Fishkill, NY, September 1981.
4. D. P. Seraphim and I. Feinberg, "Electronic Packaging Evolution in IBM," *IBM J. Res. Develop.* **25**, 617-630 (1981).
5. *IBM Series/1 Digest*, Order No. G360-0061; available through IBM branch offices.
6. *Event Driven Executive Language Reference*, Order No. SC34-0314; available through IBM branch offices.
7. *System/7 Event Driven Executive*, Order No. SB30-0812; available through IBM branch offices.
8. R. M. Fiorenza, A. J. Fleming, R. J. Gerlach, L. W. Holmstrom, and J. C. Pace, "Computer Controlled Distribution Apparatus for Distributing Transactions to and from Controlled Machine Tools," U.S. Patent 4,069,488, 1978.
9. *IBM Series/1 EDX Communications Facility*, Order No. GB30-1593; available through IBM branch offices.
10. *IBM/VS Version 1 General Information Manual*, Order No. GH20-1260; available through IBM branch offices.
11. *Introduction to IBM 4300 and DOS/VSE Facilities*, Order No. GR20-4666; available through IBM branch offices.
12. *Series/1 Local Communication Controller*, Feature No. 1400 of an IBM 4955 Series/1 Processor.
13. *Teletypewriter Adapter*, Feature No. 7850 of an IBM 4955 Series/1 Processor.
14. *Engineering Education Guide*, Mid-Hudson Education, IBM Data Systems Division, Poughkeepsie, NY, April 1981.

Received October 6, 1981; revised February 16, 1982

The authors are located at the IBM General Technology Division laboratory, East Fishkill Facility, Hopewell Junction, New York 12533.