Volume 26, Number 1, 1982

System Development and Technology Aspects of the IBM 3081 Processor Complex *by M. S. Pittler, D. M. Powers, and D. L. Schnabel, p. 2.* 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.

IBM 3081 Processor Unit: Design Considerations and Design Process *by R. N. Gustafson and F. J. Sparacio, p. 12.* Significantly new challenges were presented for the design of the 3081 Processor Unit since it was the first IBM large system implemented in LSI technology. Solutions had to be found for a new set of problems in order to achieve the required product objectives while maintaining an acceptable development cost and schedule. In this paper, the design aspects and the characteristics of the 3081 Processor Unit are described and the tradeoffs that were made due to the implementation of LSI are presented. A design strategy was chosen that included tradeoffs covering the areas of machine organization, performance level, implementation costs, testing and servicing aids, and development schedules. An innovative verification effort was introduced into the design process, capitalizing on a hardware flowcharting discipline and rigorous design rules. On the basis of this development experience, some thoughts concerning VLSI implementations are explored.

Processor Controller for the IBM 3081 *by J. Reilly, A. Sutton, R. Nasser, and R. Griscom, p. 22.* The IBM 3081 represents an important step in the evolution of large-scale data processing systems. The incorporation of LSI technology in its design has resulted in a departure from the use of traditional support tools and techniques. Its increased processing capabilities and its dyadic organization have further accented the requirements for high availability and ease of operation. This paper describes the Processor Controller unit of the 3081 Processor Complex, which handles the functions of maintenance, monitoring, and control of the system. It discusses how this unit has dealt with these changes to keep pace with the demands of advanced technology and improved system availability. In providing the necessary test and support functions of reset and manual control, the Processor Controller exploits the level-sensitive scan design capability of the 3081 to obtain read/write access to all latches and arrays. This design approach, coupled with significantly expanded capabilities for error recovery, configuration control, and diagnostics, has significantly affected the structure and capabilities of the Processor Controller.

Thermal Conduction Module: A High-Performance Multilayer Ceramic Package *by A. J. Blodgett and D. R. Barbour, p. 30.* Innovations in package design coupled with major advances in multilayer ceramic (MLC) technology provide a high-performance LSI package for the IBM 3081 Processor Unit. The thermal conduction module (TCM) utilizes a 90 × 90-mm MLC substrate to interconnect up to 118 LSI devices. The substrate, which typically contains 130 m of impedance-controlled wiring, provides an array of 121 pads for solder connections to each device and an array of 1800 pins for interconnection with the next-level package. A unique thermal design provides a cooling capacity of up to 300 W. This paper describes the TCM design and outlines the processes for fabrication of these modules. The TCM is compared to prior technologies to illustrate the improvements in packaging density, reliability, and performance.

A New Set of Printed-Circuit Technologies for the IBM 3081 Processor Unit *by D. P. Seraphim, p. 37.* A new set of printed-circuit technologies have been developed which permit construction of printed-circuit panels with several kilometers of controlled-impedance interconnections. Communications between internal layers of signal planes are achieved through small plated vias (drilled with a laser), while plated through-holes are used for the logic service terminals for cable terminations and module terminals. The panels are the largest currently known in the industry, 600 × 700 mm, and have the most layers, 20. This paper describes new LSI package designs which are achievable with the exceptional versatility that the new technologies provide. These technologies encompass vacuum lamination, electroless plating, photosensitive dielectric, laser drilling, automatic twisted-pair wire bonding, and other new approaches to printed circuits.

Conduction Cooling for an LSI Package: A One-Dimensional Approach *by R. C. Chu, U. P. Hwang, and R. E. Simons, p. 45.* The introduction of LSI packaging has significantly increased the number of circuits per silicon chip, and at the same time has greatly increased their heat flux density. In comparison to earlier MST (monolithic systems technology) products, the heat flux which must be removed from the new multi-chip substrates (100 or more chips) has

increased by an order of magnitude or more. This paper discusses an innovative conduction-cooling approach using He gas encapsulation which has been developed in response to the new LSI technology requirements. Background is provided on the liquid-encapsulated-module technology which preceded the new approach, and the basic challenges encountered in building a thermal bridge from individual chips to the module and cold plate are described. The underlying theory of operation is presented using one-dimensional mathematical and discrete analog models. The effects of various factors such as geometry, chip tilt, He concentration, air leakage, and materials are illustrated using these models. A thermal sensitivity analysis is performed to determine variations in junction temperatures and the contributions of the major parameters. The companion paper by Oktay and Kammerer which follows this one treats the more general "multi-dimensional" approach using numerical analysis techniques.

A Conduction-Cooled Module for High-Performance LSI Devices *by S. Oktay and H. C. Kammerer, p. 55.* The advent of LSI chip technology makes possible significantly increased performance and circuit densities by means of large-scale packaging of multiple devices on a single multi-layer ceramic (MLC) substrate. Integration at the chip and module levels has resulted in circuit densities as high as 2.5×10^7 circuits per cubic meter, with the necessity of removing heat fluxes on the order of 100 kW/m^2. This paper describes the development and implementation of a novel packaging concept which meets the stringent and highly interactive demands on cooling, reliability, and reworkability of LSI technology. These requirements resulted in an innovative packaging approach, referred to as the thermal conduction module (TCM). The TCM uses individually spring-loaded "pistons" that contact each chip with helium gas, the conducting medium for removing heat efficiently. A dismountable hermetic seal makes multiple access possible for device and substrate rework, while ensuring mechanical and environmental protection of critical components. A wide range of thermal, mechanical, and environmental experiments are described with analytical and computer models. The one-dimensional approach used in the previous paper by Chu et al. is extended to three-dimensional computer modeling. Simulations of expected chip temperature distributions in the IBM 3081 Processor Unit are discussed. Enhanced thermal performance of the advanced packaging concept for future applications is also indicated.

Model for Transient and Permanent Error-Detection and Fault-Isolation Coverage *by D. C. Bossen and M. Y. Hsiao, p. 67.* As computer technologies advance to achieve higher performance and density, intermittent failures become more dominant than solid failures, with the result that the effectiveness of any diagnostic procedure which relies on reproducing failures is greatly reduced. This problem is solved at the system level by a new strategy of dynamic error detection and fault isolation based on error checking and analysis of captured information. The model developed in this paper allows the system designer to project the dynamic error-detection and fault-isolation coverages of the system as a function of the failure rates of components and the types and placement of error checkers, which has resulted in significant improvements to both detection and isolation in the IBM 3081 Processor Unit. The model has also resulted in new probabilistic isolation strategies based on the likelihood of failures. Our experiences with this model on several IBM products, including the 3081, show good correlation between the model and practical experiments.

Automated Diagnostic Methodology for the IBM 3081 Processor Complex *by N. N. Tendolkar and R. L. Swann, p. 78.* The concepts of automated diagnostics that were developed for and that are implemented in the IBM 3081 Processor Complex are presented in this paper. Significant features of the 3081 diagnostics methodology are the capability to isolate intermittent as well as solid hardware failures, and the automatic isolation of a failure to the failing field-replaceable unit (FRU) in a high percentage of the cases. These features, which permit a considerable reduction in the time to repair a failure as compared to previous systems, are achieved by designing a machine which has a very high level of error-detection capability as well as special functions to facilitate fault isolation using Level-Sensitive Scan Design (LSSD), and which includes a Processor Controller to implement diagnostic microprograms. Intermittent failures are isolated by analyzing data captured at the detection of the error, and the analysis is concurrent with customer operations if the error is recoverable. A further improvement in the degree of isolation is achieved for solid failures by using automatically generated validation tests which detect and isolate stuck faults in the logic. The diagnostic package was designed to meet a specified value of isolation effectiveness, stated as the average number of FRUs replaced per failure. The technique used to estimate the isolation effectiveness of the diagnostic package and to evaluate proposals for improving isolation is described. Testing of the diagnostic package by hardware bugging indicates very good correlation between projected and measured effectiveness.

Design Verification System for Large-Scale LSI Designs *by M. Monachino, p. 89.* This paper describes the changing environment of large-scale hardware designs as influenced by technology advancements and the growing use of design verification in the design implementation process. The design verification methodology presented here saved some 66% from the 3081 product schedule, when compared with a schedule utilizing a conventional verification method, on almost 800,000 LSI logic circuits. The paper discusses the use of software modeling techniques to verify LSI hardware designs, methods used for deciding when modeling should be stopped and hardware can be built with sufficient assurance to permit additional verification to continue on the hardware, methods for testing the hardware as it is assembled into a very large processor complex, and the organization of the design

verification system to avoid duplicate creation of test cases for different stages of the design process. Experiences encountered in designing and verifying the 3081 system, a discussion of some shortcomings, and an endorsement of certain techniques and improvements for use in future designs are also presented.

Timing Analysis of Computer Hardware *by R. B. Hitchcock, Sr., G. L. Smith, and D. D. Cheng, p. 100.* Timing Analysis is a design automation program that assists computer design engineers in locating problem timing in a clocked, sequential machine. The program is effective for large machines because, in part, the running time is proportional to the number of circuits. This is in contrast to alternative techniques such as delay simulation, which requires large numbers of test patterns, and path tracing, which requires tracing of all paths. The output of Timing Analysis includes "slack" at each block to provide a measure of the severity of any timing problem. The program also generates standard deviations for the times so that a statistical timing design can be produced rather than a worst case approach. This system has successfully detected all but a few timing problems for the IBM 3081 Processor Unit (consisting of almost 800,000 circuits) prior to the hardware debugging of timing. The 3081 is characterized by a tight statistical timing design.

Boolean Comparison of Hardware and Flowcharts *by G. L. Smith, R. J. Bahnsen, and H. Halliwell, p. 106.* Boolean comparison is a design verification technique in which two logic networks are compared for functional equivalence using analysis rather than simulation. Boolean comparison was used on the IBM 3081 project to establish that hardware flowcharts and the detailed hardware logic design were functionally equivalent. Hardware flowcharts are a graphic form of a hardware description language which describes the logical behavior of the machine in terms of the inputs, outputs, and latches. The logical correctness of the hardware flowcharts was previously established via cycle simulation. The concepts and techniques of Boolean comparison as used on the IBM 3081 project are described.

Volume 26, Number 2, 1982

Laser-Enhanced Plating and Etching: Mechanisms and Applications *by R. J. von Gutfeld, R. E. Acosta, and L. T. Romankiw, p. 136.* We have developed experimental electroplating, electrodeless plating, and etching techniques that use a focused laser beam to define the localized plating or etching region. Enhancements in plating (etching) rates up to $\approx 10^3$ to 10^4, compared to background rates, have been observed in the region of laser irradiation. A thermal model has been developed to describe the observed effects over the entire overpotential (polarization) curve. In the low overpotential region the enhancement is dominated by the increase in the local charge-transfer kinetics due to the local increase in temperature produced by absorption of the laser energy by the cathode (anode). At higher overpotentials, in the mass-transport-limited region, the main enhancement occurs due to hydrodynamic stirring caused by the large local temperature gradients. Examples of gold, nickel, and copper electroplating are described to illustrate the value of this technique for micron-sized circuit personalization and repair. Additional examples of electroless laser-enhanced plating and exchange plating are also described.

Laser-Enhanced Chemical Etching of Solid Surfaces *by T. J. Chuang, p. 145.* The laser-enhanced chemical etching of Si, SiO_2, Ta, and Te films with halogen-containing gases excited by a pulsed CO_2 laser and a continuous-wave (cw) Ar^+ laser has been studied. Detailed measurements of the etch rates as functions of the laser frequency, the laser intensity, and the gas pressure have been performed for some of the gas-solid systems. The enhanced surface reactions have been classified into three categories: those activated by the vibrational excitation of the etchant molecules, those with radicals generated by photodissociation, and those induced by laser excitation of solid substrates. Examples which illustrate the effects of laser radiation on these surface photochemical processes are given. Achievable etch rates and spatial resolutions for the various reaction mechanisms are also examined.

Ultrafast High-Resolution Contact Lithography with Excimer Lasers *by K. Jain, C. G. Willson, and B. J. Lin, p. 151.* A new technique for speckle-free, fine-line high-speed lithography using high-power pulsed excimer lasers is described and demonstrated. Use of stimulated Raman shifting is proposed for obtaining the most desirable set of spectral lines for any resist. This permits, for the first time, the optimization of the exposure wavelengths for a given resist, rather than the reverse situation. Excellent-quality images are obtained in 1-μm-thick diazo-type photoresists such as ®AZ-2400 and a diazonaphthoquinone-®Novolak resist system by means of contact printing with a XeCl laser at 308 nm and a KrF laser at 248 nm. Resolution down to 1000 line pairs per millimeter is experimentally demonstrated. These images are comparable to state-of-the-art contact lithography obtained with conventional lamps. The major difference is that the excimer laser technique is approximately two orders of magnitude faster. Tests on reciprocity failure in several resists indicate a decrease in sensitivity by only a factor of three, despite the $\approx 10^8$ times larger power density used in the laser exposures. The possibility of photochemical reactions being different from those taking place in the case of lamp exposures is discussed in view of these results.

Image Projection with Nonlinear Optics *by M. D. Levenson and K. Chiang, p. 160.* The nonlinear optical process of conjugate wavefront generation by degenerate four-wave mixing can project images in a manner akin to conventional optical systems. The underlying physics is quite different, and the difference allows higher resolution over larger useful fields. This article reviews the basics of conjugate wavefront generation and its relationship to holography, and proposes

applications in fine-line lithography. Initial experiments confirm the predicted advantages but also point out inadequacies in present-day nonlinear optics technology.

Modification of Semiconductor Device Characteristics by Lasers *by Y. C. Kiang, J. R. Moulic, W.-K. Chu, and A. C. Yen, p. 171.* The paper discusses an emerging area of laser-semiconductor processing: the effect of laser irradiation on the electrical parameters of diffused-junction devices. Shallow diffusions of 0.2- to 0.5-μm depths have been achieved. Transistor current gains have been modified and the results agree with the theoretical analysis.

The Application of GaAlAs Lasers to High-Resolution Liquid-Crystal Projection Displays *by A. G. Dewey and J. D. Crow, p. 177.* In this review we describe a high-resolution liquid-crystal (LC) display that uses multiple GaAlAs lasers to write on the LC cell. Arrays of high-powered lasers were needed to meet the writing speed requirements of the display. A fiber optic delivery system was developed to allow high-resolution, multi-spot writing on the cell, regardless of the laser configuration used. The flexibility and unconstrained length of the fiber optics also permitted simplification in the design of the scanner system which moved the optical beams over the LC cell. Because of the proximity of the injection laser's wavelength to the visible spectrum used in the display projection system, and because of divergence of the beams in the optical delivery system, a reflective LC cell was designed. Almost total absorption of the GaAlAs laser beam was achieved. This paper discusses the thermal and optical properties of the cell and their effect on the writing characteristics and the quality of the projected image (contrast and brightness). An off-axis projection system, which was designed for use with the reflective cell, creates a high-resolution image on the screen.

Lasers in Electrophotography *by A. C. Tam and R. D. Balanson, p. 186.* We review the recent advances in the use of laser sources in the field of electrophotography. New experimental investigations of the photogeneration and transport processes utilized in electrophotography are now possible because of the short pulse duration, high peak power, high collimation, monochromaticity, and/or coherence of a laser source. Also, the increasing availability of reliable and inexpensive lasers has resulted in many new technological uses of lasers in electrophotographic applications. In addition to the continued use of HeNe and HeCd lasers, new products have been configured with GaAs lasers.

Multiple Photochemical Hole Burning in Organic Glasses and Polymers: Spectroscopy and Storage Aspects *by A. R. Gutiérrez, J. Friedrich, D. Haarer, and H. Wolfrum, p. 198.* A scheme for optical information storage using photochemical hole burning (PHB) in amorphous systems is evaluated. Limits imposed by the nature of PHB in polymers and glasses and its dependence on temperature are discussed. It is demonstrated that optical information storage can be multiplexed by a factor of 10^3 using the frequency dimension and PHB.

Integrated Optics and Raman Scattering: Molecular Orientation in Thin Polymer Films and Langmuir-Blodgett Monolayers *by J. F. Rabolt, R. Santo, N. E. Schlotter, and J. D. Swalen, p. 209.* Studies of submicron films and molecular monolayers with infrared and Raman spectroscopy have been hampered by the inability of current spectroscopic techniques to detect the minute amount of material present in these thin-film assemblies. A method for overcoming this problem by using integrated optics has been successfully demonstrated. In the case of Raman studies, the material whose spectrum was desired was made into an asymmetric slab waveguide or a composite waveguide structure in which both the optical field intensity of the in-coupled laser source and the scattering volume of the sample have been significantly increased. Using this technique we have obtained Raman spectra of thin polymer films (<80 nm) and the resonant Raman spectra of single dye monolayers (2.7 nm). Estimates of molecular orientation within the two-dimensional films have been made based on the results of polarized Raman measurements. In addition, the results of overcoating experiments illustrate the versatility and applicability of this technique to a wide variety of surface and thin-film studies.

A New Class of Materials for Holography in the Infrared *by C. Bräuchle, U. P. Wild, D. M. Burland, G. C. Bjorklund, and D. C. Alvarez, p. 217.* A new class of holographic materials is described. These materials undergo four-level two-photon photochemistry. The holograms recorded in these materials are self-developing and are not erased during reading. Furthermore, the recording process may be gated off and on by using an auxiliary incoherent light source. As a specific example of such a system, holograms formed using samples of an α-diketone dissolved in poly(cyanoacrylate) are described. Gated holograms have been recorded in this material at 752 and 1064 nm.

Holography in the IBM 3687 Supermarket Scanner *by L. D. Dickson, G. T. Sincerbox, and A. D. Wolfheimer, p. 228.* The IBM 3687 Supermarket Scanner is described, with emphasis on the holographic deflector disk used to create the scan pattern. The scanner exploits the functional advantages of holography to create a dense, multiple-focal-plane scan pattern with small spot size and large depth of field. The optical design of the holographic disk is discussed and basic disk fabrication concepts are introduced.

Laser-Induced Arcing in Cathode Ray Tubes *by R. T. Hodgson, A. B. Minn, and J. D. Rockrohr, p. 235.* Laser-induced arcing in cathode ray tubes has been used to study the effect of spontaneous internal high-voltage arcs under normal operating conditions. Ruby and Nd:YAG laser systems were compared as laser sources for the breakdown.

Various arc-suppression schemes for CRT systems were evaluated for future use.

Performance Analysis of Suspend Locks in Operating Systems *by J. Hofmann and H. Schmutz, p. 242.* Performance models of suspend locks in operating systems are developed and analyzed. Analytical expressions and algorithms for numerical results have been obtained for an arbitrary number of processors, an arbitrary number of tasks, and one suspend lock. The results are discussed and important dependencies among the major characteristic quantities such as queue length, processor speed, number of processors, dispatching overhead, and processor degradation are shown. Expressions are derived permitting the control program designer to estimate the system impact of locking during the early design phase.

Volume 26, Number 3, 1982

Cost/Performance Single-Chip Module *by D. J. Bendz, R. W. Gedney, and J. Rasile, p. 278.* The introduction of the high performance T^2L device technology for the IBM 4300 Systems and the System/38 required extensions of the 24-mm metallized ceramic module used in prior IBM systems. The extension of the metallized ceramic technology into a physically larger 28-mm module, electrically enhanced with a reference plane and with fine line (0.025-mm) wiring, is described. The reliability of this module was evaluated with particular emphasis on corrosion and metal migration in humid environments and on exposure to atmospheric contaminants.

The Thin-Film Module as a High-Performance Semiconductor Package *by C. W. Ho, D. A. Chance, C. H. Bajorek, and R. E. Acosta, p. 286.* This paper discusses a multichip module for future VLSI computer packages on which an array of silicon chips is directly attached and interconnected by high-density thin-film lossy transmission lines. Since the high-performance VLSI chips contain a large number of off-chip driver circuits which are allowed to switch simultaneously in operation, low-inductance on-module capacitors are found to be essential for stabilizing the on-module power supply. Novel on-module capacitor structures are therefore proposed, discussed, and evaluated. Material systems and processing techniques for both the thin-film interconnection lines and the capacitor structures are also briefly discussed in the paper. Development of novel defect detection and repair techniques has been identified as essential for fabricating the Thin-Film Module with practical yields.

Advanced Printed-Circuit Board Design for High-Performance Computer Applications *by R. F. Bonner, J. A. Asselta, and F. W. Haining, p. 297.* A new integrated circuit packaging structure was needed to support the new 90-mm multilayer ceramic modules, known as Thermal Conduction Modules (TCMs), used in the IBM 3081 computers. The structure developed eliminates one level of packaging (the card level) and allows up to nine TCMs to be plugged directly into a large multilayer printed-circuit board using a new zero-insertion-force connector system. The board has 18 internal circuit planes for signal and power distribution and accommodates new signal cabling, power bus, terminating resistors, decoupling capacitors, and cooling hardware, forming a packaged unit of up to a quarter million logic gates and half a million bits of memory. This paper focuses on the detailed design of the printed-circuit board and on its signal and power transmission characteristics.

High-Density Board Fabrication Techniques *by J. R. Bupp, L. N. Chellis, R. E. Ruane, and J. P. Wiley, p. 306.* A variety of construction methods can be used to form high-density printed-circuit boards. The electrical and mechanical requirements of a design strongly influence the choice of processes used to produce the finished product. The introduction of physically large high-density boards required the development of many new processes, several of which are critical to the electrical performance of the composite. In this paper the process sequence employed in the fabrication of the 3081 high-density board is described, with particular emphasis on the selection of the critical processes used in its manufacture. The 3081 registration system is also discussed, and a new method for merging through-hole location data with the data representing the location of the conducting layers is presented.

Development of Interconnection Technology for Large-Scale Integrated Circuits *by R. Babuka, G. J. Saxenmeyer, Jr., and L. K. Schultz, p. 318.* The size, cost, reliability, and serviceability of modern large data processing systems are influenced by the electrical interconnections required among the hundreds or thousands of chips they contain. The development of the complex interconnection design scheme is profoundly influenced by the interactions among the component and subassembly designs: modules, printed-circuit boards, cables, connectors, etc. The design scheme is also strongly affected by the constraints imposed by manufacturing process limitations and by environmental factors, such as cooling and corrosion. This paper deals with these interactions and constraints as encountered in the interconnection design of the IBM 3081 system and the design discipline required to deal with them.

Influence on LSI Package Wireability of Via Availability and Wiring Track Accessibility *by J. H. Koch III, W. F. Mikhail, and W. R. Heller, p. 328.* In this paper, experiments are reported which use automatic global and line-packing wiring routines, supplemented by a restricted maze runner, to evaluate the overall influences of several important physical variables upon the wireability of several logic-circuit package types. The logic circuits are contained in subpackages (e.g., modules carrying chips), which are inserted, using pins, into the carrier package in a regular array of holes otherwise available as vias interconnecting the wiring planes. Over a

range of connection counts from several hundred to several thousand, it is found that "overflows" (connections not wired by the program) amounting to as much as 10 or 15% of the wires can be substantially reduced in number by careful design. This can be done by using a sufficient number of programmable vias (greater than one per used pin) and by using a track grid ensuring maximum global track accessibility to all pins, or by a combination of both of these tactics in conjunction with suitable wiring algorithms. Some simple theoretical arguments are given which characterize the design problem in the light of the results.

Lead Reduction Among Combinatorial Logic Circuits *by W. V. Vilkelis, p. 342.* The paper provides a description of the behavior of lead reduction among combinatorial logic circuits. Methods by which one may design the lead reduction architecture for modular packaging schemes are developed. The methods are then applied to the design of the lead reduction architecture of an integrated circuit chip and a nine-chip cell.

Electrical Design of a High Speed Computer Package *by E. E. Davidson, p. 349.* A methodology for optimizing the design of an electrical packaging system for a high speed computer is described. The pertinent parameters are first defined and their sensitivities are derived so that the proper design trade-offs can ultimately be made. From this procedure, a set of rules is generated for driving a computer aided design system. Finally, there is a discussion of design optimization and circuit and package effects on machine performance.

A VLSI Bipolar Metallization Design with Three-Level Wiring and Area Array Solder Connections *by L. J. Fried, J. Havas, J. S. Lechaton, J. S. Logan, G. Paal, and P. A. Totta, p. 362.* The ability to interconnect large numbers of integrated silicon devices on a single chip has been greatly aided by a three-level wiring capability and large numbers of solderable input/output terminals on the face of the chip. This paper describes the design and process used to fabricate the interconnections on IBM's most advanced bipolar devices. Among the subjects discussed are thin film metallurgy and contacts, e-beam lithography and associated resist technology, a high temperature lift-off stencil for metal pattern definition, planarized rf sputtered SiO_2 insulation/passivation, the "zero-overlap" via hole innovation, in situ rf sputter cleaning of vias prior to metallization, and area array solder terminals.

Optimization of Indium-Lead Alloys for Controlled Collapse Chip Connection Application *by R. T. Howard, p. 372.* Indium-lead solders are used for IBM controlled collapse chip connections (C-4s) to improve fatigue life in temperature cycling for large chip applications. Using 50% In-Pb alloy, which is expensive, has posed a number of manufacturing and reliability concerns. This paper presents the results of development studies leading to the use of a low-indium solder alloy for C-4 applications. This alloy overcomes all previous concerns while exceeding the fatigue life specification of the high-indium alloy. Also described are the variables and tests used to evaluate C-4 performance of In-Pb alloys over the 5% to 50% range. Results are presented graphically and mathematically to show the improvement obtained with indium-content solders over the conventional tin-lean alloys.

Immersion Wave Soldering Process *by J. R. Getten and R. C. Senger, p. 379.* Immersion wave soldering is a new technique for soldering advanced printed circuit boards. In this process, an assembled and fixtured printed circuit board is immersed in a fluxing fluid and preheated to the soldering temperature. The fluxing fluid is glycerine activated with ethylenediamine tetra-acetic acid (EDTA). After preheating and while still immersed in the fluxing fluid, the fixtured board is passed over a tin-bismuth eutectic (42% Sn, 58% Bi) solder wave. This technique has several advantages over conventional soldering processes, including elimination of solder dross formation, improved control over solder deposition, reduced thermal shock, easier cleaning after soldering, and improved flux effectiveness.

Volume 26, Number 4, 1982

Preface *by E. J. Galli, p. 399.*

Microprocessor Implementation of Mainframe Processors by Means of Architecture Partitioning *by P. W. Agnew and A. S. Kellerman, p. 401.* The benefits of Large-Scale Integration (LSI) implementations have applied quite naturally to processors with relatively low performances and simple architectures; e.g., the one-chip microprocessors used in personal computers contain several thousand logic gates. Mainframe processors, however, have so far been limited to using logic chips that contain several hundred logic gates. The best use of LSI logic employs microprocessors to keep critical paths on chip, thus keeping pin counts and power dissipations within reasonable limits. Microprocessors have been extensively used to implement peripheral functions, such as I/O device control. However, as of this writing, a single state-of-the-art microprocessor cannot contain a mainframe processor function. Therefore, new machine organizations are needed to use today's state-of-the-art microprocessors to implement a mainframe processor. This paper examines several methods for applying LSI and microprocessors to the design of processors of increasing performance and complexity, and describes a number of specific approaches to microprocessor-based LSI implementation of System/370 processors. The most successful approaches partition the System/370 instruction set into subsets, each of which can be implemented by microcode on a special microprocessor or by programs written for an off-the-shelf microprocessor.

A Microprocessor for Signal Processing, the RSP *by F. Mintzer and A. Peled, p. 413.* Signal processing is a data processing domain that contains a diversity of applications, including speech processing, image processing, radar, sonar,

medical imaging, data communications, seismic processing, and many others. Despite the diversity of the applications, this processing domain has a very structured set of characteristics. These include real-time operation, dominance of arithmetic operations, and well-structured data flows. The Real-Time Signal Processor (RSP) is a microprocessor architecture that was created to exploit these characteristics in order to provide an expeditious and economical way to implement signal processing applications. In this paper, the organization and architecture of the RSP are described. Features of the RSP, such as the instruction pipeline and the fractional fixed-point arithmetic, which exploit the characteristics of signal processing to provide additional computational power, are emphasized. Other features, such as the powerful indexing, the saturation arithmetic, the guard bits, and the double-word-width accumulator, which add much to the processor's versatility and programmability, are also highlighted. The performance of the RSP is illustrated through examples.

Rectangular Transforms for Digital Convolution on the Research Signal Processor *by J. W. Cooley, p. 424.* The Rectangular Transform (RT) method for computing convolutions belongs to a family of Reduced Computational Complexity (RCC) algorithms. Convolution calculations by the RT method were programmed for the Research Signal Processor (RSP) and run on the RSP simulator, giving tabulations of numbers of RSP machine cycles. One of the original objectives was to see how well the original RSP architecture was suited to the RCC algorithms and to be able to make suggestions for possible changes. The results are also intended to demonstrate the efficiency of the RT convolution algorithms on a microprocessor with a limited instruction set and to show how to construct efficient RT programs for digital convolution. All results are given for the original RSP, as it was before the modification which resulted in the Real-time Signal Processor described in another paper in this issue.

Real-Time Signal Processor Software Support *by K. Davies and F. Ris, p. 431.* The Real-Time Signal Processor (RSP) is a microprocessor optimized to provide fast, cost-efficient processing for signal processing applications. In order for the RSP to become fully useful, a complete set of software support tools needed to be developed. The hardware design and software development, which took place between 1978 and 1980, resulted in many architectural features which minimized hardware complexity at the expense of programmability. This paper describes the tools that were developed and the decisions that were involved, and includes hindsight comments on what was done. Particular emphasis is placed on the most interesting aspects of the software development, i.e., how the special architectural features of the RSP were handled to make the overall hardware/software system more programmable.

Common Chip for Use in Disk and Diskette Controllers *by G. L. Dix and M. D. Brown, p. 440.* The advent of LSI technology makes common microprogrammable controllers very cost-effective today. This paper focuses on the application of microcontrollers for disk and diskette control functions and describes a custom-designed FET chip which is being developed for use in these types of controllers. The architecture, the functional organization, and the physical design of this chip are presented, and the requirement of matching a microcontroller to the application is discussed.

Physical Design of a Custom 16-Bit Microprocessor *by A. Correale, p. 446.* The physical chip design aspects of a 16-bit, single-chip, custom-macro-designed microprocessor are described. This microprocessor represents the IBM System Products Division's highest-density VLSI FET processor design to date. The chip is a complex arrangement of over 6500 VLSI circuits utilizing a state-of-the-art polysilicon-gate HMOS-1 (high-performance MOS) technology. The physical design of this chip required the use of a comprehensive methodology, from conception through completion. The methodology used in the design of the microprocessor was based on a hierarchical approach and is presented in this paper.

Design Considerations for a VLSI Microprocessor *by J. E. Campbell and J. Tahmoush, p. 454.* The machine architecture and design considerations for an interrupt-driven bipolar VLSI Microprocessor are presented. The processor is designed to a complex architecture and includes an integrated channel and a flexible storage interface. Floating-point functions are optional. A 3-ns custom bipolar technology was developed for the microprocessor, resulting in a very high circuit density package. The 50-mm four-chip air-cooled microprocessor module is packaged on a printed-circuit card with associated repowering circuits and high-speed random-access memory. Important design considerations and tradeoffs associated with the development of this machine, within specific cost, performance, reliability, and schedule objectives, are discussed. Various processor design techniques are described which minimize hardware where performance is not critical. A degree of functional parallelism is utilized, as well as timing flexibility, to attain the required performance.

Bipolar Chip Design for a VLSI Microprocessor *by K. F. Mathews and L. P. Lee, p. 464.* In this paper, a pseudo-custom approach to bipolar VLSI chip design is presented, and a hierarchical structure of logic macros assembled from building blocks is described. A strategy of placing the logic macros along with algorithmically designed PLA structures and ROS with a placement aid, and of wiring the placement with an automatic wiring program, is discussed. The paper also focuses on the implementation of this strategy in terms of technology, chip structure, and chip design methodology. In addition, chip statistics are presented and their implications are discussed.

A VLSI Design Verification Strategy *by A. S. Tran, R. A. Forsberg, and J. C. Lee, p. 475.* With the ever-increasing

density, development cost, and turn-around time of VLSI chips it becomes increasingly important to have a design verification methodology which enables first-pass chips to be fully functional. The strategy discussed in this paper exploits the best attributes of the two traditional methods of design verification (i.e., software simulation and hardware modeling). Software simulation was chosen for its capability in the area of delay analysis and early functional checking. An automatically generated nodal-equivalent hardware model was built to provide the vehicle on which exhaustive functional checking could be performed. The model also operated as early user hardware on which functions such as operating systems, I/O adapters, and a floating-point feature could be tested. A technique known as interface emulation was used on certain well-defined subsystems to facilitate a shorter verification schedule through parallel debug efforts.

A Bipolar VLSI Custom Macro Physical Design Verification Strategy *by J. F. McCabe and A. Z. Muszynski, p. 485.* The level of complexity and the turn-around time associated with the development of custom bipolar VLSI chips have defined the need for a highly structured physical and electrical design validation approach which can guarantee fully functional first-pass chips, yet be flexible enough to allow logical and physical designers the latitude necessary to achieve specified cost and performance objectives. This paper describes such a design verification strategy and its implied constraints on chip design. The rationale for comparing the logic equivalence of the high-level logical models to the low-level-device physical models is presented, a description of the hierarchical logical-to-physical and electrical checking is given, and its impact on cost and complexity is examined.

Plant Automation in a Structured Distributed System Environment *by N. C. Mescia and C. D. Woods, p. 497.* 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.

Construction of Bounded Delay Codes for Discrete Noiseless Channels *by P. A. Franaszek, p. 506.* Algorithms are described for constructing synchronous (fixed rate) codes for discrete noiseless channels where the constraints can be modeled by finite state machines. The methods yield two classes of codes with minimum delay or look-ahead.

Volume 26, Number 5, 1982

Semiconductor Manufacturing Technology at IBM *by H. Carre, R. H. Doxtator, and M. C. Duffy, p. 528.* IBM's current semiconductor manufacturing processes are summarized. Briefly outlined are highlights of salient features of semiconductor manufacturing that established IBM leadership and direction for embarking on VLSI production. Some of these features are taken from papers already published; others are from articles in this issue of the IBM Journal of Research and Development ranging from silicon crystal growth through quality control and automated logistics and dispositioning of tested product.

Evolution and Accomplishments of VLSI Yield Management at IBM *by C. H. Stapper, P. P. Castrucci, R. A. Maeder, W. E. Rowe, and R. A. Verhelst, p. 532.* The methods developed at IBM to manage and improve the yield of some of its newer FET semiconductor products are described. A number of visual inspection and electric monitoring techniques have evolved since discrete semiconductors were manufactured. The data obtained with these techniques are used in self-checking yield models to give the relative yields for all the yield components. The results are applied not only to day-to-day control of the manufacturing lines, but also in the long-range forecasting and planning of future semiconductor integrated circuit products. An example is given comparing the actual and planned yield of a 64K-bit random access memory chip as a function of time. The results show the yield enhancement that was obtained with redundant circuits and additionally with the use of partially functional products. Another example shows the decrease in fault levels over a span of more than ten years.

Oxygen Incorporation and Precipitation in Czochralski-Grown Silicon *by A. Murgai, W. J. Patrick, J. Combronde, and J. C. Felix, p. 546.* Oxygen incorporation is examined for growth of large-diameter (80–130 mm) silicon single crystals by the Czochralski method. The primary growth parameters affecting the oxygen concentration in the crystals are shown to be the crystal-melt interface position within the hot zone and the rate of crucible rotation used. Tight control of the oxygen concentration [±1.5 parts per million atomic (ppma) in the range of 25–40 ppma (ASTM)] has been reproducibly attained by programmed variation of these growth parameters. The attainable oxygen concentrations may be extended over a wider range (20–45 ppma) through slight modifications of the hot zones. Wafers from the uniform-oxygen-concentration crystals are subjected to single-step and two-step annealing procedures (700–1100°C) for oxygen precipitation studies. The rate of precipitation is shown to depend on the initial oxygen content and on the number of initially unpopulated nucleation sites present.

Metrology in Mask Manufacturing *by H. R. Rottmann, p. 553.* A major source of registration failure in microlithography was found to be due to variations in optical field sizes defined by the dies, caused by unsatisfactory focus control. Two methods for determining variations in optical field sizes are described. Both allow measurements of selected registration errors with an uncertainty of ±0.01 μm (1 σ) under manufacturing conditions using commercially available measuring microscopes. The long-term registration stability of stepped-mask exposure systems was also investigated. It is concluded that maintenance of registration over long periods of time can be improved through accurate focus control along the optical axis. Finally, it is suggested that no single set of measurements at one point in time can completely characterize a microlithographic system. Periodic monitoring of key measurable parameters during use is advisable.

Feature Size Control in IC Manufacturing *by P. Frasch and K. H. Saremski, p. 561.* Until recently, the ever-increasing demand for higher device density in integrated-circuit (IC) designs has been satisfied mainly through device and circuit design ingenuity, increased chip size, and dimensional reduction. To keep pace with dimensional reduction, control of dimensional variations has assumed a more significant role. New IC designs are very dense and performance-oriented, requiring 2-μm lithography ground rules with less than ±0.2-μm variations in circuit-feature dimensions at 3σ. Such small variations are difficult to detect and control with the present-day routine in-line-inspection optical tools. In this article, we present a description of the primary causes of dimensional variations in a typical manufacturing environment and proposals for their control.

Electron-Beam Proximity Printing—A New High-Speed Lithography Method for Submicron Structures *by H. Bohlen, J. Greschner, J. Keyser, W. Kulcke, and P. Nehmiz, p. 568.* A laboratory prototype of an electron-beam proximity printer is described which shadow-projects patterns of chip-size transmission masks onto wafers. Electron-beam transmission masks with physical holes at transparent areas have been fabricated with the smallest structures down to 0.3 μm. Experiments to replicate mask patterns were directed at demonstrating the applicability of this lithographic method to high-speed printing of repetitive patterns on wafers. Linewidth resolution and positional accuracy, as well as exposure speed, meet the requirements for micron and submicron lithography.

Optimization of Plasma Processing for Silicon-Gate FET Manufacturing Applications *by A. S. Bergendahl, S. F. Bergeron, and D. L. Harmon, p. 580.* The development and implementation of plasma processing techniques for silicon-gate FET manufacturing applications is described. Process requirements are discussed for the stripping of photoresist and the isotropic etching of thin films. A systematic approach for the optimization of these processes, involving the use of statistically designed multiparametric experimentation (MPE), is presented. This approach, in combination with the use of response-surface analysis techniques, is illustrated by examples of its application to typical processing problems. In addition, the multiparametric optimization of anisotropic etching is presented for potential plasma processing enhancements.

The Mechanism of Single-Step Liftoff with Chlorobenzene in a Diazo-Type Resist *by R. M. Halverson, M. W. MacIntyre, and W. T. Motsiff, p. 590.* The mechanism of the chlorobenzene single-step liftoff process is defined. The chlorobenzene penetrates to some depth into the resist film during the soaking cycle, extracting residual casting solvent and low-molecular-weight resin species. The chlorobenzene is subsequently removed by a rinse cycle. The penetrated layer of resist develops at a slower rate than the underlying bulk resist, producing the liftoff structure.

Process Control of the Chlorobenzene Single-Step Liftoff Process with a Diazo-Type Resist *by G. G. Collins and C. W. Halsted, p. 596.* Introduction of the single-step chlorobenzene liftoff process using a diazo-type resist to manufacturing lines produced problems not encountered during development and pilot-line work. Variances in the structure of the photoresist liftoff image are the result of complex interactions among exposure, chlorobenzene soaking, development, and post-application baking conditions. Effects produced by these variables can be controlled by monitoring the linewidth, overhang, and height of the liftoff resist structure using a scanning electron microscope (SEM). Loss of resist thickness during the chlorobenzene soak is used instead of penetration, as measured on SEM photographs, to monitor the soaking process. Data are presented on the creation and stability of the overhang structure, the process controls required to achieve that stability, and the interactions among the process variables. The process, as practiced in a manufacturing mode, was found to have greatest reproducibility at low exposure, with a combination of long soaking times and high post-application baking temperatures.

Semiconductor Final Test Logistics and Product Dispositioning Systems *by R. M. Burgess, K. B. Koens, and E. M. Pignetti, Jr., p. 605.* As product lines at the IBM East Fishkill plant have expanded in the last few years, and as the number of technologies and semiconductor wafer and module volumes have increased, more sophisticated software systems were introduced which not only drove test times downward, but also reduced the development time previously required to accommodate testing new technologies. This paper discusses two such systems developed for the bipolar semiconductor line. They are a final test logistics system which provides for computerized tracking of each device in the production line through both testing and diagnostic analysis, and an automatic product dispositioning system which immediately identifies shippable product batches after flow through the test sector or schedules the batches for additional diagnostic analysis.

Quality and Reliability Assurance Systems in IBM Semiconductor Manufacturing *by E. H. Melan, R. T. Curtis, J. K. Ho, J. G. Koens, and G. A. Snyder, p. 613.* Soon after semiconductor manufacturing began it was realized that classical process control techniques were needed for the control of quality, reliability, and yield. The discovery and control of yield, quality, and reliability detractors have been pursued continually by IBM manufacturing engineers ever since, and the resulting evolution of process control techniques has grown into a highly disciplined state. Inferential methods were added later to augment the classic techniques. This paper, in addition to providing a brief overview of semiconductor manufacturing control techniques and placing them into historical perspective, discusses a method of feed-forward control based on statistical distributions which is used in the VLSI FET memory device line. This is followed by a description of a process profile technique which is used in bipolar logic manufacturing. The importance of the system aspects in both techniques is emphasized.

Regenerative Simulation of Networks of Queues with General Service Times: Passage Through Subnetworks *by G. S. Shedler and J. Southard, p. 625.* A linear "job stack," an enumeration by service center and job class of all the jobs, is an appropriate state vector for simulation of closed networks of queues with priorities among job classes. Using a representation of the job stack process as an irreducible generalized semi-Markov process, we develop a regenerative simulation method for passage times in networks with general service times. Our estimation procedure avoids Cox-phase representation of general service time distributions and is applicable to networks with "single states" for passage times. Based on a single simulation run, the procedure provides point estimates and confidence intervals for characteristics of limiting passage times.

Volume 26, Number 6, 1982

Document Analysis System *by K. Y. Wong, R. G. Casey, and F. M. Wahl, p. 647.* This paper outlines the requirements and components for a proposed Document Analysis System, which assists a user in encoding printed documents for computer processing. Several critical functions have been investigated and the technical approaches are discussed. The first is the segmentation and classification of digitized printed documents into regions of text and images. A nonlinear, run-length smoothing algorithm has been used for this purpose. By using the regular features of text lines, a linear adaptive classification scheme discriminates text regions from others. The second technique studied is an adaptive approach to the recognition of the hundreds of font styles and sizes that can occur on printed documents. A preclassifier is constructed during the input process and used to speed up a well-known pattern-matching method for clustering characters from an arbitrary print source into a small sample of prototypes. Experimental results are included.

Automatic Scaling of Digital Print Fonts *by R. G. Casey, T. D. Friedman, and K. Y. Wong, p. 657.* New raster-based printers form character patterns using carefully designed matrices of dots. It is desirable to be able to use fonts designed for one printer on a different machine, but to do so the dot matrix patterns should first be scaled to the second printer's resolution. If the scaling is carried out as a simple interpolation, however, severe degradation in the appearance of the characters may occur. A new algorithm reduces such degradation by recognizing attributes associated with print character quality in the original patterns and then correcting the scaled patterns in order to maintain those attributes. Attributes that are detected and preserved during scaling include local and global symmetries, stroke width, sharpness of corners, and smoothness of contour. The method has been used both to scale low-resolution fonts to a finer representation and to reduce the scale of high-resolution photocomposer fonts for output on an office-type printer.

Analysis of Linear Interpolation Schemes for Bi-Level Image Applications *by I. E. Abdou and K. Y. Wong, p. 667.* In the office, it is often necessary to scan a picture at a certain resolution and then reproduce it at a different (usually higher) resolution. This conversion can be achieved by interpolating the scanned signal between the sample intervals. This paper discusses a class of linear interpolating methods based on resampling polynomial functions. In addition, we introduce new methods to compare the performance of these interpolating schemes. The signal models used are one-dimensional step and pulse functions. These bi-level models are sufficient to describe many black/white documents. The performance of the linear interpolators is determined by evaluating their accuracy in reconstructing the original bi-level signal. The analysis considers the effects of the coarse scan and fine print intervals as well as the quantization effects. Experiments using the IEEE facsimile chart as input verify the analytical findings. The results show the advantage of using odd-order polynomials, such as the first order and TRW cubic. Also, we discuss the relationship between the interpolating ratio and the number of quantization levels needed to represent the scanned signal.

Word Autocorrelation Redundancy Match (WARM) Technology *by N. F. Brickman and W. S. Rosenbaum, p. 681.* Word Autocorrelation Redundancy Match (WARM) is an intelligent facsimile technology which compresses the image of textual documents at nominally 145:1 by use of complex symbol matching on both the word and character level. At the word level, the complex symbol match rate is enhanced by the redundancy of the word image. This creates a unique image compression capability that allows a document to be scanned for the 150 most common words, which make up roughly 50% of the text by usage, and upon their match the words are replaced for storage/transmission by a word identification number. The remaining text is scanned to achieve compaction at the character level and compared to both a previously stored library and a dynamically built library of complex symbol

(character) shapes. Applying the complex symbol matching approach at both the word and character levels results in greater efficiency than is achievable by state of the art CCITT methods.

Digital Halftoning of Images *by D. Anastassiou and K. S. Pennington, p. 687.* Most printers and some display devices are bilevel (black or white) and therefore not capable of reproducing continuous tone pictures. Digital halftoning algorithms transform digital gray scale images into bilevel ones which give the appearance of containing various shades of gray. A halftoning algorithm is presented in which novel concepts are combined resulting in an output image in which moiré patterns are suppressed and, at the same time, the edges are enhanced. Various other artifacts associated with the halftoning process, such as contouring due to coarse quantization or to textural changes, are also absent from the output images in the proposed scheme. The algorithm separates the image into many small clusters which are processed independently and, therefore, it is capable of parallel implementation.

An Improved Segmentation and Coding Algorithm for Binary and Nonbinary Images *by P.-E. Danielsson, p. 698.* This paper presents a new segmentation and coding algorithm for nonbinary images. The algorithm performs contour coding of regions of equally valued and connected pixels. It consists of two distinct phases: raster scanning and border following. In this sense it is similar to algorithms presented by Kruse. However, the algorithm of this paper is considerably improved since it correctly segments truly nonbinary images. The basic idea of the algorithm is to "coat" (color, label) the borders (the cracks) between the regions from both sides in two separate border-following procedures called island following and object following. Thus, all adjacencies between the objects are systematically explored and noted. Furthermore, the raster scanner, which exhaustively searches the image for new regions, can easily determine from existing/nonexisting coating which boundaries have been traced out and which have not. The algorithm can be considerably simplified for the binary image case.

Reduced Data Re-Order Complexity Properties of Polynomial Transform 2D Convolution and Fourier Transform Methods *by T. A. Kriz, p. 708.* This paper presents new results concerning the matrix data re-order requirements of polynomial-transform-based 2D convolution and 2D Fourier Transform methods which can be employed in digital processing of images and other 2D problems. The results indicate that several power-of-2 length-modified ring polynomial transform methods developed by Nussbaumer allow the total avoidance of the row-column matrix transpose commonly encountered in other algorithmic approaches, while also providing a number of other computational advantages. It is demonstrated that this property can be the source of significantly improved throughput on a number of existing data processing structures. An execution time comparison with an efficient Fast Fourier Transform algorithm base is made assuming the use of general register architecture and array processor units. It is also assumed that one makes use of recently developed efficient matrix transpose methods by Eklundh and Ari to support 2D FFT data re-order requirements. These comparisons demonstrate a two to four times throughput improvement for the use of the polynomial transform method in place of the 2D FFT approach to circularly convolve or generate 2D Fourier transforms for large 2D fields in the range 1024×1024 to 8192×8192.

Importance of Higher-Order Components to Multispectral Classification *by J. V. Dave, R. Bernstein, and H. G. Kolsky, p. 715.* A Landsat multispectral image was combined with the corresponding digital terrain elevation data to study several information extraction procedures. Principal component and limited multispectral classification procedures were conducted on 1024×1024 four-band Landsat and five-band (Landsat plus terrain data) images, and color composites as well as quantitative information were generated. Selected results of this preliminary investigation confirm the usefulness of the principal component analysis in a qualitative presentation of the multi-band data and its association with a significant reduction in dimensionality. However, unlike some other investigators, we found that the full dimensionality must be retained when the information content of the data has to be preserved quantitatively.

Some Experiments in Image Vectorization *by J. Jimenez and J. L. Navalon, p. 724.* The application of vectorization algorithms to digital images derived from natural scenes is discussed. It is argued that the fractal nature of these scenes precludes some of the savings in storage expected from vector over raster representation, although considerable savings still result. Experimental results are given. Algorithms for contour following, line thinning, and polygonal approximation well adapted to complex images are presented. Finally, the Map Manipulation System, an experimental program package designed to explore the interaction between vector and raster information, is described briefly.

Digital Multi-Image Analysis: Application to the Quantification of Rock Microfractography *by L. Montoto, p. 735.* The microfissuration of a rock sample is analyzed using a multi-image formed of micrographs which were obtained under fluorescence and polarizing microscopy of the same sample area. Image analysis methods are applied to obtain descriptions of each type of picture, one showing the microfissure network and the other the texture of the rock. Descriptions in the form of tables of coordinates are used to quantify the features contained in the pictures. Finally, it is shown that relationships between these descriptions can result in the integration of the available information, providing more knowledge about microfissuration in the sample, including characterization and quantification of microcrack types according to their position with respect to the texture of the rock.

Backscatter and Attenuation Imaging from Ultrasonic Scanning in Medicine *by E. J. Farrell, p. 746.* Images of backscatter, attenuation, and frequency dependence of attenuation are obtained based on a three-parameter model and computation techniques described in this paper. There are several critical sources of error: backscatter speckle, beamwidth distortion, and cross-coupling artifacts between attenuation and backscatter. An iterative method of imaging and filtering is developed which effectively reduces these errors. Stability of the numerical solution involving the large number of unknowns is obtained by image iteration as opposed to parameter iteration along individual transmitter rays. This method incorporates three basic functional aspects: (1) multiple scans to reduce speckle, beamwidth distortions, and certain cross-coupling artifacts, (2) pre-image filtering to decrease beam distortions and post-image filtering to reduce cross-coupling artifacts, and (3) proper sequencing of image reconstruction and filtering. Backscatter images formed by this image iteration method are significantly superior to standard B-scan images. Further, the image iteration method yields three images of the same scan field. The present investigation is based on simulated echo data from cyst-like and complex targets.

An Algorithm for Separating Patterns by Ellipsoids *by E. R. Barnes, p. 759.* We give an algorithm for finding the ellipsoid of least volume containing a set of points in a finite-dimensional Euclidean space. Such ellipsoids have been proposed for separating patterns in a feature space.

Cursive Script Recognition by Elastic Matching *by C. C. Tappert, p. 765.* Dynamic programming has been found useful for performing nonlinear time warping for matching patterns in automatic speech recognition. Here, this technique is applied to the problem of recognizing cursive script. The parameters used in the matching are derived from time sequences of x-y coordinate data of words handwritten on an electronic tablet. Chosen for their properties of invariance with respect to size and translation of the writing, these parameters are found particularly suitable for the elastic matching technique. A salient feature of the recognition system is the establishment, in a training procedure, of prototypes by each writer using the system. In this manner, the system is tailored to the user. Processing is performed on a word-by-word basis after the writing is separated into words. Using prototypes for each letter, the matching procedure allows any letter to follow any letter and finds the letter sequence which best fits the unknown word. A major advantage of this procedure is that it combines letter segmentation and recognition in one operation by, in essence, evaluating recognition at all possible segmentations, thus avoiding the usual segmentation-then-recognition philosophy. Results on cursive writing are presented where the alphabet is restricted to the lower-case letters. Letter recognition accuracy is over 95 percent for each of three writers.

Volume 27, Number 1, 1983

Product Quality Level Monitoring and Control for Logic Chips and Modules *by D. S. Cleverley, p. 4.* In the manufacture of chips and modules, it is important to minimize defects in order to maximize quality levels and product reliability at each level of product assembly (chips, module, card, system). These objectives are best achieved by controlling defects through manufacturing process controls and testing at the lowest possible level of assembly. Defective product remaining after test and inspection must be repaired or discarded. The ability to detect and reduce or eliminate these defects is crucial to ensuring maximum product quality. The amount of such defective product is typically described quantitatively in terms of statistical sampling plans. The problem with such approaches is that the absolute defect level is imprecisely defined. This paper defines an absolute number for such product defects, which we will call "product quality level" (PQL). PQL categories found in logic chips and modules after completion of electrical testing are described and a methodology for the monitoring and control of the PQL in chips is presented. The impact of chip defects on module, card, and system performances is discussed with the aid of examples. By using the described comprehensive design, process control, testing, and user-feedback approach at each assembly level, final product can be manufactured with the lowest possible level of defects that must then be repaired at the machine level.

Multi-Layer Ceramics Manufacturing *by W. G. Burger and C. W. Weigel, p. 11.* Multi-layer ceramics (MLC) is an advanced packaging technology developed by IBM and used in the IBM 4300 and IBM 3081 processor models to significantly improve the module circuit density, reliability, and performance over those of previous packages. A key element in this package is the MLC substrate, which is capable of supporting more than 100 semiconductor chips in some design uses. The MLC manufacturing facility requires that sophisticated processing and inspection/test equipment operate within a hierarchically integrated system. This paper provides an overview of the manufacturing process that has been implemented to satisfy the high-volume manufacturing needs of this complex package. It also describes the data handling systems.

Precise Numerical Control for the Thermal Conduction Module *by M. A. Sanborn, p. 20.* Automated manufacturing processes have been developed and put into practice for the high-volume production of the thermal conduction module, the high-density circuit package used in the IBM 3081 processor models. The very precise work required on the ceramic surface required the solution of many problems. The small but significant residual distortions (warpage) which resulted from firing the multi-layer ceramic substrates made it difficult to locate tooling precisely for subsequent processes. A unique scheme for measuring the module and precalculating numerical control data has made it possible to achieve full automation. The nature of the problem, the attempts at a simple solution, and the algorithms finally used for numerical control data calculation are presented.

Multi-Chip Module Test and Diagnostic Methodology *by J. J. Curtin and J. A. Waicukauski, p. 27.* The development of a manufacturing test and diagnostic methodology for multi-chip modules as used in the IBM 4300 processor models involves determining the most attractive compromise among a number of conflicting factors: a) high test coverage, b) high diagnostic resolution, c) test generation, d) test equipment, and e) test application and diagnosis. This paper describes a set of solutions which were developed to create a high-volume, low-cost manufacturing test operation for the product in question. This paper examines the role of the testing methodology in productivity and product quality, details the diagnostic approach chosen, and provides an example of the overall manufacturing system performance achieved by analyzing a large module production sample.

The LT1280 for Through-the-Pins Testing of the Thermal Conduction Module *by R. L. Pierson and T. B. Williams, p. 35.* Testing the thermal conduction module (TCM), the high-density field-replaceable unit (FRU) used in the IBM 3081 processor models, and diagnosing the faults encountered to a minimal repairable set of entities posed a new problem for the IBM engineers. The requirement and the economic necessity of thoroughly exercising the entire TCM logic and random access memory (RAM) array structure through the input/output pins of the TCM are discussed. This is followed by a description of the test system alternatives and the LT1280 [logic tester having 1280 input/output (I/O) pins] as the selected TCM manufacturing test system. The TCM logic density and high I/O count required new concepts of test system organization, size, and complexity to achieve a test and diagnostic system with high flexibility and high throughput capability.

Failure Diagnosis on the LT1280 *by P. L. Barry, p. 41.* The high-density circuitry and I/O pin population of the thermal conduction module (TCM), the VLSI package used in the IBM 3081 processor models, dictates that there be a precise and cost-effective method of detecting and diagnosing TCM defects. This paper describes the challenge faced in testing the logic and random access memories of the TCM and the diagnostic approach used in the LT1280 test system for testing the TCM through its I/O pins. The generation and application of tests are discussed, and the automated multiple-defect diagnostic (AMDD) algorithm is presented in detail.

Computer-Controlled Optical Testing of High-Density Printed-Circuit Boards *by M. A. West, S. M. DeFoster, E. C. Baldwin, and R. A. Ziegler, p. 50.* The increased densities of multi-layer printed-circuit boards have required development of unique approaches to product testing. An optical automatic inspection system developed to test interplanes of the printed-circuit board used in the IBM 3081

processor (TCM board) is described. This system scans the features of the product and locates surface defects of 25.4 μm (1 mil) or larger through changes in reflectivity. It is capable of finding over 90% of the errors on subtractively plated printed-circuit interplanes, as well as shorts and opens of 50 μm or more on glass masters of printed-circuit boards. Alternative approaches to the inspection problem are discussed, together with the technical trade-offs which were made that led to the final system configuration. The tester theory and some hardware/software trade-offs are also covered.

A Model for the Prediction of Assembly, Rework, and Test Yields *by B. J. Dooley, p. 59.* The increase in density and complexity of computer components used as field-replaceable units of the IBM 3081 processors has required more sophisticated and capital-intense manufacturing and test lines than heretofore seen. To help in the manufacturing planning effort, a simulation model based on the concepts described in this paper was implemented as a means for studying the behavior of different assembly/test methodologies and to predict the throughput capabilities of various manufacturing configurations. By adopting a different simulation philosophy, the model was greatly simplified, and by taking advantage of the matrix processing features of APL, modeling changes required due to procedure or test-equipment changes could easily be implemented. This paper reviews the philosophy of the simulation of computer part assembly, takes a new look at modeling, and describes the architectural concepts embodied in the implementing program. It also details some of the more important concepts needed to complete the simulator.

Technologies for Network Architecture and Implementation *by F. D. Smith and C. H. West, p. 68.* Systems Network Architecture (SNA) provides a framework for constructing networks of distributed processors and terminals. This paper discusses some of the fundamental properties of network architectures such as SNA, and the evolution of formal descriptive methods that provide precise, complete definitions of the architecture. This has culminated in the development of a programming language, Format and Protocol Language (FAPL), tailored for programming a reference model or meta-implementation of an SNA node. In this form, the architecture specification is itself machine-executable. This property has led to new software technologies that improve quality and productivity in the processes for developing a network architecture and the product implementations derived from it. Automated protocol validation provides the tool necessary to ensure a correct and internally consistent definition of the architecture. This definition can then be used as a standard for testing products to determine compliance with the architecture. Direct implementation of network software by compiling the meta-implementation program is another emerging technology. This paper reviews the current state of work in these areas.

***IBM Journal of Research and Development*: Information for Authors** *by H. Freitag, p. 89.* This paper combines background information about the *IBM Journal of Research and Development* with guidelines for the preparation of *Journal* manuscripts. The purpose is to acquaint authors with the *Journal* as a primary, professional publication and to present suggestions to ease the work of author and editor in preparing clear, concise, and useful manuscripts.

Volume 27, Number 2, 1983

Arithmetic Codes for Constrained Channels *by G. N. N. Martin, G. G. Langdon, Jr., and S. J. P. Todd, p. 94.* Arithmetic codes have been studied in the context of *compression coding*, i.e., transformations to code strings which take up less storage space or require less transmission time over a communications link. Another application of coding theory is that of noiseless *channel coding*, where constraints on strings in the channel symbol alphabet prevent an obvious mapping of data strings to channel strings. An interesting duality exists between compression coding and channel coding. The *source* alphabet and *code* alphabet of a compression system correspond, respectively, to the *channel* alphabet and *data* alphabet of a constrained channel system. The decodability criterion of compression codes corresponds to the representability criterion of constrained channel codes, the generalized Kraft Inequality has a dual inequality due to the senior author.

A General Fixed Rate Arithmetic Coding Method for Constrained Channels *by S. J. P. Todd, G. G. Langdon, Jr., and G. N. N. Martin, p. 107.* This paper extends the result of earlier work on the application of arithmetic codes to the constrained channel problem. We specifically present a general length-based fixed rate implementation technique which performs the arithmetic coding recursions during each channel time unit. This technique is superior to an earlier unpublished code for general constrained channels. The approach permits the design of codes for sophisticated channel constraints.

Design Issues and Architecture of HACIENDA, an Experimental Image Processing System *by P. Franchi, J. Gonzalez, P. Mantey, C. Paoli, A. Parolo, and J. Simmons, p. 116.* The paper provides the rationale for the architecture and design of a color image display and processing system called HACIENDA. The system was heavily influenced by one of its most important intended applications, the processing of LANDSAT data, including that to be provided by LANDSAT D. Also considered in the paper are the trade-offs involved in making the system suitable for a broader range of image processing work without unduly adding to cost or complexity.

The Perceptual Color Space of Digital Image Display Terminals *by A. Santisteban, p. 127.* The geometric properties of the set of colors that can be displayed on the TV

monitor of a digital image processing terminal are discussed in the framework of the Commission Internationale de l'Eclairage (CIE) 1976 ($L^* u^* v^*$) color space. Quantitative results are presented for the HACIENDA image processing system. The use of lightness, chroma, and hue angle for the representation of multi-band images is briefly discussed.

A Fair Carpool Scheduling Algorithm *by R. Fagin and J. H. Williams, p. 133.* We present a simple carpool scheduling algorithm in which no penalty is assessed to a carpool member who does not ride on any given day. The algorithm is shown to be fair, in a certain reasonable sense. The amount of bookkeeping grows only linearly with the number of carpool members.

A "Zero-Time" VLSI Sorter *by G. Miranker, L. Tang, and C.-K. Wong, p. 140.* A hardware sorter suitable for VLSI implementation is proposed. It operates in a parallel and pipelined fashion, with the actual sorting time absorbed by the input/output time. A detailed VLSI implementation is described which has a very favorable device count compared to existing static RAM.

OYSTER: A Study of Integrated Circuits as Three-Dimensional Structures *by G. M. Koppelman and M. A. Wesley, p. 149.* This paper presents a design for a software system (OYSTER) for the parametric simulation and analysis of the fabrication steps of very large scale integrated circuit devices. The system is based on a solid geometric modeling approach in which the component parts of an integrated circuit are represented at any step as three-dimensional solid objects in a geometric data base. The simulation of a fabrication step transforms the data base representation of the geometry and the relations among component parts from their state before the step to their state after the step. At any step, and particularly after the final step, the component parts may be analyzed automatically to determine geometric, mechanical, thermal, and electrical properties. Statistical effects may be incorporated to allow investigation of alignment tolerance build-up and yield. A prototype study is described in which an existing geometric modeling system is used to transform a set of planar masks for an FET device through 28 process steps into 3-D models which are used to compute device capacitances.

Fractal Nature of Software-Cache Interaction *by J. Voldman, B. Mandelbrot, L. W. Hoevel, J. Knight, and P. Rosenfeld, p. 164.* This paper uses fractals to model the clustering of cache misses. The clustering of cache misses can be quantified by a single number analog to a fractional dimension, and we are intrigued by the possibility that this number can be used as a measure of software complexity. The essential intuition is that cache misses are a direct reflection of changes in locality of reference, and that complex software requires more frequent (and larger) changes in this locality than simple software. The cluster dimension provides a measure (and perhaps the basis for a model) of the intrinsic differences between workloads. In this paper, we focus on cache miss activity as a discriminate between interactive and batch environments.

Axisymmetric Motion of Radially Polarized Piezoelectric Cylinders Used in Ink Jet Printing *by N. Bugdayci, D. B. Bogy, and F. E. Talke, p. 171.* An analysis of the low frequency response of piezoelectric squeeze tubes used in drop-on-demand ink jet printing is carried out. The displacements at inner and outer boundaries are determined as a function of voltage and fluid pressure. The results are used to obtain fluid pressure per applied voltage as a function of inner and outer radius. Numerical computations are carried out for PZT-5H, and results are presented in graphical form which can be used to optimize design dimensions. The effect of finite electrode thickness is also studied.

An Improved Regional Correlation Algorithm for Signature Verification Which Permits Small Speed Changes Between Handwriting Segments *by J. S. Lew, p. 181.* If two nearly coincident accelerometers on a pen axis measure orthogonal acceleration components perpendicular to this axis, then the *regional correlation algorithm* for signature verification divides these data into plausible segments, it compares each segment with a corresponding reference, and it combines the results into a global similarity index. The presently used intersegment distance permits certain natural data transformations: (1) translating a segment by small integer multiples of the sampling interval; (2) moving the pen with uniformly larger amplitude throughout a segment; and (3) rotating the pen about its axis between any two segments. We propose a new intersegment distance which permits further natural transformations: (4) translating a segment by a fraction of the sampling interval; and (5) writing at slightly different uniform speed within each segment. The new distance, like the old one, is the minimum of a certain function. We describe an algorithm which computes this minimum.

Volume 27, Number 3, 1983

System/370 Extended Architecture: Design Considerations *by A. Padegs, p. 198.* This paper reviews the overall objectives and the design considerations that led to the System/370 Extended Architecture (370-XA). It presents an overview of the differences between the System/370 and 370-XA architectures and summarizes all architectural extensions, deletions, and changes. Then it describes in more detail the extensions for interpretive execution, 31-bit addressing, protection, tracing, inter-CPU communications, and extended-precision floating-point division.

System/370 Extended Architecture: The Channel Subsystem *by R. L. Cormier, R. J. Dugan, and R. R. Guyette, p. 206.* The 370-XA channel subsystem architecture represents a significant extension of the corresponding System/370 architecture. This paper examines the need for

these extensions, discusses the important features and facilities of the new architecture, and provides comparisons with its predecessor, the System/370 channel architecture. It also describes, from an operational viewpoint, how these new concepts affect I/O processing and how they relate to the current trend toward using multiple CPUs, increasing CPU execution speed, and increasing the number of I/O attachments.

A Fault-Tolerant System Architecture for Navy Applications *by W. T. Comfort, p. 219.* This paper describes the architecture of a computer system, being designed and built for the U.S. Navy, that is expected to be the standard Navy shipboard computer for the next twenty years or so. It has a requirement for very high system reliability, which is addressed by a multiprocessor system configuration that can recover dynamically from hardware faults and support on-line repair of failed hardware elements. Successfully accomplishing this requires various types of redundant hardware elements and special system architecture features, as well as intelligent fault-recovery software. This also requires that the application programs be designed to participate fully in the recovery and reconfiguration process. This paper presents the overall system architecture and discusses a number of significant new features designed to support fault-tolerant operation, including a duplex control bus, a computer interconnection system, a technique for remote diagnostics, a single-button maintenance procedure, and special fault-handling software.

The 801 Minicomputer *by G. Radin, p. 237.* This paper provides an overview of an experimental system developed at the IBM Thomas J. Watson Research Center. It consists of a running hardware prototype, a control program, and an optimizing compiler. The basic concepts underlying the system are discussed, as are the performance characteristics of the prototype. In particular, three principles are examined: (1) system orientation towards the pervasive use of high-level language programming and a sophisticated compiler, (2) a primitive instruction set which can be completely hard-wired, and (3) storage hierarchy and I/O organization to enable the CPU to execute an instruction at almost every cycle.

Integration of Machine Organization and Control Program Design—Review and Direction *by G. S. Rao and P. L. Rosenfeld, p. 247.* This paper discusses the relationship between machine organization and control program design in high-end commercial computer systems. The criterion is cost/performance, subject to achieving an acceptable performance level. A brief discussion of the environment expected for the design and operation of high-end commercial computer systems is outlined, followed by a discussion of machine organization techniques which are classified and reviewed to permit a qualitative evaluation of the degree to which control program intent is exploited in machine organization. The thesis is developed next, using a hierarchical model which illustrates the contention that architecture has acted as a barrier to communication between the control program and machine organization. Examples of techniques that exploit knowledge of the intent of the control program and comments on the methodology that might be used to investigate such techniques follow. Directions for further research are then proposed.

Documenting a Computer Architecture *by R. E. Wright, p. 257.* The documentation of a computer architecture requires that special conditions be satisfied in its preparation to ensure exactness, consistency, and completeness. One group in IBM, Central Systems Architecture, has spent many years refining procedures for producing quality documentation of this type, most notably for the System/370 architecture and the System/370 Extended Architecture (370-XA). This paper describes the steps this group performs in documenting a computer architecture, the requirements identified for the finished description, and the procedures followed to satisfy those requirements.

Random-Pattern Coverage Enhancement and Diagnosis for LSSD Logic Self-Test *by E. B. Eichelberger and E. Lindbloom, p. 265.* Embedded linear feedback shift registers can be used for logic component self-test. The issue of test coverage is addressed by circuit modification, where necessary, of random-pattern-resistant fault nodes. Also given is a procedure that supports net-level diagnosis for structured logic in the presence of random test-pattern generation and signature analysis.

Bounce and Chaotic Motion in Impact Print Hammers *by F. Hendriks, p. 273.* The basis of this paper is a lumped-parameter description of an impact printer actuator of the stored-energy type. All constants necessary to describe the actuator and the ribbon/paper pack are derived from measurements. The equations of motion are integrated both for single- and multiple-current pulse excitation. The numerical results show that for low repetition rates, each impact is distinct and independent, but at higher rates the impacts interact. The interaction manifests itself initially as flight-time and print-force variations: Strict periodicity of the actuator motion is lost, as shown in Poincaré plots for the actuator motion, and randomness sets in. At extremely high repetition rates, the actuator "hangs up" and the backstop no longer participates in the actuator dynamics. During settle-out the actuator motion is extremely sensitive to the timing of the current excitation. This fact can, in principle, be exploited to achieve extremely fast cycle times. However, without knowledge of the state of the actuator, as is commonly the case, this sensitivity is detrimental to print quality.

Digital Simulation of Magnetic Czochralski Flow Under Various Laboratory Conditions for Silicon Growth *by W. E. Langlois and K.-J. Lee, p. 281.* Previous digital simulations have suggested that an axial magnetic field in the 0.1-T (1000-Gs) range can effectively suppress convection in Czochralski growth of silicon. The present paper treats the

matter more quantitatively by investigating the convection in a variety of flow conditions corresponding to typical Czochralski growth of silicon on a laboratory scale.

Large Multi-Layer Panel-Drilling System *by W. R. Wrenner, p. 285.* This paper describes the mechanical, electrical, and control features of an advanced four-station, twelve-spindle, large printed-circuit panel-drilling system. The system was designed as a production tool to produce over 35,000 0.4-mm-diameter holes in a panel containing twenty layers of copper and nineteen layers of epoxy glass. Built around a precision four-station *x-y* positioning system with three motor-driven spindles over each work station, the system includes a minicomputer, a printer, a display station, and a controller/interface station. Drilling parameters, such as infeed, outfeed, top of stroke, bottom of stroke, and spindle revolution rates, are all computer-controlled.

Volume 27, Number 4, 1983

High-Resolution Carbon-13 NMR of Polymers in the Solid State *by J. R. Lyerla and C. S. Yannoni, p. 302.* This paper reviews the salient features of the methods used for resolution and sensitivity enhancement in ^{13}C NMR spectra of solids. Performance characteristics of the unique variable-temperature cross-polarization magic angle spinning (VT-CPMAS) apparatus developed in our laboratory are given. We present results of structural studies on fluoropolymers using this technique, and discuss the issue of intensities in cross-polarized spectra. In addition, results of an ongoing VT-CPMAS study of polypropylene dynamics are given.

Magic Angle Spinning NMR of Conducting Polymers *by T. C. Clarke, J. C. Scott, and G. B. Street, p. 313.* Because of their typically intractable nature, conducting polymers as a class of materials have proved particularly difficult to characterize by the conventional techniques of polymer analysis. We present here examples of the application of a powerful new tool, cross-polarization magic angle spinning (CPMAS) ^{13}C NMR spectroscopy, to the investigation of the structure and reactions of the conducting polymers polyacetylene and polypyrrole.

Polymer Structure Determination Using Electron Diffraction Techniques *by R. H. Geiss, G. B. Street, W. Volksen, and J. Economy, p. 321.* The crystallographic structure of organic crystals is most commonly studied using x-ray diffraction (XRD) techniques. Unfortunately, rather large crystals, at least 10^6 μm^3, are required for XRD analysis, and it is often quite difficult and sometimes impossible to prepare such large crystals. On the other hand, electron diffraction techniques, although not nearly as precise as XRD, do afford the capability of studying much smaller crystals. The minimum size for electron diffraction is on the order of 10^{-3} μm^3 (0.1 μm^2 area by 0.01 μm thick). Since most polymer crystals are very sensitive to radiation damage caused by the beam in the electron microscope, special precautions must be taken to minimize beam damage to the specimen. Our approach to minimizing radiation damage, while still obtaining usable diffraction data, is described in terms of using the condenser-objective lens optics of the Philips 301 S(TEM) electron microscope. Three examples of the application of electron diffraction structure analysis are given. These include the structures of halogenated polysulfur nitride $(SN)_x$, neutral α,α'-polypyrrole, and poly(p-hydroxybenzoic acid) (PHBA).

Electrochemical Synthesis of Electrically Conducting Polymers from Aromatic Compounds *by J. Bargon, S. Mohmand, and R. J. Waltman, p. 330.* Amorphous, electrically conducting polymeric films can be deposited from acetonitrile solutions of specific aromatic compounds containing an appropriate electrolyte. Free-standing films peeled off a platinum electrode have electrical conductivities σ between 10^{-3} and 1 Ω^{-1}-cm^{-1}. Polymers resembling the better-known polypyrrole have been obtained from benzenoid, nonbenzenoid, and heterocyclic monomers. A few characteristic examples are discussed to highlight the characterization efforts. Also discussed are the chemical investigations which have provided some insight into the mechanism of formation and the structure of these conducting polymers.

Mechanical Properties of Electrochemically Prepared Polypyrrole Films *by A. F. Diaz and B. Hall, p. 342.* The mechanical and conducting properties of free-standing films of polypyrrole toluenesulfonate depend on the preparation conditions. Films prepared in aqueous ethylene glycol solvent mixtures show a variation of two orders of magnitude in the conductivity and of a factor of three in the tensile strength.

Electron Microscopy of Carbon-Loaded Polymers *by V. E. Hanchett and R. H. Geiss, p. 348.* It is very difficult to fully characterize the spatial distribution of particles in carbon-loaded polymers using traditional methods. One approach has been to cross-section the polymers, using ultramicrotomy techniques, to a thickness of 80–100 nm. These sections are then examined in a transmission electron microscope (TEM). Unfortunately, the information on particle dispersion obtained from such images is essentially two-dimensional in nature, and therefore does not lend itself easily to a three-dimensional interpretation. By increasing the thickness of the cross sections to 0.5 μm or more, one is able to utilize stereoscopic imaging techniques and obtain a three-dimensional image of the carbon particle dispersion. In this way, the characteristic dispersion of the carbon particles may be fully evaluated along a particular direction in the polymer film. Examples are given of the three-dimensional analyses of polymer films containing different carbon-particle loadings.

Color Display and Interactive Interpretation of Three-Dimensional Data *by E. J. Farrell, p. 356.* Three-dimensional results from engineering and scientific

computations often involve the display and interpretation of a large volume of complex data. A method is developed for color display of 3D data with several interactive options to facilitate interpretation. The method is based on representing points whose values fall within a specified range as a single hue. An image is formed by overlaying successive 2D frames with increasing hue lightness towards the front. Interactive options to aid interpretations are viewpoint, contour lines, multiple range display, slicing, veiled surfaces, and stereo image pairs. The display method is successfully applied to several types of data. The overall structure and variations of the 3D data are observable, as well as transients which may be overlooked in a large input data set.

Row-by-Row Dynamic Image Analysis of a Matrix of Scanned Points *by D. McAuley, p. 367.* A practical image analysis algorithm is described that incorporates features to permit it to carry out its analysis on a row-by-row basis, in contrast to a full image matrix basis. A significant storage saving is realized, since the total image to be scanned is substantially larger than that used by this method. The dynamic build-up of the image on a scan-by-scan or row-by-row basis is carried out by a series of connectivity, pivot, chain set-up, and chaining algorithms used in conjunction with a dynamic memory allocation strategy.

A Hybrid Optical-Digital Image Processing Method for Surface Inspection *by F. Wahl, S. So, and K. Wong, p. 376.* A hybrid measurement technique is proposed for high-precision surface inspection. The technique uses an interferometer to image microscopic surface defects. In order to quantify the degree of various surface defects, the interferograms are scanned, digitized, and subsequently converted to a binary image by using an adaptive thresholding technique which takes into account the inhomogeneity of the imaging system. A new misalignment measure for binary patterns identifies the "straightness" of the fringe lines. It is shown that the resulting percentages of misaligned picture elements conform fairly well with the degree of various surface defects.

A Processor-Based OCR System *by R. G. Casey and C. R. Jih, p. 386.* A low-cost optical character recognition (OCR) system can be realized by means of a document scanner connected to a CPU through an interface. The interface performs elementary image processing functions, such as noise filtering and thresholding of the video image from the scanner. The processor receives a binary image of the document, formats the image into individual character patterns, and classifies the patterns one-by-one. A CPU implementation is highly flexible and avoids much of the development and manufacturing costs for special-purpose, parallel circuitry typically used in commercial OCR. A processor-based recognition system has been investigated for reading documents printed in fixed-pitch conventional type fonts, such as occur in routine office typing. Novel, efficient methods for tracking a print line, resolving it into individual character patterns, detecting underscores, and eliminating noise have been devised. A previously developed classification technique, based on decision trees, has been extended in order to improve reading accuracy in an environment of considerable character variation, including the possibility that documents in the same font style may be produced using quite different print technologies. The system has been tested on typical office documents, and also on artificial stress documents, obtained from a variety of typewriters.

Image Thresholding for Optical Character Recognition and Other Applications Requiring Character Image Extraction *by J. M. White and G. D. Rohrer, p. 400.* Two new, cost-effective thresholding algorithms for use in extracting binary images of characters from machine- or hand-printed documents are described. The creation of a binary representation from an analog image requires such algorithms to determine whether a point is converted into a binary one because it falls within a character stroke or a binary zero because it does not. This thresholding is a critical step in Optical Character Recognition (OCR). It is also essential for other Character Image Extraction (CIE) applications, such as the processing of machine-printed or handwritten characters from carbon copy forms or bank checks, where smudges and scenic backgrounds, for example, may have to be suppressed. The first algorithm, a nonlinear, adaptive procedure, is implemented with a minimum of hardware and is intended for many CIE applications. The second is a more aggressive approach directed toward specialized, high-volume applications which justify extra complexity.

Exclusive-OR Representations of Boolean Functions *by H. Fleisher, M. Tavel, and J. Yeager, p. 412.* With the goal of making exclusive-OR formulations of switching functions more readily available to designers for implementation in LSI and VLSI technologies, we introduce the concept of an exclusive-OR space in which an exclusive-OR normal form is defined to correspond to the conventional disjunctive normal form. A geometrical representation of exclusive-OR space is described, and its various bases are listed and discussed.

Volume 27, Number 5, 1983

Periodic Sequences with Optimal Properties for Channel Estimation and Fast Start-Up Equalization *by A. Milewski, p. 426.* The problems of fast channel estimation and fast start-up equalization in synchronous digital communication systems are considered from the viewpoint of the optimization of the training sequence to be transmitted. Various types of periodic sequences having uniform discrete power spectra are studied. Some of them are new and may be generated with data sets commonly used in phase modulation systems. As a consequence of their power spectra being flat, these sequences ensure maximum protection against noise when initial

equalizer settings are computed via channel estimates and noniterative techniques.

An Analysis of the Tolerance to Crosstalk Noise of a Pulse Width Modulation System *by D. I.-H. Park and R. D. Love, p. 432.* This paper reports the results of an investigation to determine the drive degradation caused by random noise in a pulse width modulation (PWM) system originally designed for communicating on coaxial cable but using instead twisted-pair cables. Presented are the analysis of the driver/receiver circuit, the theoretical modeling of the transmitted waveforms over the twisted-pair medium using ASTAP, and a methodology for trading off rms noise for transmission capability. The analysis of the effect of noise on drive distance as a function of allowable error rate is derived, and an example is provided. Finally, a simplified analysis technique is proposed to allow rapid calculation of approximate (yet conservative) results.

A DC-Balanced, Partitioned-Block, 8B/10B Transmission Code *by A. X. Widmer and P. A. Franaszek, p. 440.* This paper describes a byte-oriented binary transmission code and its implementation. This code is particularly well suited for high-speed local area networks and similar data links, where the information format consists of packets, variable in length, from about a dozen up to several hundred 8-bit bytes. The proposed transmission code translates each source byte into a constrained 10-bit binary sequence which has excellent performance parameters near the theoretical limits for 8B/10B codes. The maximum run length is 5 and the maximum digital sum variation is 6. A single error in the encoded bits can, at most, generate an error burst of length 5 in the decoded domain. A very simple implementation of the code has been accomplished by partitioning the coder into 5B/6B and 3B/4B subordinate coders.

Some Methods for Providing OSI Transport in SNA *by P. François and A. Potocki, p. 452.* This paper is made up of three parts: first, a review of open systems interconnection (OSI) objectives, followed by a comparison/contrasting of OSI and systems network architecture (SNA) objectives; second, a description of techniques for providing communication between programs residing in systems based on different architectures; and third, the possible application of these techniques to the four lower OSI layers, demonstrating some ways which could be used to provide connectivity between SNA end users and other heterogeneous system end users through OSI protocols.

Address-Independent Routing for Local Networks *by P. A. Franaszek, p. 464.* A routing methodology is introduced which permits messages to be propagated throughout a network without recourse to destination or origin addresses. Two classes of networks, bidirectional trees and augmented rings, are analyzed from this point of view. An optimality property is proved for the bidirectional tree, and three types of address-independent routing strategies are derived. It is shown that augmented loops, a class of structures incorporating redundant links, may be rerouted to compensate for the failure of any single node or link.

Simulation of Non-Markovian Systems *by D. L. Iglehart and G. S. Shedler, p. 472.* A generalized semi-Markov process provides a stochastic process model for a discrete-event simulation. This representation is particularly useful for non-Markovian systems where it is nontrivial to obtain recurrence properties of the underlying stochastic processes. We develop "geometric trials" arguments which can be used to obtain results on recurrence and regeneration in this setting. Such properties are needed to establish estimation procedures based on regenerative processes. Applications to modeling and simulation of ring and bus networks are discussed.

A Local Communications Network Based on Interconnected Token-Access Rings: A Tutorial *by N. C. Strole, p. 481.* Local area networks are expected to provide the communications base for interconnecting computer equipment and terminals over the next decade. The primary objective of a local area network (LAN) is to provide high-speed data transfer among a group of nodes consisting of data-processing terminals, controllers, or computers within the confines of a building or campus environment. The network should be easily accessible, extremely reliable, and extendible in both function and physical size. The rapid advances in computing and communications technology over the last two decades have led to several different transmission schemes and media types that could be used in these networks. The star/ring wiring topology with token-access control has emerged as a technology that can meet all of these objectives. The requirements of small networks with just a few nodes, as well as those of very large networks with thousands of nodes, can be achieved through this one architecture. This paper is a tutorial of the fundamental aspects of the architecture, physical components, and operation of a token-ring LAN. Particular emphasis is placed on the fault detection and isolation capabilities that are possible, as well as the aspects that allow for network expansion and growth. The role of the LAN relative to IBM's Systems Network Architecture (SNA) is also discussed.

Optimization of Relational Expressions Using A Logical Analogon *by G. A. Blaauw, A. J. W. Duijvestijn, and R. A. M. Hartmann, p. 497.* An expression applying to a relational database is optimized by mapping the expression upon set expressions which, in turn, are transformed into logical expressions. These logical expressions then are optimized, taking into account the constraints that are inherent in relational expressions and the costs of those expressions. Subsequently a reverse transformation to relational expressions is applied. The method is developed for the traditional relational operators and is applicable to a variety of cost criteria. Common subexpressions as well as redundant expressions are optimized. A new relational operation "*split*" is proposed that may be used effectively in an optimized

expression. Results obtained with a model for the optimization method are presented.

Volume 27, Number 6, 1983

System/370 Extended Architecture: Facilities for Virtual Machines *by P. H. Gum, p. 530.* This paper describes the evolution of facilities for virtual machines on IBM System/370 computers, and presents the elements of a new architectural facility designed for the virtual-machine environment. Assists that have been added to various System/370 models to support the use of virtual machines are summarized, and a general facility for this purpose which was introduced with the System/370 Extended Architecture (370-XA) is described. A new instruction of the 370-XA architecture places the machine in a specific mode in which several special capabilities are enabled. These allow the machine to provide execution in the virtual-machine environment of most of the instructions (including many privileged instructions) and most of the facilities (such as dynamic address translation) of both the System/370 and the 370-XA architectures. The major features of this new facility are individually discussed and summarized.

On Murphy's Yield Formula *by B. Meister, p. 545.* Some properties of yield are presented, and one lower and three upper bounds for yield are derived. Some of these bounds represent yield formulas already known as useful approximations. Pure Poisson statistics for defect density provides the lower bound. The upper bounds are obtained with mixtures of Poisson distributions and a formula of Price and Stapper.

Modeling of Integrated Circuit Defect Sensitivities *by C. H. Stapper, p. 549.* Until now only cursory descriptions of mathematical models for defect sensitivities of integrated circuit chips have been given in the yield literature. This paper treats the fundamentals of the defect models that have been used successfully at IBM for a period of more than fifteen years. The effects of very small defects are discussed first. The case of photolithographic defects, which are of the same dimensions as the integrated circuit device and interconnection patterns, is dealt with in the remainder of the paper. The relationships between these models and test sites are described. Data from measurements of defect sizes are discussed.

SLAN-4: A Language for the Specification and Design of Large Software Systems *by F. Beichter, O. Herzog, and H. Petzsch, p. 558.* The language SLAN-4 has been defined in view of the need for formal tools supporting the specification and design of large software systems. It offers its users language constructs for algebraic and axiomatic specifications as well as for design in pseudocode. One of its major design goals has been to ease subsequent refinements of a (given) specification. The user can start his development with an informal high-level specification which can be formalized and implemented at a later date by using lower-level concepts. This paper provides the formal definitions of the SLAN-4 language, discusses the design decisions, and presents examples for the use of the syntactic constructs.

Replacing Square Roots by Pythagorean Sums *by C. Moler and D. Morrison, p. 577.* An algorithm is presented for computing a "Pythagorean sum" $a \oplus b = \sqrt{a^2 + b^2}$ directly from a and b without computing their squares or taking a square root. No destructive floating point overflows or underflows are possible. The algorithm can be extended to compute the Euclidean norm of a vector. The resulting subroutine is short, portable, robust, and accurate, but not as efficient as some other possibilities. The algorithm is particularly attractive for computers where space and reliability are more important than speed.

A Class of Numerical Methods for the Computation of Pythagorean Sums *by A. A. Dubrulle, p. 582.* Moler and Morrison have described an iterative algorithm for the computation of the Pythagorean sum $(a^2 + b^2)^{1/2}$ of two real numbers a and b. This algorithm is immune to unwarranted floating-point overflows, has a cubic rate of convergence, and is easily transportable. This paper, which shows that the algorithm is essentially Halley's method applied to the computation of square roots, provides a generalization to any order of convergence. Formulas of orders 2 through 9 are illustrated with numerical examples. The generalization keeps the number of floating-point divisions constant and should be particularly useful for computation in high-precision floating-point arithmetic.

Bending and Stretching an Elastic Strip Around a Narrow Cylindrical Drum *by H. C. Lee and H. D. Conway, p. 590.* An analysis is made of the deflected forms assumed by an elastic strip when it is bent and stretched around a rigid drum. A summary is first given of the method of analysis used and the results obtained in a previous study where the drum width is equal to or greater than that of the strip. This work is then extended to the case where the strip width exceeds that of the drum. Deflected forms and contact regions are delineated for a strip-width/drum-width ratio of two and for various values of two parameters: anticlastic deformation and tension.

Automated Twisted-Pair Wire Bonding *by D. E. Houser and K. J. Lubert, p. 598.* A need for highly dense and reworkable external interconnections with controlled characteristic impedance on high-performance printed-circuit boards led to the development of the process and equipment described in this paper. These interconnections are required to implement engineering changes, to install precise circuit delays, to repair printed-circuit defects, and to complete interconnections not possible with printed circuitry alone. As an economical way to meet the characteristic impedance requirement of $80 \pm 10\ \Omega$ and the density requirement for terminations on a 2.5-mm "staggered" grid (twice as dense as a 2.5-mm-square grid), 39-gauge twisted-pair wire is used.

A solder-reflow process bonds the wires to the printed circuitry and meets the need for reworkability. Computer-controlled equipment, developed to install twisted-pair wires at production rates, is now being used in several IBM plants worldwide.

Volume 28, Number 1, 1984

Optimizing Preventive Service of Software Products *by E. N. Adams, p. 2.* The implementer of a large, complex software system cannot make it completely defect free, so he must normally provide fixes for defects found after the code is put into service. A system user may do preventive service by installing these fixes before the defects cause him problems. Preventive service can benefit both the software developer and the software user to the extent that it reduces the number of operational problems caused by software errors, but it requires the expenditure of the resources required to prepare, disseminate, and install fixes; and it can be the cause of additional software problems caused by design errors introduced into the code by fixes. The benefit from removing a given defect depends on how many problems it would otherwise cause. Benefits may be estimated by modeling problem occurrence as a random process in execution time governed by a distribution of characteristic rates. It is found that most of the benefit to be realized by preventive service comes from removing a relatively small number of high-rate defects that are found early in the service life of the code. For the typical user corrective service would seem preferable to preventive service as a way of dealing with most defects found after code has had some hundreds of months of usage.

Combined Network Complexity Measures *by N. R. Hall and S. Preiser, p. 15.* Most of the considerable work that has been done in the measurement of software complexity during the past several years has addressed complexity measurement of source code or design languages. Here we describe techniques to measure the complexity of large (>100,000 source lines of code) systems during the software architecture phase, before major design decisions have been made. The techniques to measure and reduce complexity are intuitively reasonable, easy to apply, and produce consistent results. Methods developed include (1) an extension of the graph-theoretic measure developed by McCabe to software architecture, as represented by networks of communicating modules, (2) a general technique that allows the complexity associated with allocation of resources (CPU, tape, disk, etc.) to be measured, and (3) a method that combines module complexity and network complexity, so that design trade-offs can be studied to determine whether it is advantageous to have separate modules for service functions, such as mathematical subroutines, data management routines, etc.

Interactive Language Implementation System *by J. F. Sowa, p. 28.* The Interactive Language Implementation System (ILIS) is a tool for implementing language processors. It is fast enough for conventional compilers and general enough for processing natural languages. ILIS is built around a language for writing grammars. Unlike most compiler-compilers, the language includes a full range of semantic operators that reduce or eliminate the need for invoking other programming languages during a translation. ILIS is also highly interactive: It has facilities for tracing a parse and for adding or deleting grammar rules dynamically. This paper describes the features of ILIS and its use in several different projects.

Experience with Access Functions in an Experimental Compiler *by F. N. Ris, p. 40.* This paper describes an access function subsystem embedded in portions of an experimental microcode compiler which was built and used during 1973–6 using the IBM PL/I optimizing compiler under VM/370 and CMS. The use of the access function subsystem in this context was itself an experiment, performed by a group for all of whom PL/I was a new language and VM/370 a new operating system. The implementation of the subsystem was done strictly within the confines of the PL/I language. The basic objectives were ease of use, provision of a focal point for global storage management, extensive run-time validity checking with appropriate diagnostics, and data protection. Beyond satisfying these objectives, the subsystem proved more valuable than anticipated due to positive contributions made to debugging code in the VM/370 interactive development environment.

A New Programming Methodology for Long-Lived Software Systems *by R. Strom and N. Halim, p. 52.* A new software development methodology based on the language NIL is presented. The methodology emphasizes (1) the separation of program development into functional specification and tuning phases, (2) the use of a fully compilable and executable design, (3) an interface definition and verification mechanism. This approach reduces life-cycle costs and improves software quality because (a) errors are detected earlier, and (b) a single functional design can be re-used to produce many implementations.

A Program Development Tool *by C. N. Alberga, A. L. Brown, G. B. Leeman, Jr., M. Mikelsons, and M. N. Wegman, p. 60.* In this paper we describe how we have combined a number of tools (most of which are tailored to a particular programming language) into a single system to aid in the reading, writing, and running of programs. We discuss the efficacy and the structure of two such systems, one of which has been used to build several large application programs. We report some of the experience we have gained in evolving these systems. We first describe the system components which users have found most important; some of the facilities described here are new in the literature. Second, we attempt to show how these tools form a synergistic union, and we illustrate this point with a number of examples. Third, we illustrate the use of various system commands in the development of a simple program. Fourth, we discuss the implementation of the system components and indicate how some of them have been generalized.

Managing Multi-Version Programs with an Editor *by V. Kruskal, p. 74.* When more than one version of a program must be maintained, generally much of the code is repeated unchanged in many versions. Techniques such as "deltas" and

conditional compilation are commonly used to avoid duplicating these common parts. In addition to saving storage, these methods aid the programmer greatly in managing updates to the versions. Unfortunately, these representations of multi-version programs can appear very unlike a program, making them difficult to edit. Described here is a new method of automating much of the bookkeeping involved in dealing with multi-version programs. It entails use of a special editor that enables a multi-version program to be seen and modified in a fashion that is far closer to that normally permitted for a single-version program.

Mapping Uninterpreted Schemes into Entity-Relationship Diagrams: Two Applications to Conceptual Schema Design *by M. A. Casanova and J. E. Amaral de Sa, p. 82.* A method of mapping sets of uninterpreted record or relation schemes into entity-relationship diagrams is described and then applied to two conceptual design problems. First, the method is applied to the design of relational databases. It is shown that the method can be interpreted as a normalization procedure that maps a given relational schema into a new schema that represents an entity-relationship diagram. That is, the original schema has an interpretation in terms of higher-order concepts, which helps in understanding the semantics of the database it describes. The second design problem is related to the conversion of conventional file systems to the database approach. The method is used in this context to obtain a database conceptual schema from the description of the conventional system, which is one of the fundamental steps of the conversion process.

Performance Analysis of Future Shared Storage Systems *by A. Goyal and T. Agerwala, p. 95.* This paper deals with the analysis and design of two important classes of computer systems: BIP (Billion Instructions Per Second) systems consisting of a few very high performance processors and KMIP (K Million Instructions Per Second) systems with hundreds of low speed processors. Each system has large, shared semiconductor memories. Simple analytic models are developed for estimating the performance of such systems. The models are validated using simulation. They can be utilized to quickly reduce the design space and study various trade-offs. The models are applied to BIP and KMIP systems and their use is illustrated using examples.

***IBM Journal of Research and Development*: Information for Authors** *by H. Freitag, p. 119.* This paper combines background information about the *IBM Journal of Research and Development* with guidelines for the preparation of *Journal* manuscripts. The purpose is to acquaint authors with the *Journal* as a primary, professional publication and to present suggestions to ease the work of author and editor in preparing clear, concise, and useful manuscripts.

Volume 28, Number 2, 1984

Error-Correcting Codes for Semiconductor Memory Applications: A State-of-the-Art Review *by C. L. Chen and M. Y. Hsiao, p. 124.* This paper presents a state-of-the-art review of error-correcting codes for computer semiconductor memory applications. The construction of four classes of error-correcting codes appropriate for semiconductor memory designs is described, and for each class of codes the number of check bits required for commonly used data lengths is provided. The implementation aspects of error correction and error detection are also discussed, and certain algorithms useful in extending the error-correcting capability for the correction of soft errors such as α-particle-induced errors are examined in some detail.

An Introduction to Arithmetic Coding *by G. G. Langdon, Jr., p. 135.* Arithmetic coding is a data compression technique that encodes data (the data string) by creating a code string which represents a fractional value on the number line between 0 and 1. The coding algorithm is symbolwise recursive; i.e., it operates upon and encodes (decodes) one data symbol per iteration or recursion. On each recursion, the algorithm successively partitions an interval of the number line between 0 and 1, and retains one of the partitions as the new interval. Thus, the algorithm successively deals with smaller intervals, and the code string, viewed as a magnitude, lies in each of the nested intervals. The data string is recovered by using magnitude comparisons on the code string to recreate how the encoder must have successively partitioned and retained each nested subinterval. Arithmetic coding differs considerably from the more familiar compression coding techniques, such as prefix (Huffman) codes. Also, it should not be confused with error control coding, whose object is to detect and correct errors in computer operations. This paper presents the key notions of arithmetic compression coding by means of simple examples.

A Universal Reed-Solomon Decoder *by R. E. Blahut, p. 150.* Two architectures for universal Reed-Solomon decoders are given. These decoders, called time-domain decoders, work directly on the raw data word as received without the usual syndrome calculation or power-sum-symmetric functions. Up to the limitations of the working registers, the decoders can decode any Reed-Solomon codeword or BCH codeword in the presence of both errors and erasures. Provision is also made for decoding extended codes and shortened codes.

Implementation and Evaluation of a (b,k)-Adjacent Error-Correcting/Detecting Scheme for Supercomputer Systems *by J. Arlat and W. C. Carter, p. 159.* This paper describes a coding scheme developed for a specific supercomputer architecture and structure. The code considered is a shortened (b,k)-adjacent single-error-correcting

double-error probabilistic-detecting code with $b = 5$, $k = 1$, and code group width = 4. An evaluation of the probabilistic double-error-detection capability of the code was performed for different organizations of the coding/decoding strategies for the codewords. This led to the selection of a system organization encompassing the traditional feature of memory data error protection and also providing for the detection of major addressing errors that may result from faults affecting the interconnection network communication modules. The cost of implementation is a limited amount of extra hardware and a negligible degradation in the double-error-detection properties of the code.

Fault Alignment Exclusion for Memory Using Address Permutation *by D. C. Bossen, C. L. Chen, and M. Y. Hsiao, p. 170.* A significant improvement in memory fault tolerance, beyond what is already provided by the use of an appropriate error-correcting code (ECC), can be achieved by electronic chip swapping, without any compromise of data integrity as large numbers of faults are allowed to accumulate. Since most large and medium-sized semiconductor memories are organized so that each bit position of the system word (ECC codeword) is fed from a different chip, and quite often from a different array card, or at least from distinct partitions of an array card, the various bit positions have separate address circuitry on the array cards. This fact is important, and can be exploited to provide effective address permutation capability, which allows the realignment of faults which would otherwise have caused an uncorrectable multiple error in an ECC codeword. When faults occur in a codeword to produce an uncorrectable error (UE), the addressing within one of the error bit position array cards can be altered using simple EX-OR circuitry and storage latches. The content of the latches is computed using a fault map of the memory together with an algorithm. These techniques are referred to as Fault Alignment Exclusion (FAE) using address permutation. Practical considerations as to the complexity of the fault map, the number of storage latches per bit position, and the overall effectiveness of the permutation to disperse the expected numbers of errors are presented in this paper.

Fault-Tolerant Design Techniques for Semiconductor Memory Applications *by F. J. Aichelmann, Jr., p. 177.* Advances in semiconductor memory technology towards higher-density and higher-performance memory chips have created new reliability challenges for the memory system designer. An example would be the multiple-bit-per-chip organization with the chip outputs used in the same word. This design structure would be prone to uncorrectable errors with conventionally implemented single-error-correcting double-error-detecting codes. With these newer chips, memory system designers will have to give special attention not only to the types of failures but to ways of minimizing the system impact of reliability defects. In this paper, a number of design approaches are presented for minimizing the effects of chip failures through the use of organizational techniques and through enhancements to conventional error checking and correction facilities. The fault-tolerant design techniques described are compatible with most existing memory designs. An evaluative comparison of these techniques is included, and their application and utility are discussed.

Fault-Tolerant Memory Simulator *by C. L. Chen and R. A. Rutledge, p. 184.* Memory systems in modern computers employ a variety of methods to achieve fault tolerance, such as single- or double-error correction, page deallocation, or the use of spare chips or cells. Such methods ensure that the failure rate of the system is considerably less than the sum of the failure rates of the components. However, these methods also complicate the task of evaluating system reliability. The memory reliability function is too intractable to handle analytically, and we must turn to Monte Carlo methods. This paper describes the Fault-Tolerant Memory Simulator (FTMS), an interactive APL program which uses Monte Carlo simulation to evaluate the reliability of fault-tolerant memory systems.

A General-Purpose Memory Reliability Simulator *by M. R. Libson and H. E. Harvey, p. 196.* With rapid advances in computer memory capacity and performance, coupled with corresponding increases in the expense of field service calls, memory reliability and optimal maintenance strategies have become more and more important in terms of customer satisfaction and field service cost. At the same time, significant improvements in error correction and recovery over recent years have made the prediction of uncorrectable error and field service frequency much more difficult. This paper describes a Monte Carlo simulator which can predict uncorrectable error rates and field-replaceable-unit replacement rates for a wide range of memory architectures and under a variety of maintenance strategies. The model provides valuable information for performing sensitivity studies of intrinsic failure rates for memory components, for performing tradeoff studies of alternative storage module and card organizations, for evaluating system functions, and for establishing optimum maintenance strategies.

Analysis of Correctable Errors in the IBM 3380 Disk File *by T. D. Howell, p. 206.* A method of analyzing the correctable errors in disk files is presented. It allows one to infer the most probable error in the encoded-data stream given only the unencoded readback and error-correction information. This method is applied to the errors observed in seven months of operation of four IBM 3380 head-disk assemblies. It is shown that nearly all the observed errors can be explained as single-bit errors at the input to the channel decoder. About 90 percent of the errors were related to imperfections in the disk surfaces. The remaining 10 percent were mostly due to heads which were unusually susceptible to random noise-induced errors.

Iterative Exhaustive Pattern Generation for Logic Testing *by D. T. Tang and C. L. Chen, p. 212.* Exhaustive pattern logic testing schemes provide all possible input patterns with

respect to an output in the set of test patterns. This paper is concerned with the problem that arises when this is to be done simultaneously with respect to a number of outputs, using a single test set. More specifically, in this paper we describe an iterative procedure for generating a test set consisting of n-dimensional vectors which exhaustively covers all k-subspaces simultaneously, i.e., the projections of n-dimensional vectors in the test set onto any input subset of a specified size k contain all possible patterns of k-tuples. For any given k, we first find an appropriate N ($N>k$) and generate an efficient N-dimensional test set for exhaustive coverage of all k-subspaces. We next develop a constructive procedure to expand the corresponding test matrix (formed by taking test vectors as its rows) such that a test set of N^2-dimensional vectors exhaustively covering the same k-subspaces is obtained. This procedure may be repeated to cover arbitrarily large n ($n = N^{2^i}$ after i iterations), while keeping the same k. It is shown that the size of the test set obtained this way grows in size which becomes proportional to (log n) raised to the power of [log $(q+1)$], where q is a function of k only, and is approximated (bounded closely below) by $k^2/4$ in binary cases. This approach applies to nonbinary cases as well except that the value of q in an r-ary case is approximated by a number lying between $k^2/4$ and $k^2/2$.

Volume 28, Number 3, 1984

Preface *by E. E. Lucente, p. 233.*

Introduction to Topical Issue on Non-Impact Printing Technologies *by R. A. Myers and J. C. Tamulis, p. 234.* In this paper, we present an overview of the technologies of importance in computer output printing, providing a background for the succeeding papers in this issue. IBM workers have played a key role in the evolution of computer printers from a few marginally reliable complex mechanisms to an industry which today counts annual sales of over ten billion dollars worldwide. Thus, in this necessarily brief description of the different technologies we mention those in which the IBM role was particularly noteworthy, with specific mention of machines which brought the technologies to the marketplace.

Technology Trends in Electrophotography *by M. H. Lee, J. E. Ayala, B. D. Grant, W. Imaino, A. Jaffe, M. R. Latta, and S. L. Rice, p. 241.* Selected innovations and advances in electrophotographic printing are reviewed. Technological advances in photoreceptors, imaging techniques, toner development and cleaning, fusing, sensors, and new machine configurations and functions are discussed with respect to their possible applicability to future electrophotographic imaging systems. In general, improvements in print quality and reliability remain the main goals. In specific applications, added functions and an increased range of usable media are highly desired.

Introduction to the IBM 3800 Printing Subsystem Models 3 and 8 *by R. C. Miller, Jr., p. 252.* The IBM 3800 Printing Subsystem Models 3 and 8 are high-speed, non-impact, electrophotographic printers which are designed for system printing, text, and graphics applications requiring flexibility and high print quality. Features such as their print density of 240×240 pels/in.^{2} and all-point addressability extend substantially their printing capabilities beyond those of the Models 1 and 2.

Technology of the IBM 3800 Printing Subsystem Model 3 *by D. McMurtry, M. Tinghitella, and R. Svendsen, p. 257.* IBM has introduced the 3800 Model 3 electrophotographic printer, which is a modified version of the previously developed Model 1. The Model 3 improves the print density of the 3800 Model 1 from 180×144 pels/in.^{2} to 240×240 pels/in.^{2} and permits each pel to be addressed individually. The laser print head was modified to create the higher density by using the same laser, focusing to a smaller spot size, slowing down the speed of the rotating mirror, and developing a dual-beam system. Beam-power balance and beam separation were selected for optimum print quality. The dual-beam print head required the development of a new photoconductor with improved sensitivity. That photoconductor also displayed a significantly increased lifetime. Improved manufacturing techniques were developed to reduce photoconductor defects. A digital voltmeter and a new processor were introduced that reduced the cost of the process servos by eliminating previously hard-wired elements.

Electrophotographic Printer Control as Embodied in the IBM 3800 Printing Subsystem Models 3 and 8 *by C. Barrera and A. V. Strietzel, p. 263.* Control of the IBM 3800 Models 3 and 8 electrophotographic printers is achieved by use of a fundamentally different control system than was used in their predecessors, the Models 1 and 2. As a result, printing of composed pages or electronic overlays can include text of many different font sizes and styles printed in multiple orientations, as well as raster images up to a full page in size. The printers manage stored resources, including fonts, segments of pages, and electronic overlays. Pages are composed inside the printers in a logical sequence, instead of by the more traditional line-by-line sequence. This, as well as the capability to position text and images at any addressable point, enhances usability. A high-speed, table-driven character generator, a new command set, and a microcoded control unit make all of this possible.

Print Quality Measurements for High-Speed Electrophotographic Printers *by J. L. Crawford, C. D. Elzinga, and R. Yudico, p. 276.* Described here are some of the characteristics that make electrophotographic printing esthetically pleasing, and the use of a recently developed computer-controlled scanner for measuring those characteristics. New print quality measurements are described (for example, measurements of modulation, image gray-scale fidelity, and tangential-edge roughness) that allow the

monitoring of the advanced printing functions made possible by all-point addressability. The requirements for and implementation of the scanner are discussed. Also discussed are the effects and limitations on print quality measurement resolution imposed by the algorithms used and the scanner; the effect of light scatter by the paper; and the usefulness of the print quality measurements as an aid in making design trade-off decisions and in manufacturing control.

Paper Material Considerations for System Printers *by J. Borch and R. G. Svendsen, p. 285.* Laser printing of continuous forms by means of electrophotographic imaging requires that the paper on which the printing is carried out show good handling characteristics at fast machine speeds and be compatible with toner transfer and fixing. Toner transfer requires suitable electrical paper characteristics, similar to those of xerographic cut sheets intended for use in electrophotographic cut-sheet copiers and printers. Hot-pressure fusing in the IBM 3800 printing subsystem has been shown to require adequate paper-fiber wettability by the hot molten toner. Therefore, characteristics of paper chemistry, and in particular paper sizing, are essential in creating good image fixing (fusing). In addition, the temperature and pressure in the fusing assembly affect the thermal and dimensional stability of the paper. For the 3800, this necessitated modifications in roll design within the fusing station in order to ensure satisfactory paper handling.

Flash Fusing in Electrophotographic Machines *by G. W. Baumann, p. 292.* A theoretical model of the flash fusing process for electrophotographic machines was developed using the joint solution of a nonlinear circuit equation and the one-dimensional thermal diffusion equation. Numerous experiments were run using different toners according to the size of toner particles and the pulse width in order to determine the minimum energy that was required for fusing. The experiments confirm that this model predicts reasonably well what was observed in the lab. The melt depth required for good fusing is slightly less than mean particle size. At that depth the temperature is somewhat greater than the temperature required in the nip of a hot roll fuser for the same toner. Under typical flash fusing, the top surface of the toner is subjected to considerably higher temperatures than the melt temperature of the toner. From the combined analytical and experimental results, the proper compromises can be made for efficiency and volatiles.

Multiple-Nozzle Ink Jet Printing Experiment *by R. H. Darling, C.-H. Lee, and L. Kuhn, p. 300.* An experimental printer is described which utilizes a densely packed array of ink jets operated in a binary, pressurized, asynchronous mode. The printer configuration, the ink jet head structure, the operating point, and the maintenance approach are all discussed.

The Application of Drop-on-Demand Ink Jet Technology to Color Printing *by F. C. Lee, R. N. Mills, and F. E. Talke, p. 307.* The application of drop-on-demand ink jet technology to high-frequency and high-resolution color printing is investigated using an experimental multi-nozzle print head. Cross talk and drop misregistration are examined as a function of drop ejection frequency. Typical print samples from scanned and computer-generated data are obtained and the dependence of print quality on various parameters such as drop ejection frequency, ink composition, and paper quality is discussed.

Experimental and Theoretical Study of Wave Propagation Phenomena in Drop-on-Demand Ink Jet Devices *by D. B. Bogy and F. E. Talke, p. 314.* This paper presents experimental observations and a theoretical analysis of the operation of drop-on-demand piezoelectric ink jet devices. By studying experimentally the dependence of several operating characteristics on the length of the cavity in the nozzle of an ink jet device, we have gained insight into the physical phenomena underlying the operation of such a device. It is concluded that drop-on-demand ink jet phenomena are related to the propagation and reflection of acoustic waves within the ink jet cavity. A simple analysis is carried out on the basis of linear acoustics which is in good agreement with the experimental observations.

Numerical Calculation of the Fluid Dynamics of Drop-on-Demand Jets *by J. E. Fromm, p. 322.* A numerical method that makes use of the complete incompressible flow equations with a free surface is discussed and used to study an impulsively driven laminar jet. Flow behavior dependence upon fluid properties (characterized by a Reynolds number over Weber number nondimensionalization) is compared for drop integrity purposes. Several variations of square wave pressure history applied at a nozzle inlet are discussed in relation to drop velocities produced and structure of ejected drops. Timewise development of flow both interior and exterior to the nozzle is illustrated through computed contour sequences.

Volume 28, Number 4, 1984

Statistical failure analysis of system timing *by D. R. Tryon, F. M. Armstrong, and M. R. Reiter, p. 340.* Techniques are developed that quantify the probability that large computer systems will meet their cycle time objectives. Both approximation techniques and rigorous multivariate statistical techniques are described. A method is developed that enumerates the cycle-limiting paths so that these approaches can be utilized. The results of these techniques enable system designers to ensure that performance and reliability objectives are met.

Measures of ideal execution architectures *by M. J. Flynn and L. W. Hoevel, p. 356.* This paper is a study in ideal computer architectures or program representations. We define measures of "ideal" architectures that are related to the higher-level representation used to describe a program at the

source language level. Traditional machine architectures name operations and objects which are presumed to be present in the host machine: a memory space of certain size, ALU operations, etc. An ideal language-based architecture is based on a specific higher-level (source) language, and uses the operations in that language to describe transformations over objects in that language. The notion of ideal is necessarily constrained. The object program representation must be easily decompilable (i.e., the source is readily reconstructable). It is assumed that the source itself is a good representation for the original problem; thus any nonassignment operation present in the source program statement appears as a single instruction (operation) in the ideal representation. All named objects are defined with respect to the scope of definition of the source program. For simplicity of discussion, statistical behavior of the program and language is assumed to be unknown; Huffman codes are not used. From the above, canonic interpretive (CI) measures are developed. CI measures apply to both the space needed to represent a program and the time needed to interpret it. Example-based CI measures are evaluated for a variety of contemporary architectures, both host-and language-oriented, as well as a CI-derived language-oriented architecture.

Design considerations of a static LSSD polarity hold latch pair *by A. Correale, p. 370.* There are many considerations relating to the design of a static LSSD polarity hold latch pair. High performance, low power dissipation, small size, and stability are some of the major requirements for a good design. The engineering trade-offs needed to ensure that all goals are met are discussed.

All points addressable raster display memory *by R. Matick, D. T. Ling, S. Gupta, and F. Dill, p. 379.* This paper discusses display designs which store the image point by point in random access memory, so that independent update of every pixel is possible. A frequent bottleneck in the design of high performance displays of this type is the available bandwidth of the memory subsystem. In this paper, we focus on this issue and present features of a customized dynamic RAM chip which can readily provide the necessary bandwidth and thus greatly simplify the design of very high performance APA raster scan displays. The customized RAM chip is quasi-two-ported. After briefly introducing APA raster displays, we discuss display memory system design and the design of the proposed custom memory chip. We describe the second port for the video refresh, which makes the primary port available for update almost continuously. We also discuss modifications to the existing primary port to make it easily usable for the parallel update required for high update performance as well as for other applications.

A mapping and memory chip hardware which provides symmetric reading/writing of horizontal and vertical lines *by D. L. Ostapko, p. 393.* This paper describes a mapping and memory chip hardware for enhancing the performance of an APA display. The approach describes a modification to the primary port of a quasi-two-ported memory. This modification allows several contiguous horizontal or vertical bits to be read or written in one cycle. The number of bits that can be stored is given by the number of memory chips. The hardware modifications can be on or off chip, and if on chip, the chip can still be used as a conventional memory chip. Simple modifications to the hardware will support different screen sizes.

Reduction of random noise from multiband image data using phase relationships among their Fourier coefficients *by J. V. Dave and J. Gazdag, p. 399.* The problem of reducing the random noise from multiband image data is examined, taking a nine-band image data set of a terrestrial scene and adding a significant amount of noise to the original pixel value of each band. For each of these data sets, data of three nearby bands generally exhibit a high correlation among them. Selected data for several sets of three nearby bands are transformed from the spatial domain to the frequency domain to study phase relationships (i.e., coherency) among their Fourier coefficients of various frequencies. It is shown that, in general, the addition of random noise results in the rapid change, with frequency, of a quantity called the coherency measure. (This is a quantitative measure of the phase agreement among the phases of various Fourier coefficients at a given frequency.) The coherency measure vs. frequency curve for a given data line is then used to attenuate various Fourier coefficients of the corresponding nearby bands of that line. It is then shown that the inverse transformation of such modified Fourier coefficients results in a statistically significant reduction of noise from the data of single lines, or from those data of some finite areas of the image. Results of a supervised boxcar multispectral classification with the original as well as various modified data sets of the selected image are also presented to provide additional guidance in the use of such a sophisticated analytic procedure in image processing.

Empty arrays in extended APL *by D. L. Orth, p. 412.* During the past several years considerable work has been done on extending APL in three areas: operators, heterogeneous data, and nested data. In each area a proposed extension must treat empty arrays consistently. In this paper various possibilities for providing consistent behavior are presented. The new proposals possess at least one of two important qualities in which older proposals tend to be deficient: consistent behavior is independent of the structural properties of rank and nesting, and the user has control over the behavior when he wants it.

Software reliability analysis models *by M. Ohba, p. 428.* This paper discusses improvements to conventional software reliability analysis models by making the assumptions on which they are based more realistic. In an actual project environment, sometimes no more information is available than reliability data obtained from a test report. The models described here are designed to resolve the problems caused

by this constraint on the availability of reliability data. By utilizing the technical knowledge about a program, a test, and test data, we can select an appropriate software reliability analysis model for accurate quality assessment. The delayed S-shaped growth model, the inflection S-shaped model, and the hyperexponential model are proposed.

The digital data exchange—A space-division switching system *by E. Hopner and M. A. Patten, p. 444.* The user requirements for an automated switch for computer terminals in a large establishment or laboratory environment are discussed. This is followed by an explanation of the design requirements, and finally by a description of the significant features of the digital data exchange (DDEX), a microprocessor-controlled user-transparent space-division digital data switch of modular design which is capable of connecting 512 to 2048 lines from various types of IBM terminals to different control units, thereby substantially saving interconnection resources.

Resist profile control in E-beam lithography *by S. J. Gillespie, p. 454.* Imaging studies have confirmed that a desired resist profile can be obtained by selecting the appropriate combination of process parameters: dose, interrupted development, pattern bias, and resist thickness. Bias sensitivity of the resist image to process parameters was measured using a positive diazo resist with nonlinear development characteristics on an IBM EL-3 E-beam tool. Because of superior bias stability, top-edge imaging with undercut profiles in a single-layer resist was found to provide many of the imaging advantages of a multilayer system. Sufficient resolution and image quality are obtained to extend the application of a single-layer resist system to 1-μm lithography.

Modeling of defects in integrated circuit photolithographic patterns *by C. H. Stapper, p. 461.* In a previous paper by the same author the foundation was laid for the theory of photolithographic defects in integrated circuits. This paper expands on the earlier one and shows how to calculate the critical areas and probability of failure for dense arrays of wiring. The results are used to determine the nature of the defect size distribution with electronic defect monitors. Several statistical techniques for doing this are described and examples are given.

Aspects of the traveling salesman problem *by M. Held, A. J. Hoffman, E. L. Johnson, and P. Wolfe, p. 476.* For fifty years the traveling salesman problem has fascinated mathematicians, computer scientists, and laymen. It is easily stated, but hard to solve; it has become the prototypical hard problem in theoretical computer science. A large part of the extensive research conducted by IBM in the broad area of optimization, or mathematical programming, contributed to or was inspired by aspects of this challenging problem. This article reviews some of that work as well as recent developments in techniques that were used on the largest traveling salesman problem ever solved.

Volume 28, Number 5, 1984

A software architecture for a mature design automation system *by R. L. Taylor, p. 501.* Design automation systems are groups of programs used to aid the design of the electronic portions of computers. IBM has used such systems for over twenty-five years, and has a large number of experienced users who place severe requirements upon their design automation system. The essential requirement is that design programs must be easy to use in the way that a particular development location wishes to use them. Design programs which cannot be changed to meet local requirements are not acceptable. The software structure portion of an architecture for a design automation system which meets these requirements is described; interactive and foreground design applications are stressed. The structure involves a uniform set of services, such as command languages and terminal access methods, needed to support the design application, and a formal, two-part partitioning of the design application itself. Modularity and good interfaces are important parts of the software structure which permit a development laboratory to change the application enough for it to be easily used. These changes are the rule, not the exception. Examples of such changes are given, based upon early user experience. The applicability of this architecture to future system needs such as distributed data and distributed processing is discussed.

A device-independent graphics package for CAD applications *by R. B. Capelli and G. C. Sax, p. 512.* GSSP (Graphics Support Subroutine Package) is a device-independent two-dimensional graphics package developed by Engineering Design Systems (EDS) to support several major electronic and mechanical computer-aided design applications within IBM. Graphics systems supported range from interactive, high-function, distributed-graphics workstations to passive graphic-output devices. GSSP provides many of the functions usually found in other device-independent graphics packages, with additional support for hierarchically structured display files and distributed graphics systems. The major functions provided by GSSP are described, and an overview of the implementation is presented to show how issues such as interactive performance and human factors are addressed.

An interactive system for VLSI chip physical design *by W. H. Elder, P. P. Zenewicz, and R. R. Alvarodiaz, p. 524.* The Federal Systems Division has developed a structured design methodology and a companion chip physical design system that has been used to build seven large VLSI chips (ranging in size from 7K to 36K logic primitives). Using the MVISA system, a logic designer has complete control and responsibility for the total chip design. Our experience has been that when this highly interactive software and methodology is used, chip physical design requires less than

two weeks. This is a significant savings in design time; but more importantly the designer can allocate more schedule for logic design and simulation. This paper describes how FSD's unique interactive physical design system has improved productivity of VLSI design.

LSS: A system for production logic synthesis *by J. A. Darringer, D. Brand, J. V. Gerbi, W. H. Joyner, Jr., and L. Trevillyan, p. 537.* For some time we have been exploring methods of transforming functional specifications into hardware implementations that are suitable for production. The complexity of this task and the potential value have continued to grow with the increasing complexity of processor design and the mounting pressure to shorten machine design times. This paper describes the evolution of the Logic Synthesis System from an experimental tool to a production system for the synthesis of masterslice chip implementations. The system was used by one project in IBM Poughkeepsie to produce 90 percent of its more than one hundred chip parts. The primary reasons for this success are the use of local transformations to simplify logic representations at several levels of abstraction, and a highly cooperative effort between logic designers and synthesis system designers to understand the logic design process practiced in Poughkeepsie and to incorporate this knowledge into the synthesis system.

Automated technology mapping *by J. L. Gilkinson, S. D. Lewis, B. B. Winter, and A. Hekmatpour, p. 546.* Logic "mapping," or "transformation," refers to the process of converting a logic design from one form of specification to another. The output is usually a specific technology implementation and the input could range from a previous technology implementation to a high-level design language. Motivated initially by the problem of test case generation for new technologies, a logic transformation system, known as the Technology Mapping System (TMS), was developed. This system has focused on the problem of technology-to-technology mapping involving gate array or standard cell logic families. TMS makes use of an intermediate notation, called GLN, and uses several forms of "rules" to control the mapping process. This paper discusses the history and general operation of TMS, and makes a comparison of transformations from different types of sources.

Hardware design and description languages in IBM *by L. I. Maissel and H. Ofek, p. 557.* Hardware design languages (HDLs) allow computer hardware to be described in sufficient detail to be simulated and built, such a description being at a sufficiently high level of abstraction to make the complete design readily intelligible to anyone skilled in that language. A number of HDLs have been developed and are in use in IBM. To date, no overwhelming case can be made for choosing any one HDL over the others. The major trends in HDL are discussed. Several examples of HDLs are presented in some detail. VHDL, the yet-to-be released HDL which is to serve as a front end to the U.S. Government's Very High Speed Integrated Circuits program, is among these.

Using a hardware simulation engine for custom MOS structured designs *by Z. Barzilai, D. K. Beece, L. M. Huisman, and G. M. Silberman, p. 564.* Mixed-level simulation techniques are widely used in VLSI designs for verification and test evaluation. In this paper we indicate how to perform mixed-level simulation on structured MOS designs using the Yorktown Simulation Engine (YSE), a hardware simulator developed at IBM. On the YSE, simulation can be done at the functional, gate, and transistor levels. The design specification used by the YSE is well suited for mixed-level simulation, particularly with regard to interfacing the different levels. We apply our techniques to an nMOS design to show the important features of our approach.

PSI: A symbolic layout system *by R.-D. Fiebrich, Y.-Z. Liao, G. Koppelman, and E. Adams, p. 572.* A symbolic layout tool, PSI, is described for use with IBM circuit technology. Significant features of PSI are used with multiple circuit technologies, adaptation to rapid changes of technology design rules, creation of nested designs, and extensive designer control over the spacing process.

Constraint solver for generalized IC layout *by P. W. Cook, p. 581.* This paper presents a constraint solver suitable for use in a general symbolic IC layout system. The essential features of the constraint solver, which is intended to place few restrictions on the source of the constraints to be solved, are that it accommodate mixed equality and inequality constraints, that it allow selective "maximization" of variables, that it proceed with any number of variables given user-defined values, and that it fail to produce a solution only when no solution exists. These features all flow from the desire to provide a constraint solver suitable for use in an "open" system, in which there are no restrictions on the form or order of the constraints. The algorithm presented meets these objectives while remaining reasonable in its use of storage and time. An extension to the class of constraints acceptable by the constraint solver is presented; the extension of the system to this added constraint class has yet to be done.

Custom Chip/Card Design System *by A. M. Barone and J. K. Morrell, p. 590.* The Custom Chip/Card Design System (CCDS) is a set of software applications, tied together via a common data interchange, that is used for the design, analysis, and checking of custom electronic circuits. CCDS is intended to unite the separate electrical, logical, and physical design phases into a single design process. The underlying principle of the system rests on the idea of describing the product as a generalized network, with multiple overlapping views that correspond to the different design phases. In addition to the applications contained within CCDS, there are also links to other software tools for such things as circuit simulation.

ACORN: A system for CVS macro design by tree placement and tree customization *by P. S. Hauge and E. J. Yoffa, p. 596.* ACORN is a system for the physical design of

cascode voltage switch (CVS) macros which utilizes tree placement and tree customization to improve macro wirability. The results obtained by designing a 43-tree differential (DCVS) macro on a masterslice chip image are presented to illustrate the design improvements. In this example, tree placement reduces wire length and via count by 12 percent relative to a transistor-pair placement design. Tree customization increases this improvement to 25 percent, and increases the porosity of the wired macro to vertical global wires from 8 percent for tree placement alone to 15 percent.

KWIRE: A multiple-technology, user-reconfigurable wiring tool for VLSI *by P. C. Elmendorf, p. 603.* In a VLSI design environment where a range of chip technologies are available and concurrent chip designs are commonplace, it is not feasible to build a wiring program for each technology. Additionally, a chip's design methodology may demand specific abilities from a wiring program. KWIRE was developed to meet the needs of a multiple-technology, multiple-methodology VLSI design community. It has been used on a range of chips, from small designs to custom microprocessors. Modeling the users' designs and design rules in geometric terms allows KWIRE to handle such a diversity of chip designs. This paper describes the KWIRE system and router.

An iterative-improvement penalty-function-driven wire routing system *by R. Linsker, p. 613.* A wire routing system (VIKING) has been developed for interconnection packages. It uses iterative-improvement methods that allow "illegalities" (such as wire crossings within a plane) at intermediate stages of the routing, eliminates some drawbacks of conventional sequential routers, and extends the range of penalty functions with respect to which a wiring configuration can be optimized. Efficient routing in directionally uncommitted planes is provided; specification of preferred-direction (x and y) planes is optional but not required. Significant reductions in required manual embedding effort, number of vias required, routed wire length, and the number of signal planes required to wire a package, have been found, compared with sequential routers that have been used. Improved automatic control of electrical crosstalk noise has also been provided. In addition to presenting VIKING methods and results, we discuss other issues relating to wiring methods and global optimization. Application of these methods to chip design is also discussed.

A CMOS LSSD test generation system *by D. Leet, P. Shearon, and R. France, p. 625.* Automatic test pattern generators based on the stuck-fault concept are theoretically inadequate in their ability to generate test patterns for CMOS circuits. A new set of pin faults, called CMOS faults, is discussed that can represent the necessary test pattern sequences for these circuits. Processing of these faults by a new test pattern generator, called the Enhanced Test Generator (ETG), is also described.

Yield model for fault clusters within integrated circuits *by C. H. Stapper, p. 636.* Generalized negative binomial statistics turns out to be a model of the fault distribution in very large chips or wafers with internal defect clusters. This is expected to influence large chip and full wafer redundancy requirements. Furthermore, the yield appears to be affected by an experimental dependence of the average number of faults on chip area.

Volume 28, Number 6, 1984

Preface: Advances in materials and processes for printed circuit packaging technology *by D. P. Seraphim, P. A. Toole, W. T. Chen, and R. Rosenberg, p. 652.*

Moisture solubility and diffusion in epoxy and epoxy-glass composites *by L. L. Marsh, R. Lasky, D. P. Seraphim, and G. S. Springer, p. 655.* The solubility and kinetics of moisture transport mechanisms in epoxy-type resin and resin-glass composites have been investigated over a range of partial pressure and temperature. Moisture absorption-desorption in these systems is a quasi-reversible process, the kinetics of which are non-Fickian (Type II) and dependent on prior history. The multistaged sorption and transport behavior are interpreted in terms of multiphase models.

Bending-cantilever method for the study of moisture swelling in polymers *by B. S. Berry and W. C. Pritchet, p. 662.* Self-induced bending of a bilayer strip is shown to be a simple but sensitive method for the study of water absorption and swelling in polymers. Expressions for both the time-dependent and equilibrium curvature of the strip have been derived, enabling both diffusional and dilatational parameters to be extracted from experimental data. To illustrate the technique, it is shown that a vacuum-dried epoxy swells linearly with water content, at a rate of 0.93% in volume per weight percent of water. Desorption into vacuum has been observed under diffusion-controlled conditions and the diffusion coefficient for water in the epoxy found to be 2.5×10^{-9} cm²/s at 295K.

Mechanisms of electroless metal plating: I. Mixed potential theory and the interdependence of partial reactions *by P. Bindra, D. Light, and D. Rath, p. 668.* Electroless plating reactions are classified according to four overall reaction schemes in which each partial reaction is either under diffusion control or electrochemical control. The theory of a technique, based on the observation of the mixed potential as a function of agitation, concentration of the reducing agent, and concentration of metal ions, is presented. By using this technique it is shown that in electroless copper plating the copper deposition reaction is diffusion-controlled, while the formaldehyde decomposition reaction is activation-controlled. Values of the kinetic and mechanistic parameters for the partial reactions obtained by this method and by other electrochemical methods indicate that the two partial reactions are not independent of each other.

Mechanisms of electroless metal plating: II. Decomposition of formaldehyde *by P. Bindra, J. M. Roldan, and G. V. Arbach, p. 679.* A detailed investigation of the decomposition of formaldehyde was carried out to account for the fact that formaldehyde decomposition on Group VIII metals, e.g., Pd, occurs without simultaneous hydrogen generation, while on Group IB metals, e.g., Cu, formaldehyde decomposition is accompanied by hydrogen evolution. It was found that in principle metals may be divided into three main classes: (a) metals with positive free energy of hydrogen adsorption, (b) metals with free energy of hydrogen adsorption close to zero, and (c) metals with negative free energy of hydrogen adsorption. In the case of class (a) metals formaldehyde oxidation is accompanied by hydrogen evolution; for class (b) metals there is no simultaneous hydrogen evolution; and class (c) metals show low catalytic activity for formaldehyde oxidation. Hence, formaldehyde cannot be used as a reducing agent for electroless plating of class (c) metals.

Initiation of electroless Cu plating on nonmetallic surfaces *by J. Horkans, C. Sambucetti, and V. Markovich, p. 690.* Electroless plating of a metal on a dielectric substrate requires the prior deposition of a catalyst such as a Pd-Sn colloid consisting of a metallic Pd core surrounded by a stabilizing layer of Sn ions. The activation step (deposition of the colloid) is usually followed by an acceleration step (removal of excess ionic tin). Adhesion of the deposit to the substrate is improved by mechanical and chemical pretreatment steps. An electrochemical method has been developed for assessing the catalytic activity of Pd-Sn colloids. Hydrogen sorption in the Pd in the colloid can be correlated with catalytic activity, since Pd accessible for the H-sorption reaction is also accessible as the catalyst for the electroless deposition reaction. These conclusions have been confirmed by surface analytical techniques and by functional tests. The efficacy of various accelerating solutions has also been assessed.

Microstructure evolution during electroless copper deposition *by J. Kim, S. H. Wen, D. Y. Jung, and R. W. Johnson, p. 697.* A study using transmission and scanning electron microscopy was made of the evolution of the microstructure of electroless plated Cu on activated amorphous substrates and on single-crystal Cu grains. On amorphous substrates activated in a $PdCl_2$–$SnCl_2$ colloidal solution, Sn atoms dissolved into the plating solution concurrently with Cu deposition on the substrate during the initial stage of deposition. The very small face-centered-cubic grains of Cu-Pd solid solution agglomerated into much larger particles and later coalesced into spherical grains. As the grains grew, they developed crystallographic facets, impinged upon one another, and finally covered the entire substrate. Grains of energetically favorable crystallographic orientation selectively developed into the columnar structure. These columnar grains contained subgrains, dislocations, and twins. Remarkably different structures were observed for the Cu grown on large single-crystal grains. In this case epitaxial growth dominated the plating process. Low-surface-energy (111) Cu planes were frequently observed on plated Cu surfaces. Growth rates were a function of substrate orientation.

Micromechanics of multilayer printed circuit boards *by L. C. Lee, V. S. Darekar, and C. K. Lim, p. 711.* Analytical and experimental techniques are reported for the evaluation of micromechanical components in multilayer printed circuit boards. The concern in this investigation comes from the *Z*-axis thermal mismatch between epoxy-glass and copper that generates stresses when the board is subject to a temperature change. Finite-element modeling for both plated through-holes (PTH) and buried via (PV) structures is used to calculate the stresses in the copper barrel and at the via junctions. A simple experiment is designed to measure the thermomechanical strain in the PTH barrel. Also discussed are the PTH peel and PV pull techniques which have been used to characterize the barrel-laminate adhesion and the via junction strength.

Optimization of interconnections between packaging levels *by J. H. Kelly, C. K. Lim, and W. T. Chen, p. 719.* In large-scale integrated circuits, the interface between ceramic modules and the next level—epoxy-glass circuit boards or cards—contains a large number of pin arrays. Because the modules and the card are usually quite rigid and mechanically strong, the interface between the module and the card is commonly the weakest region in the assembly system. This interface is where the differential deformations between the two levels of packaging are accommodated. This paper describes a theoretical and experimental program to understand the loadings and stresses present, and to optimize the design of the connecting pin in order to distribute the stresses more evenly across the surfaces of the braze joint that connects the pin to the ceramic module. This work was done jointly by members of the East Fishkill and Endicott laboratories.

Immersion tin: Its chemistry, metallurgy, and application in electronic packaging technology *by Z. Kovac and K.-N. Tu, p. 726.* The surfaces in copper-plated through-holes in printed circuit boards for complex electronic packaging can be made solderable by immersion deposition of tin. The properties of the prepared surfaces vary from those of "white immersion tin," which is easily wetted by molten tin solder, to those of "gray immersion tin," which is nearly nonwettable. The wettability affects the electrical contact between the printed circuit board and the modules on it. In this paper, the rate law for tin deposition in the tin immersion plating bath is studied. Certain effects of chemical composition of the plating bath upon the character of the tin layer are investigated and the effects of thermal annealing of the plated surface upon the composition of the tin layer are determined. The differences in composition of white and gray immersion tin surface coatings are revealed by X-ray diffraction Rutherford backscattering spectroscopy and Auger electron spectroscopy.

Solderability tests on Sn, Cu, Cu_3Sn, and immersion tin surfaces are included.

High-temperature stability of a polyimide film *by R. Ginsburg and J. R. Susko, p. 735.* Polyimide (PMDA-ODA) films were analyzed by mass spectroscopy to determine their high-temperature stability. Using a high-resolution instrument, the identity of the low-molecular-weight evolved gases was confirmed. With a semiquantitative technique, the effect of a vacuum pre-bake was shown to reduce outgassing appreciably during subsequent treatment at high temperature. Subjection of the films to moisture did not affect their thermal stability. Low-temperature processing (240°C vs 400°C) reduced gaseous evolution by an order of magnitude.

Determination of Gafac in complex solution matrices *by S. A. Schubert, p. 741.* A new analytic method is described for determining the concentration of Gafac Re-610 (trademark of the GAF Corporation, New York, NY) in multicomponent solutions. This method utilizes a simple methylene chloride extraction to separate the Gafac from interfering chemical species, such as cupric sulfate. The ultraviolet absorbance of the methylene chloride extract is then measured at 276 nm and is shown to be proportional to the concentration of Gafac over the range of 1–170 ppm. However, this relationship is nonlinear except for concentrations less than 15 ppm. The limit of detection is 0.6 ppm and the relative precision at the 10-ppm level is ± 6%. Experiments to optimize and characterize various aspects of the analytical procedure are described, including determining the absorptivity of Gafac, measuring the distribution ratio, calculating extraction efficiencies, optimizing the extraction pH, and evaluating selected spectral interferences.

Volume 29, Number 1, 1985

The dynamical equations governing a lubricating film consisting of a gas film overlying a liquid film *by W. E. Langlois, p. 2.* The dynamical equations governing a two-phase lubricating configuration are derived. It is assumed that a gas film overlies a liquid film, both thin enough that the lubrication approximation may be used. The analysis leads to coupled Reynolds equations governing the pressure and the relative thickness of the two films. The coupling, which is determined by continuity of tangential stress across the gas-liquid interface, is considerably simplified if the shear viscosity of the liquid greatly exceeds that of the gas.

The stability of a colloidal suspension of coated magnetic particles in an aqueous solution *by D. Y. C. Chan, D. Henderson, J. Barojas, and A. M. Homola, p. 11.* Expressions for the magnetic, electrostatic and van der Waals interactions between isolated magnetic spheres which are coated with an inert material and immersed in an aqueous electrolyte solution are obtained and used to study the stability of a colloid of spheres in an electrolyte. Use is made of a simplified version of the theory of colloid stability of Derjaguin, Landau, Verwey, and Overbeek. We find that the colloidal dispersion becomes more stable as 1) the electrolyte concentration is decreased, 2) the radius of the magnetic spheres is decreased, or 3) the thickness of the inert layer is increased. In order to obtain stability with uncoated spheres, the spheres should have radii of about 5 nm. Such radii are typical of ferrofluids.

A dielectric loss investigation of moisture in epoxy-glass composites *by L. L. Marsh, D. C. Van Hart, and S. M. Kotkiewicz, p. 18.* Dielectric loss has been used to study moisture absorption-desorption equilibria and kinetics in epoxy-glass composites. Measurements were made at a temperature of 90°C and a frequency of 200 Hz to maximize sensitivity to interfacial or Maxwell-Wagner polarization. Exposure to various partial pressures of moisture at that temperature permitted a kinetics analysis from which an estimate was made of the diffusivity of water in the composites. Values of dielectric loss at saturation were used to establish a moisture/dielectric loss calibration at 90°C which can be used to estimate the macroscopic internal moisture content of coupon samples and printed circuit cards and boards. Dielectric loss measurements offer promise as a screening technique for resin-glass coupler development.

Cathodic delamination of methyl methacrylate-based dry film polymers on copper *by J. M. Atkinson, R. D. Granata, H. Leidheiser, Jr., and D. G. McBride, p. 27.* Studies of the bond degradation between a laminated organic coating and a copper substrate have been carried out using electrochemical techniques. The failure of the bond is attributed to a cathodic reaction which occurs under the coating. The rates of delamination are shown to be affected by delay time after exposure, temperature, applied potential, composition of the electrolyte, and surface abrasion prior to application of the coating.

A development of *APL2* syntax *by J. A. Brown, p. 37.* This paper develops the rules governing the writing of *APL2* expressions and discusses the principles that motivated design decisions.

A hardware sort-merge system *by N. Takagi and C.-K. Wong, p. 49.* The hardware sort-merge system which can sort large files rapidly is proposed. It consists of an initial sorter and a pipelined merger. In the initial sorter, record sorting is divided into two parts: key-pointer sorting and record rearranging. The pipelined merger is composed of several intelligent disks each of which has a simple processor and some buffers. The hardware sort-merge system can sort files of any size by using the pipelined merger repeatedly. The key-pointer sorting circuit in the initial sorter requires only unidirectional connections between neighboring cells, instead of the usual bidirectional ones. The initial sorter can also generate sorted sequences longer than its capacity so that the number of merging passes can be reduced. A new data management scheme is proposed to run all merging passes in a pipelined fashion.

Adaptive clustering algorithm *by L. V. O'Malley, p. 68.* Output data from many types of sensor systems (radar, radar warning, sonar, electro-optical, etc.) must be associated with one or more possible sources based on multiple observations of the data. This paper presents an algorithm that associates data with their source by simultaneous n-dimensional clustering of multiple data observations. The algorithm first orders the observations by successive nearest neighbor, in the n-dimensional Euclidean sense, from a defined starting point. Clusters are then isolated using a method derived from statistical decision theory. The algorithm's primary feature is its ability to perform clustering adaptively without any assumptions about the size, number, or statistical characteristics of the clusters. Since the algorithm was developed for radar warning system processing, a performance comparison with a well-known algorithm used in that field is included.

Novel method for analysis of printed circuit images *by J. R. Mandeville, p. 73.* To keep pace with the trend towards increased circuit integration, printed circuit patterns are becoming denser and more complex. A variety of automated visual inspection methods to detect circuit defects during manufacturing have been proposed. This paper describes a method which is a synthesis of the *reference-comparison* and the *generic-property* approaches that exploits their respective strengths and overcomes their respective weaknesses. It is based on the observation that the local geometric and global topological correctness of a printed circuit can be inferred from the correctness of simplified, skeletal versions of the circuit in a test image. These operations can be realized using

simple *processing elements* which are well suited for implementation in hardware.

The effects of wafer to wafer defect density variations on integrated circuit defect and fault distributions *by C. H. Stapper, p. 87.* A method for modeling the variations in defect levels in circuits produced on modern integrated circuit manufacturing lines is described in this paper. The effects on defect and fault distributions are derived. A deficiency in some previous yield models is eliminated.

Volume 29, Number 2, 1985

Microprocessors in brief *by R. C. Stanley, p. 110.* This paper presents a tutorial overview of the past, present, and future of microprocessors and describes the key elements of their structure and operation. It is intended to serve as a technical introduction to the rapidly expanding field of microprocessor and microcomputer technology and to provide an overview of what these elements are, what they can do, and how they do it. The origin and evolution as well as the basic principles of operation are discussed. Several different types of microprocessor are considered and examples of their application in the solution of real-world problems are given.

Architecture of a digital signal processor *by G. Ungerboeck, D. Maiwald, H.-P. Kaeser, P. R. Chevillat, and J. P. Beraud, p. 132.* A digital signal processor (DSP) is described which achieves high processing efficiency by executing concurrently four functions in every processor cycle: instruction prefetching from a dedicated instruction memory and generation of an effective operand, access to a single-port data memory and transfer of a data word over a common data bus, arithmetic/logic-unit (ALU) operation, and multiplication. Instructions have a single format and contain an operand, index control bits, and two independent operation codes called "transfer" code and "compute" code. The first code specifies the transfer of a data word over the common data bus, e.g., from data memory to a local register. The second determines an operation of the ALU on the contents of local registers. A fast free-running multiplier operates in parallel with the ALU and delivers a product in every cycle with a pipeline delay of two cycles. The architecture allows transversal-filter operations to be performed with one multiplication and ALU operation in every cycle. This is accomplished by a novel interleaving technique called ZIP-ing. The efficiency of the processor is demonstrated by programming examples.

Signal processor chip implementation *by J. P. Beraud, p. 140.* Very large microprocessors can now be integrated on a single chip; the integration eliminates packaging delays and is especially attractive for performance-oriented processors such as signal processors. This paper describes a semicustom signal processor chip designed jointly by the IBM France Essonnes and La Gaude laboratories. The logic is implemented from an optimized library of bipolar circuits. Layouts are compatible and designed to map data flow structures efficiently. Chip design time has been greatly reduced through the use of developed CAD tools tailored to our methodology. The design achieves twice the density that would be possible (with the same technology) with a masterslice. The chip's high performance has been verified with hardware; it provides enough computation power for 125 second-order filters with 8-kHz sampling of the input signal.

Voice-excited predictive coder (VEPC) implementation on a high-performance signal processor *by C. Galand, C. Couturier, G. Platel, and R. Vermot-Gauchy, p. 147.* In this paper, we discuss the implementation of a medium-bit-rate linear prediction baseband coder on an IBM bipolar signal processor prototype having a high processing capacity. We show that the implementation of our algorithm requires a processing load of 5 MIPS, with a program size of 5K instructions. We then discuss the application of our coder in a normal telephone environment, which requires mu-law to linear PCM conversion and other signal processing functions such as voice activity detection, automatic gain control, echo control, and error recovery. Quality evaluation tests are also reported which show that this type of coder, operating at 7.2 kbps, allows the transmission of telephone speech with communications quality. Moreover, obtained intelligibility scores and speaker recognition levels are high enough to demonstrate that this coder is a good candidate for telephony applications such as digital trunk transmissions, satellite speech communications, secure voice communications, and audio distribution systems.

Personal Instrument (PI)—A PC-based signal processing system *by G. Shichman, p. 158.* The Personal Computer (PC) technology has seen an enormous growth in the last two years. Although increasingly viewed as a major productivity tool, the PC is likely to be limited for computation-intensive tasks such as telecommunications and improved human-factors I/O. At the same time, there has been another evolving technology—VLSI realization of general-purpose signal processor (SP) engines which are capable of boosting the performance levels of standard PCs by almost two orders of magnitude. With SPs in PCs, we now see tremendous opportunities for distributing computation-intensive tasks away from high-performance mainframe computers; previously formidable tasks such as speech coding and recognition, pattern and scene analysis, spectral analysis, high bit-rate communication, and the like are now all computable by utilizing a single VLSI module embedded in any standard personal computer. This combination of a general-purpose CPU and superfast real-time coprocessor is likely to be key to the future functions and success of advanced workstations. A signal processing subsystem with real-time data acquisition and control capabilities has been developed for the IBM PC at the IBM Thomas J. Watson Research Laboratory and is the topic of this paper.

An approach to DFT calculations using standard microprocessors *by J. T. Rayfield and H. F. Silverman, p. 170.*

The use of the DFT as an everyday tool is now commonplace, principally due to advances in hardware technology. Special-purpose VLSI chips for signal processing are available. In this paper, we describe an approach which marries the Winograd Fourier Transform Algorithm (WFTA) with a state-of-the-art 16-bit general-purpose microprocessor for the purpose of DFT calculation. The heart of the approach is the real-input 240-point WFTA, which has been carefully optimized for time and space. In particular, an implementation for the 10-MHz M68000 executes in 10.8 ms. A simple hardware module is described which implements the optimized software. The use of the module for the inverse transform and for the complex case is also discussed. Advantages to the approach taken in this paper are the low cost/performance ratio and the general-purpose nature of the system that allows many non-signal-processing functions to be performed by the microprocessor.

Image processing applications for geologic mapping *by M. Abrams, A. Blusson, V. Carrere, T. Nguyen, and Y. Rabu, p. 177.* The use of satellite data, particularly Landsat images, for geologic mapping provides the geologist with a powerful tool. The digital format of these data permits applications of image processing to extract or enhance information useful for mapping purposes. Examples are presented of lithologic classification using texture measures, automatic lineament detection and structural analysis, and use of registered multisource satellite data. In each case, the additional mapping information provided relative to the particular treatment is evaluated. The goal is to provide the geologist with a range of processing techniques adapted to specific mapping problems.

Parameter reduction and context selection for compression of gray-scale images *by S. J. P. Todd, G. G. Langdon, Jr., and J. Rissanen, p. 188.* In the compression of multilevel (color or gray) image data, effective compression is obtained economically by judicial selection of the predictor and the conditioning states or contexts which determine what probability distribution to use for the prediction error. We provide a cost-effective approach to the following two problems: (1) to reduce the number of coding parameters to describe a distribution when several contexts are involved, and (2) to choose contexts for which variations in prediction error distributions are expected. We solve Problem 1 (distribution description) by a *partition* of the range of values of the outcomes into *equivalence classes,* called *buckets.* The result is a special *decomposition* of the error range. Cost-effectiveness is achieved by using the many contexts only to predict the bucket (equivalence class) probabilities. The probabilities of the value within the bucket are assumed to be independent of the context, thus enormously reducing the number of coding parameters involved. We solve Problem 2 (economical contexts) by using the *buckets* of the surrounding pixels as components of the conditioning class. The bucket values have the desirable properties needed for the error distributions.

Regenerative simulation methods for local area computer networks *by P. J. Haas and G. S. Shedler, p. 194.* Local area computer network simulations are inherently non-Markovian in that the underlying stochastic process cannot be modeled as a Markov chain with countable state space. We restrict attention to local network simulations whose underlying stochastic process can be represented as a generalized semi-Markov process (GSMP). Using "new better than used" distributional assumptions and sample path properties of the GSMP, we provide a "geometric trials" criterion for recurrence in this setting. We also provide conditions which ensure that a GSMP is a regenerative process and that the expected time between regeneration points is finite. Steady-state estimation procedures for ring and bus network simulations follow from these results.

Volume 29, Number 3, 1985

Advanced bipolar transistor modeling: Process and device simulation tools for today's technology *by R. W. Knepper, S. P. Gaur, F.-Y. Chang, and G. R. Srinivasan, p. 218.* A series of programs have been developed and linked together for doing advanced transistor modeling. The strategy begins with a process modeling program, SAFEPRO, for predicting two-dimensional impurity profiles. These are input to a two-dimensional device physics modeling program, 2DP, for generating device electrical characteristics. A three-dimensional distributed device model is then assembled by a model generator program (MGP) which, in turn, is used to derive a lumped equivalent-circuit model for numerical circuit analysis. The tools make it possible to do process sensitivity studies, perform process and device optimization, and provide early feedback on technology performance. The approach has recently been used to examine and compare various technologies at IBM.

Two-dimensional process modeling: A description of the SAFEPRO program *by R. R. O'Brien, C. M. Hsieh, J. S. Moore, R. F. Lever, P. C. Murley, K. W. Brannon, G. R. Srinivasan, and R. W. Knepper, p. 229.* This paper describes the development, testing, and application of a finite element program which simulates the processes used in manufacturing transistors. The profiles calculated by the program can be input directly into a device analysis program. The paper includes a description of the physical phenomena modeled and explains the choice of the particular numerical methods used to solve the resulting equations. It shows an example of the application of the program to the design and sensitivity study of a submicrometer shallow-junction bipolar transistor and presents results obtained when an oxide is grown on boron-doped silicon.

Two-dimensional device simulation program: 2DP *by S. P. Gaur, P. A. Habitz, Y.-J. Park, R. K. Cook, Y.-S. Huang, and L. F. Wagner, p. 242.* Mathematical details of a two-dimensional semiconductor device simulation program

are presented. Applicability of the carrier transport model to shallow junction bipolar transistors is discussed. Use of this program to optimize device structures in new bipolar technology is illustrated by presenting calculated device characteristics for variations in a few selected process conditions. Software links that automatically transfer data from a two-dimensional process simulation program and to a quasi-three-dimensional device equivalent circuit model generation program are also discussed.

The generation of three-dimensional bipolar transistor models for circuit analysis *by F.-Y. Chang and L. F. Wagner, p. 252.* The results of a two-dimensional bipolar numerical device-analysis program are processed by identifying three regions of the transistor: the intrinsic transistor, the sidewall transistor, and the extrinsic base collector diode. The key parameters which describe each of these regions are extracted using a linking program and fed into a quasi-three-dimensional device analysis program referred to here as the Model Generation Program (MGP). The MGP first generates a large equivalent-circuit distributed network which simulates the three-dimensional geometry of an actual transistor. This distributed network is then analyzed using the existing Advanced Statistical Analysis Program (ASTAP) for circuit analysis. Finally the MGP extracts parameters from the ASTAP analysis to characterize the elements of a lumped-model equivalent circuit, which is then suitable for the circuit design of large-scale integrated circuit chips. The MGP is sufficiently flexible that transistors with a variety of geometries can be generated without repeating the two-dimensional analysis.

FEDSS—A 2D semiconductor fabrication process simulator *by L. Borucki, H. H. Hansen, and K. Varahramyan, p. 263.* The main features of the finite element semiconductor process simulator FEDSS are described, with emphasis on a recently added capability for generalized 2D oxidation with impurity redistribution in oxide and silicon. Examples are given that demonstrate the ability of the program to oxidize various structures using a model based on steady-state oxidant diffusion and incompressible viscous oxide flow. Impurity profiles and contours are also shown in both neutral and oxidizing ambients, along with several comparisons with data or with the program SUPREM II.

VLSI wiring capacitance *by P. E. Cottrell and E. M. Buturla, p. 277.* Accurate prediction of device current and the capacitance to be driven by that current is key to the design of integrated logic and memory circuits. A finite-element algorithm is described which simulates the capacitance of structures with general shape in two or three dimensions. Efficient solution of the linear equations is provided by the incomplete Cholesky conjugate gradient method. The model is used to simulate the wiring capacitance of a 1.25-micrometer VLSI technology. The predicted capacitances of closely spaced first-metal polycide-gate and second-metal conductors used in this technology agree with measured results. The simulated three-dimensional capacitance of a second-metal line crossing a first-metal line is twice that found when estimated by two-dimensional models. The effect of line-to-line capacitance on the noise margin of logic circuits and on the signal in a dynamic RAM is examined. This capacitance presents a limit to wiring density for logic circuits and is a significant signal detractor in dynamic RAMs with closely spaced metal or diffused bit lines.

Semiconductor device simulation using generalized mobility models *by S. E. Laux and R. G. Byrnes, p. 289.* A method for discretizing the semiconductor transport equations using generalized mobility models is developed as an extension of the Scharfetter-Gummel finite difference approach. The method is sufficiently general to be applicable to nearly arbitrary empirical mobility models (including those for MOS surface effects) and may be used on a variety of mesh types in two or three dimensions. The impact of generalized mobility models on the sparsity of our resulting discrete equations is discussed. Convergence rate of a Newton's method linearization of the nonlinear system of equations is measured and interpreted. Some computational results from a study of short-channel MOSFETs are presented to illustrate the approach.

Animation and 3D color display of multiple-variable data: Application to semiconductor design *by E. J. Farrell, S. E. Laux, P. L. Corson, and E. M. Buturla, p. 302.* The increasing complexity of digital simulations requires more effective techniques to display and interpret the voluminous outputs. Advanced digital processing workstations and high-resolution color monitors permit a wide range of new techniques for use in examining the global characteristics of each output variable and their interrelationships with other variables. In this investigation, animation, 3D display, and multiple-window imaging have been shown to be effective in interpreting multiple-variable data sets, both scalar and vector. These display methods are used in the solution of two specific semiconductor design problems: the avalanche breakdown of an n-MOSFET and an alpha particle hit on an npn transistor. With these techniques the user can more fully utilize the results of these long and costly computations, making these methods a powerful addition to existing techniques for imaging data.

An experimental comparison of the head/disk interface dynamics in 5¼- and 8-inch disk drives *by G. Bouchard, D. B. Bogy, and F. E. Talke, p. 316.* A laser Doppler vibrometer is used to measure the head/disk interface dynamics in computer disk drives. The stability of the head under steady operating conditions is compared between a 5¼-inch and two different 8-inch "Winchester" drives. In the larger drives, high-frequency vibrations (between 5 and 10 kHz) are detected on the slider which are not present in the smaller drive. These vibrations have amplitudes on the order of magnitude of the head/disk spacing and are related to the

rolling and pitching modes of the slider. The vibrations of the disk, suspension, and actuator arm are also investigated and correlated with the results obtained on the slider.

Volume 29, Number 4, 1985

Analysis of manufacturing systems by the Research Queueing Package *by W.-M. Chow, E. A. MacNair, and C. H. Sauer, p. 330.* Many aspects of manufacturing systems can be analyzed using simulation to model the system's behavior. The Research Queueing Package (RESQ) is a tool developed to construct and solve models of systems with jobs contending for service from many resources. The capabilities of RESQ are described in order to understand the model elements which are available for representing manufacturing systems. Then an analysis of several work-in-process (WIP) policies is presented using RESQ models solved by simulation. Four WIP management policies are analyzed and compared for a future assembly manufacturing line: (1) a push system, (2) a pull system, (3) a transfer line, and (4) a closed loop system.

Integrated Manufacturing Modeling System *by H. Engelke, J. Grotrian, C. Scheuing, A. Schmackpfeffer, W. Schwarz, B. Solf, and J. Tomann, p. 343.* The objective of the Integrated Manufacturing Modeling System (IMMS) project was to build a generic software package for manufacturing modeling. The package allows interactive model building, testing, and experimentation. IMMS integrates the fundamental components of modeling, namely system description, data acquisition and management, graphic input and output, and performance simulation. The fundamental idea is to describe the flow of information, materials, and resources by using parametric standard model building blocks. Inputting is done by placing function symbols on a graphic screen and by filling in blanks for parameter data. The output includes all common performance measures, such as utilization and work in process (WIP), in graphic form. The simulation language used to implement the model building blocks is RESQ2. The Integrated Manufacturing Modeling System combines the advantage of simulation modeling with the advantages of end user programming capability. Simulation results allow intelligent decision making based on analysis (instead of intuitive experience), and an end user package reduces the turnaround time for modeling and eliminates the need for specialized programming skill. All this translates into faster and better-quality manufacturing system design and operation.

An algorithm for carrier routing in a flexible material-handling system *by C. L. Haines, p. 356.* Flexible material-handling systems for manufacturing have the capability of moving articles or carriers between process stations in different sequences. The traditional method for controlling the routing of carriers is to determine, in advance, all of the useful paths within the system, and store the information in a central computer until needed. This article describes a routing algorithm that determines the correct turns a carrier should make while it is in motion. Making routing decisions does not require a global knowledge of the system's layout, because a method of numbering stations within the system which reflects its natural path of flow is employed. A brief survey of contemporary material-handling mechanisms is provided. The implementation of the algorithm using distributed controllers is discussed.

Precise manipulation with endpoint sensing *by R. H. Taylor, R. L. Hollis, and M. A. Lavin, p. 363.* This paper describes recent work on manipulation strategies that rely on "coarse-fine" robot hardware and direct sensing of part-workpiece relationships. The experiments reported use an extremely precise, high-bandwidth planar "wrist" and an industrial vision system to perform accurate alignment of small parts. The system architecture, experimental hardware, and programming methods employed are all discussed.

On the analysis and design of CUSUM-Shewhart control schemes *by E. Yashchin, p. 377.* In recent years cumulative sum (CUSUM) control charts have become increasingly popular as an alternative to Shewhart's control charts. These charts use sequentially accumulated information in order to detect out-of-control conditions. They are philosophically related to procedures of sequential hypothesis testing (the relation being similar to that existing between Shewhart's charts and classical procedures for hypothesis testing). In the present paper we present a new approach to design of CUSUM-Shewhart control schemes and analysis of the associated run length distributions (under the assumption that the observations correspond to a sequence of independent and identically distributed random variables). This approach is based on the theory of Markov chains and it enables one to analyze the ARL (Average Run Length), the distribution function of the run length, and other quantities associated with a CUSUM-Shewhart scheme. In addition, it enables one to analyze situations in which out-of-target conditions are not present initially, but rather appear after a substantial period of time during which the process has operated in on-target mode (steady state analysis). The paper also introduces an APL package, DARCS, for design, analysis, and running of both one- and two-sided CUSUM-Shewhart control schemes and gives several examples of its application.

Short-term production scheduling of an automated manufacturing facility *by S. B. Gershwin, R. Akella, and Y. F. Choong, p. 392.* We describe a new implementation of the Kimemia-Gershwin hierarchical policy for the real-time scheduling of flexible manufacturing systems. Major improvements result at all three levels of the policy. The algorithm simplification, resulting in substantial reductions of off-line and on-line computation time, is reported, as is the improvement in performance through the elimination of chattering. Simulation results based on a detailed model of a printed circuit card assembly facility are summarized.

Scheduling algorithms for flexible flow lines *by R. J. Wittrock, p. 401.* This paper discusses scheduling algorithms

for a certain kind of manufacturing environment, called the "flexible flow line." Two scheduling problems are considered. "Loading" decides when each part should be loaded into the system. "Mix allocation" selects the daily part mix. The goals are to maximize throughput and reduce WIP. New heuristic algorithms specially suited to solve these problems in the context of a flexible flow line are described. The paper also discusses experience with the use of an experimental implementation of these algorithms to solve such problems arising in a real production line.

A prototype manufacturing knowledge base in Syllog *by C. Fellenstein, C. O. Green, L. M. Palmer, A. Walker, and D. J. Wyler, p. 413.* This paper describes a prototype knowledge base for manufacturing planning, which we have built using a knowledge system shell called Syllog. We describe a Tester Capacity Planning and Yield Analysis task, knowledge needed for a part of the task, and the use of the knowledge in the Syllog system. We report that the sometimes difficult process of knowledge acquisition turned out, in this case, to be straightforward. Knowledge acquisition and knowledge use are done in the same language in Syllog.

Large-scale scientific application programs in chemistry and physics on an experimental parallel computer system *by G. Corongiu and J. H. Detrich, p. 422.* We present and discuss an experimental distributed system consisting of two IBM 4341s, an IBM 4381, and ten FPS-164 attached processors, configured to allow parallel execution of a single large-scale calculation on multiple processors. A number of our application programs have been converted to run on this system, and the strategy for this conversion is outlined in sufficient detail to facilitate the development of tests using other scientific and engineering computer application programs. Our tests, though limited to certain biochemical and physicochemical problems, demonstrate the versatility, flexibility, and accessibility of this system, and we find performance comparable to that of today's supercomputers, which suggests that our approach can provide a practical answer to many large-scale computer applications.

Volume 29, Number 5, 1985

Preface *by W. F. Voit, Jr., p. 442.*

Design and implementation of the SELECTRIC System/2000 *by T. H. Williams, p. 443.* This paper presents an overview of the papers in this issue of the *IBM Journal of Research and Development* covering the design and implementation of the ®SELECTRIC System/2000 Typewriter/Printer products. The SELECTRIC System/2000 products comprise a nonimpact typewriter and printer, two impact typewriters, and an impact printer which use printwheel technology. The development approach for the SELECTRIC System/2000 products, which included design for automation, introduction of new technologies, and product development concurrent with manufacturing, was accomplished by the use of common architecture, hardware, and software.

Resistive ribbon thermal transfer printing: A historical review and introduction to a new printing technology *by K. S. Pennington and W. Crooks, p. 449.* This paper describes a new high-quality thermal transfer printing process in which a printhead consisting of a linear array of small-diameter electrodes produces highly localized Joule heating of a resistive thermal transfer printing ribbon. The heat generated in the resistive ribbon results in the melting of a thermoplastic ink which is then transferred to a printable medium, such as paper, by contact. The origins of the technology in IBM are discussed, together with a description of the resistive ribbon materials and structure, the printhead, and some experimental printer performance values.

Implementation of the resistive ribbon technology in a printer and correcting typewriter *by S. Applegate, J. Bartlett, A. Bohnhoff, A. Campbell, and J. Molloy, p. 459.* This paper describes key technology implementation details and performance characteristics of a printer and typewriter using the resistive ribbon technology. The work describes the first commercial application of this print and correction technology. Key parameters necessary for proper system function such as current, various forces, velocities, and component integration are discussed. The rationales behind various compromises and problem solutions are given. A discussion of the characteristics of the print produced, along with application strengths and weaknesses, completes the paper.

Characterization of a resistive ribbon thermal transfer printing process *by T. G. Twardeck, p. 470.* Resistive ribbon thermal transfer printers transfer ink from a ribbon to paper as the result of localized Joule heating of the ribbon structure. For this printing process, this paper discusses the voltage-versus-current response of the electrode-ribbon current path, the temperature distributions throughout the ribbon structure, and the correlation of print response with electrical input power and average ribbon temperatures. Nominal input power per electrode is approximately 190 mW. For input power near this level, thermal models predict that ribbon materials which pass directly under energized electrodes reach the highest temperatures; the hottest zone in the ribbon surrounds the composite-aluminum interface. Approximately 0.1 mm downstream from the electrodes, the heated ribbon materials come to nearly constant temperature. The area of the printed image correlates with this average ribbon temperature and input power.

Development of a membrane switch-type full-travel tactile keyboard *by S. F. DeFosse, G. T. Williams, D. A. Gostomski, Jr., and R. H. Cobb, p. 478.* This paper describes the approaches to design and the rationale that successfully satisfied the requirements for a full-travel keyboard with usage exceeding 10 million actuations per character, satisfactory

N-key roll-over, and phantom key control. Emphasis was placed on technical understanding of the effects of all material and design decisions. Interactions among design, material, and processing variables were revealed through statistical parameter modeling and environmental exposure studies. This knowledge facilitated the control of critical parameters to permit an order-of-magnitude reduction of actuation forces and tolerances. Product reliability was achieved through evaluation, environmental protection features, and stringent process controls. This paper highlights the design, materials, and processing aspects of the membrane switch developed for low-force keyboard applications. It also discusses effects of environmental factors on the individual components and overall system function.

System control for a printwheel typewriter *by R. D. Mayo, p. 488.* This paper presents the design goals and architecture of the IBM ®WHEELWRITER typewriters. Some of the development efforts and resulting technical innovations, such as a unique print hammer design which minimizes sensitivities to current variations, are discussed. Included is a discussion of a variable-reluctance stepper motor driver that has selectable damping. A novel scheme for initializing the printwheel and escapement motors is given; this includes sensing the font weight. The electronic architecture of the typewriter and the design of a simplified algorithm to handle the many different keyboards that can be attached to the machine with a minimum amount of data storage are explained.

System controls for a resistive ribbon printer *by A. Bohnhoff, D. Croley, S. Dyer, T. Green, R. Maddox, and L. Struttmann, p. 494.* The system controls for a printer using the resistive ribbon print technology involve conventional requirements, such as moving the print mechanism relative to the paper, with a new requirement, controlling the electrical energy to the ribbon, an electrothermal component. Other special requirements are dictated by using the same ribbon for hard copy print/erase while ensuring that the print and erase operations are acceptable to the user. This paper discusses the design and performance of the system controls for a resistive ribbon printer that was developed for use in an interactive typewriter application and as an output printer for a personal computer.

Ink temperatures in resistive ribbon thermal transfer printing *by T. C. Chieu and O. Sahni, p. 509.* A knowledge of ink temperatures is important in thermal transfer printing technologies. This paper reports on an experimental technique which uses an infrared radiometric microscope to measure the temperature of ink deposited by the resistive ribbon process on transparent substrates. A detailed examination has been made of the spatial and temporal profiles of ink temperatures as a function of input current, printing speed, substrate materials, and number of active electrodes. The results on a ®Kapton substrate permit estimation of the ink temperatures reached during printing on paper. The peak ink temperatures are observed to depend linearly on input current and inversely on an approximately linear function of writing speed from 2 to 8 inches per second. Based on a phenomenological model, these results lead to a functional relationship among speed, print current, and ink temperature during printing. The model permits projections to be made of the current required over a wide range of printing speeds.

Electrical properties of resistive ribbon *by K. K. Shih and D. B. Dove, p. 519.* In resistive ribbon thermal transfer printing, a printhead consisting of an array of electrodes passes current into a thin ribbon to generate heat for transferring ink to paper. The ribbon is made of a polymeric material containing carbon black so as to be conducting, and has aluminum deposited on one side of the ribbon for a base contact. In this paper, the electrical conduction processes within the ribbon are discussed. Current-voltage measurements have been made with electrodes of various types in order to separate effects due to contact resistance, aluminum/resistive ribbon interfacial resistance, and bulk conduction in the resistive ribbon. Measurements have been made over a range of frequency and temperature to determine the basic conduction mechanisms. A model of conduction is presented that is in qualitative agreement with the data.

Thermal behavior of resistive ribbon for single-stylus excitation *by R. A. Laff and C. D. Makowka, p. 527.* This paper provides a quantitative description of the heating and cooling behavior of the resistive ribbon used in resistive ribbon thermal transfer printing. Since the focus is upon the ribbon, this description has been facilitated by substituting a single, tapered tungsten stylus as a model for a single printhead element. The experiment used an infrared spot pyrometer to measure ribbon surface temperatures downstream from the stylus while a laminate of ribbon and paper was continuously moved beneath the stylus and subjected to current pulses. Measured cooling rates under steady-state excitation at different ribbon velocities showed a behavior consistent with two simple analytic models which describe heat loss into the stylus during heating and two-dimensional diffusion into a half-plane during cooling. A detailed, time-dependent computer simulation using finite-element methods was used to provide a more detailed description of the process by taking into account the local geometry of the heat input distribution and the layered nature of the ribbon-paper laminate. The resulting three-dimensional temperature distributions are given for the steady-state case.

Volume 29, Number 6, 1985

Adaptive cross-parity (AXP) code for a high-density magnetic tape subsystem *by A. M. Patel, p. 546.* This paper describes an error-correction system, called adaptive cross-parity (AXP) code, for the IBM 3480, a new high-density 18-track tape storage subsystem. Redundancy is applied to two interleaved sets of nine tracks in the same proportion as that in the previous IBM 3420 tape machines.

The coding structure, however, is simpler, for it avoids the complex computations of Galois fields. The coding structure is based on a concept of interacting vertical and cross-parity checks, where the cross-parity checks span both sets of tracks and are used in either set in an adaptive manner. As a result, the overall error-correcting capability is substantially improved without increasing the redundancy. Decoding, in which simple parity equations are processed, is designed to progress iteratively. By means of adaptive use of redundancy, the new method corrects up to three known erroneous tracks in any one set of nine tracks and up to four known erroneous tracks in the two sets together. The code also identifies the first unknown erroneous track in each of the two sets, and subsequently identifies the second unknown erroneous track in one of the two sets while providing correction for all these tracks. The result is generalized for a system with any number of tracks divided into a multiple number of unequal sets.

Write equalization in high-linear-density magnetic recording *by R. C. Schneider, p. 563.* For many years, equalization has been used on the read side of a magnetic-recording channel to obtain a desired signal shape at the detector. Compensation on the write side has, for the most part, been limited to moving transition locations to offset read-signal peak shifts. This paper presents a new method of equalization on the write side through the addition of pulses at strategic locations on the write waveform. The resulting write current continues to be a two-level signal, so ac bias is not required. A linear transfer function can be derived for these write equalizers. This enables the recording-channel designer to partition the equalization more optimally between the write and read sides. The principal benefit of write equalization is that the read-flux-amplitude differences between high and low densities are significantly reduced. This permits maximum use of the linear operating region of the magnetoresistive read head. By providing high-frequency boosts on the write side, write equalization can reduce high-frequency noises at the read detector. Test results of channel linearity, as well as read signal waveshapes, are presented.

Image processing by simulated annealing *by P. Carnevali, L. Coletti, and S. Patarnello, p. 569.* It is shown that simulated annealing, a statistical mechanics method recently proposed as a tool in solving complex optimization problems, can be used in problems arising in image processing. The problems examined are the estimation of the parameters necessary to describe a geometrical pattern corrupted by noise, the smoothing of bi-level images, and the process of halftoning a continuous-level image. The analogy between the system to be optimized and an equivalent physical system, whose ground state is sought, is put forward by showing that some of these problems are formally equivalent to ground state problems for two-dimensional Ising spin systems. In the case of low signal-to-noise ratios (particularly in image smoothing), the methods proposed here give better results than those obtained with standard techniques.

Jitter accommodation in token-passing ring LANs *by R. J. S. Bates and L. A. Sauer, p. 580.* In a token-passing ring Local Area Network (LAN) each message accumulates phase jitter as it travels around the ring. Unlike typical digital transmission systems, which tend to have random data traffic, ring systems carrying computer-generated traffic may have long strings of repetitive-pattern data. This traffic produces a jitter amplitude which is a function of the message statistics and the transmission characteristics of the physical layer. For the system to be stable, this jitter must be controlled. This paper describes the repetitive-pattern jitter generation and accumulation process and gives a methodology for designing the physical layer components to accommodate it.

Analysis of the holding current in CMOS latch-up *by H. Matino, p. 588.* The holding current in CMOS latch-up with or without well and/or substrate bias has been examined. Measurements indicate that the holding current increases significantly with reverse bias and low shunting base resistance. It is shown that a previous equation for the holding current is inaccurate, and a new equation for holding current with bias is presented.

Volume 30, Number 1, 1986

Knowledge systems: Principles and practice *by A. Walker, p. 2.* We describe what is expected of a knowledge-based expert system, and the components from which such a system is constructed. We give a view of how an interplay between principles and practice has helped the knowledge system field to develop, and we give simple examples to show that reasoning techniques based on formal logic are now starting to provide a useful coupling between scientific and engineering work in the field. The examples are about logic programming, knowledge representation, judgmental reasoning, and about three methods by which a system can acquire the knowledge it needs.

A continuous real-time expert system for computer operations *by R. L. Ennis, J. H. Griesmer, S. J. Hong, M. Karnaugh, J. K. Kastner, D. A. Klein, K. R. Milliken, M. I. Schor, and H. M. Van Woerkom, p. 14.* The Yorktown Expert System/MVS Manager (or YES/MVS for short) is a continuous real-time expert system that exerts active control over a computing system and provides advice to computer operators. YES/MVS provides advice on routine operations and detects, diagnoses, and responds to problems in the computer operator's domain. This paper discusses the YES/MVS system, its domain of application, and issues that arise in the design and development of an expert system that runs continuously in real time.

Interfaces for knowledge-base builders' control knowledge and application-specific procedures *by P. Hirsch, W. Katke, M. Meier, S. Snyder, and R. Stillman, p. 29.* Expert System Environment/VM is an expert system shell—a general-purpose system for constructing and executing expert system applications. An application expert has both factual knowledge about an application and knowledge about how that factual knowledge should be organized and processed. In addition, many applications require application-dependent procedures to access databases or to do specialized processing. An important and novel part of Expert System Environment/VM is the technique used to allow the expert or knowledge-base builder to enter the control knowledge and to interface with application-dependent procedures. This paper discusses these high-level interfaces for the knowledge-base builder.

A theory for the representation of knowledge *by F. Guenthner, H. Lehmann, and W. Schönfeld, p. 39.* How to represent knowledge is one of the key questions in the construction of expert systems. Its solution depends on a number of factors, the most important of which are how knowledge is to be acquired and how it is to be used. Since we are interested in the use of natural language for communication with computers, we require from a formalism suggested for knowledge representation that it be suitable as a target for the systematic translation from natural language expressions. We want to propose a theory, called *Discourse Representation Theory (DRT),* which was originally developed by Kamp to analyze natural language discourse, as a means to represent knowledge in an expert system. With Discourse Representation Theory it has been possible to solve certain cases of contextual relations which have puzzled linguists and logicians for a long time. In this paper we give a precise definition of DRT and describe the rules used to translate from natural language to Discourse Representation Structures (DRSs), the central notion of the theory. We show how the notation used in DRT relates to standard predicate logic and define its deductive theory. We also outline ways of implementing DRT and the proof procedures we intend to use.

Implementing a semantic interpreter using conceptual graphs *by J. F. Sowa and E. C. Way, p. 57.* A parser applies grammar rules to generate a parse tree that shows the syntactic structure of a sentence. This paper describes a semantic interpreter that starts with a parse tree and generates conceptual graphs that represent the meaning of the sentence. To generate conceptual graphs, the interpreter joins *canonical graphs* associated with each word of input. The result is a larger graph that represents the entire sentence. During the interpretation, the parse tree serves as a guide to show how the graphs are joined. Both the front-end parser and the back-end semantic interpreter are written in the Programming Language for Natural Language Processing (PLNLP).

Conceptual graphs for semantics and knowledge processing *by J. Fargues, M.-C. Landau, A. Dugourd, and L. Catach, p. 70.* This paper discusses the representational and algorithmic power of the conceptual graph model for natural language semantics and knowledge processing. Also described is a Prolog-like resolution method for conceptual graphs, which allows one to perform deduction on very large semantic domains. The interpreter that we have developed is similar to a Prolog interpreter in which the terms are any conceptual graphs and in which the unification algorithm is replaced by a specialized algorithm for conceptual graphs.

Storing and evaluating Horn-clause rules in a relational database *by G. van Emde Boas and P. van Emde Boas, p. 80.* This paper describes a practical approach to storing and evaluating Horn-clause rules in a relational database system. The intention is to give a complete outline of what needs to be added to an existing relational database system to allow it to support full logic programming functions. Implementation issues for each new function are discussed. We show how Horn-clause rules can be translated into database commands without recourse to semantics and how their evaluation can be performed in the database itself. This brings the complete logic programming environment within reach of the database management system, allowing data and rule sharing, concurrency control, recovery procedures, etc., to be used. New is that the complete logic programming environment is incorporated into the database system. IBM Business System 12, extended in this way, may be a suitable vehicle for expert system applications.

Structures of rule-based belief functions *by W. F. Eddy and G. P. Pei, p. 93.* Shafer's theory of evidential reasoning has recently received much attention as a possible model for probabilistic reasoning in expert system applications. This paper discusses the particular difficulties of implementing Shafer's belief functions in the context of the most common form of expert system, rule-based systems. The two most important problems are: the representation of the expert's subjective degrees of belief corresponding to his expressed rules, and the computational complexity of the inference mechanism for combining evidence. We argue that a potential approach for dealing with both problems is given by introducing constraints on the structure of the belief functions. These constraints, along with the expressed rules and the elicited belief values, form the expert's total knowledge.

An experimental computer architecture supporting expert systems and logic programming *by H. Diel, N. Lenz, and H. M. Welsch, p. 102.* This paper describes a set of function primitives which have been designed for support of expert systems and logic programs. The functions could be offered as part of the computer architecture by implementing them in microcode and partially in hardware. The functions are primarily (but not exclusively) oriented towards support of logic programming languages such as Prolog for implementing expert systems. Particular emphasis is given to supporting the parallel execution of expert system applications by multiple processors. The concepts described are based on the Concurrent Data Access Architecture (CDAA). It is shown that OR-parallelism, as well as AND-parallelism, can be supported.

Volume 30, Number 2, 1986

New scalar and vector elementary functions for the IBM System/370 *by R. C. Agarwal, J. W. Cooley, F. G. Gustavson, J. B. Shearer, G. Slishman, and B. Tuckerman, p. 126.* Algorithms have been developed to compute short- and long-precision elementary functions: SIN, COS, TAN, COTAN, LOG, LOG10, EXP, POWER, SQRT, ATAN, ASIN, ACOS, ATAN2, and CABS, in scalar (28 functions) and vector (22 functions) mode. They have been implemented as part of the new VS FORTRAN library recently announced along with the IBM 3090 Vector Facility. These algorithms are essentially table-based algorithms. An important feature of these algorithms is that they produce bitwise-identical results on scalar and vector System/370 machines. Of these, for five functions the computed value result is always the correctly rounded value of the infinite-precision result. For the rest of the functions, the value returned is one of the two floating-point neighbors bordering the infinite-precision result, which implies exact results if they are machine-representable. For the five correctly rounded elementary functions, scalar and vector algorithms are designed independently so as to optimize performance in each case. For other functions, the bitwise-identical constraint leads to algorithms which compromise between scalar and vector performance. We have been able to design algorithms where this compromise is minimal and thus achieve very good performance on both scalar and vector implementations. For our test measurements on high-end System/370 machines, our scalar functions are always faster (sometimes by as much as a factor of 2.5) as compared to the old VS FORTRAN library. Our vector functions are usually two or three times faster than our scalar functions.

Fourier transform and convolution subroutines for the IBM 3090 Vector Facility *by R. C. Agarwal and J. W. Cooley, p. 145.* A set of highly optimized subroutines for digital signal processing has been included in the Engineering and Scientific Subroutine Library (ESSL) for the IBM 3090 Vector Facility. These include FORTRAN-callable subroutines for Fourier transforms, convolution, and correlation. The subroutines are carefully designed and tuned for optimal vector and cache performance. Speedups of up to 9½ times over scalar performance on the 3090 have been obtained.

A vectorizing Fortran compiler *by R. G. Scarborough and H. G. Kolsky, p. 163.* This paper describes a vectorizing Fortran compiler for the IBM 3090 Vector Facility. Opportunities for vectorization are investigated for eight levels of DO-loop nesting. Recurrences in inner loops do not prevent vectorization of outer loops. A least-cost analysis determines from the opportunities identified which specific vectorization will result in the fastest execution. The normal optimization phases of the compiler produce much of the information needed for the vectorization analysis.

Seismic migration on the IBM 3090 Vector Facility *by J. Gazdag, G. Radicati, P. Sguazzero, and H.-H. Wang, p. 172.* Seismic prospecting aims at determining the structure of the earth from indirect measurements. Acoustic wave fields are generated at the surface, penetrate the earth, and are backscattered by the earth's inhomogeneities. The data recorded at the surface are processed in a complex sequence of steps among which seismic migration plays an important role. This is a wave depropagation process that permits the localization in depth of the origin of the diffraction events measured (in time) at the surface. This paper presents an overview of the major wave-equation migration methods. The most frequently executed algorithms or kernels on which the execution speed depends most crucially are given particular attention. The speedup resulting from scalar-to-vector formulation is presented over wide ranges of dimensionality for linear tridiagonal equation solvers, Fourier Transforms, and convolution operations. The vectorizability and resulting speedup are also addressed in the case of migration schemes known as the Phase-Shift Method and the Phase Shift Plus Interpolation (PSPI) Method. It is shown that Fourier domain migration based on the phase-shift concept lends itself conveniently to multilevel parallelism on the 3090 Vector Facility (VF): vectorization of the innermost loops and

concurrent processing in the outer loops by means of the VS FORTRAN Version 2 Multitasking Facility.

Parallel iterative linear solvers for oil reservoir models *by I. Efrat and M. Tismenetsky, p. 184.* This paper suggests a new algorithm for solving sparse linear systems which is readily parallelized and is very efficient for matrices arising in such domains as reservoir modeling. The algorithm is, in fact, a variant of the incomplete block factorization technique, accelerated by Biconjugate Gradient iterations or by another acceleration method. Implementation on a vector computer such as the IBM 3090 Vector Facility is described. We also suggest some refinements of known preconditioning methods enabling their parallel computation. Numerical experiments are presented to display the performance of the suggested methods. Special attention is given to fully implicit, multiphase models which yield asymmetrical systems.

Monte Carlo photon transport on a vector supercomputer *by W. R. Martin, P. F. Nowak, and J. A. Rathkopf, p. 193.* The use of "event-based" algorithms for particle transport Monte Carlo methods has allowed the successful adaptation of these methods to vector supercomputers. An alternative algorithm for the specific application of photon transport in an axisymmetric inertially confined fusion plasma has been developed and implemented on a vector supercomputer. The new algorithm is described; its unique features are discussed and compared with existing vectorized algorithms for Monte Carlo. Numerical results are presented illustrating its efficiency on a vector supercomputer, relative to an optimized scalar Monte Carlo algorithm that was developed for this purpose.

Chemical and mechanical performance of flexible magnetic tape containing chromium dioxide *by R. Bradshaw, B. Bhushan, C. Kalthoff, and M. Warne, p. 203.* The development of a magnetic-recording tape using chromium dioxide suitable for use with the tape drive of the IBM 3480 Magnetic Tape Subsystem was found to require careful optimization of the chemical and mechanical properties of the coating. This paper discusses the role of chromium dioxide in the oxidative and hydrolytic degradation of the polyester-polyurethane binders used in most flexible tape coatings, and the subsequent necessity for careful binder selection to eliminate (for all practical purposes) these degradative effects. In addition to the chemical behavior, the paper discusses the role of the interaction of the chromium dioxide with the binder necessary to obtain mechanical performance suitable for the 3480 tape drive. A high modulus and a relatively high glass-transition temperature were found to be required to avoid changes in the frictional properties of the tape. The incorporation of a rough-textured back coat was found to control the change in the surface topography of the recording surface near the spool hub which arises as a result of compressive forces.

Volume 30, Number 3, 1986

Key-sequence data sets on indelible storage *by M. C. Easton, p. 230.* Methods for creating and maintaining key-sequence data sets without overwriting the storage medium are described. These methods may be applied to erasable or to write-once storage devices, and they are compatible with conventional device error-management techniques. All past values of data records are preserved in a data structure called a Write-Once B-Tree. Rapid random access is available to records by key value; rapid sequential access is available to records in key-sequence order. Moreover, queries requesting data as of a previous time are processed as rapidly as requests for current data. Access time is proportional to the logarithm of the number of *current* records in the database. Efficient methods for inserting, updating, and deleting records are described. Upper bounds for tree depth and for storage consumption are given and compared with results from simulation. It is concluded that, with rapidly improving storage technologies, indelible databases will become practical for many applications.

The average complexity of depth-first search with backtracking and cutoff *by H. S. Stone and P. Sipala, p. 242.* This paper analyzes two algorithms for depth-first search of binary trees. The first algorithm uses a search strategy that terminates the search when a successful leaf is reached. The algorithm does not use internal cutoff to restrict the search space. If N is the depth of the tree, then the average number of nodes visited by the algorithm is as low as $O(N)$ and as high as $O(2^N)$ depending only on the value of the probability parameter that characterizes the search. The second search algorithm uses backtracking with cutoff. A decision to cut off the search at a node eliminates the entire subtree below that node from further consideration. The surprising result for this algorithm is that the average number of nodes visited grows linearly in the depth of the tree, regardless of the cutoff probability. If the cutoff probability is high, then the search has a high probability of failing without examining much of the tree. If the cutoff probability is low, then the search has a high probability of succeeding on the leftmost path of the tree without performing extensive backtracking. This model sheds light on why some instances of NP-complete problems are solvable in practice with a low average complexity.

On-the-fly decoder for multiple byte errors *by A. M. Patel, p. 259.* Multiple-error-correcting Reed-Solomon or BCH codes in $GF(2^b)$ can be used for correction of multiple burst errors in binary data. However, the relatively long time required for decoding multiple errors has been among the main objections to applying these schemes to high-performance computer products. In this paper, a decoding procedure is developed for on-the-fly correction of multiple symbol (byte) errors in Reed-Solomon or BCH codes. A new decoder architecture expands the concept of Chien search of error locations into computation of error values as well, and creates a synchronous procedure for complete

on-the-fly error correction of multiple byte errors. Forney's expression for error values is further simplified, which results in substantial economies in hardware and decoding time. All division operations are eliminated from the computation of the error-locator equation, and only one division operation is required in the computation of error values. The special cases of fewer errors are processed automatically, using the corresponding smaller set of syndromes through a single set of hardware. The resultant decoder implementation is well suited for LSI chip design with pipelined data flow. The implementation is illustrated with an example.

Design and performance of a magnetic head for a high-density tape drive *by D. M. Cannon, W. R. Dempwolf, J. M. Schmalhorst, F. B. Shelledy, and R. D. Silkensen, p. 270.* The design of a magnetic head for use in the tape drive portion of the IBM 3480 Magnetic Tape Subsystem required that advances be made in track density and linear recording density while meeting requirements for signal quality, manufacturability, and reliability. The design that was developed to fulfill these needs combines ferrite pole-pieces, magnetoresistive read elements, and planar-deposited write turns. This paper describes the read and write elements of the head and the approach used in the selection of the design parameters.

Symmetric stochastic Petri nets *by L. A. Prisgrove and G. S. Shedler, p. 278.* The stochastic Petri net (SPN) model is well suited to formal representation of concurrency, synchronization, and communication. In this paper we focus on discrete event simulation methods for SPN models with special structure and define a symmetric SPN. Exploiting properties of a symmetric SPN and underlying regenerative process structure, we establish steady state estimation procedures based on independent, nonidentically distributed blocks of the marking process. We also establish estimation procedures for passage times in the symmetric SPN setting. These results lead to efficient estimation procedures for delay/throughput characteristics of ring networks with identical ports.

Compiling circular attribute grammars into Prolog *by B. Arbab, p. 294.* This paper describes an algorithm for compiling attribute grammars into Prolog. The attribute grammars may include inherited and synthesized attributes and contain recursive (circular) definitions. The semantics of the recursive definitions is defined in terms of a fixed-point finding function. The generated Prolog code stands in direct relation to its attribute grammar, where logical variables play the role of synthesized or inherited attributes. Thus an effective method for the execution of recursive attribute grammars has been defined and applied.

Least-squares storage-channel identification *by J. M. Cioffi, p. 310.* Pulse (dibit) and step (transition) responses for magnetic-storage channels are important for detection-circuitry design and for comparison of various media, heads, and other channel components. This paper presents a least-squares procedure that can be used to identify the dibit and transition responses from measurements of the read-head response to any known data sequence written on the medium. The method yields significantly higher-quality estimates for the dibit and step shapes than does determining these same characteristics by measuring the average response to isolated transition or by performing a Discrete Fourier Transform (DFT) on the response to a pseudorandom data pattern. The new method can be implemented off line but also can be made sufficiently efficient to be implemented with a microprocessor for use in self-optimizing (adaptive) channel detection circuitry.

A maximum-energy-concentration spectral window *by L. C. Barbosa, p. 321.* A time-discrete solution method is proposed for realization of a spectral window which is ideal from an energy concentration viewpoint. This window is one that concentrates the maximum amount of energy in a specified bandwidth and hence provides optimal spectral resolution.

On yield, fault distributions, and clustering of particles *by C. H. Stapper, p. 326.* Increasing the levels of semiconductor integration to larger chips with more transistors causes the fault and defect distributions of VLSI memory chips to deviate increasingly further from simple random Poisson statistics. The spatial distributions of particles on semiconductor wafers have been analyzed to gain insight into the nature of integrated circuit defect statistics. The analysis was done using grids of squares known as quadrats. It was found that the cluster parameter, which until now has been treated as a constant, did vary with quadrat area. The results also show that the deviation from Poisson statistics continues to increase into the realm of wafer-scale integration or WSI. Computer simulations were used to verify this effect.

Volume 30, Number 4, 1986

Preface *by P. Chaudhari, p. 354.*

Scanning tunneling microscopy *by G. Binnig and H. Rohrer, p. 355.* Presented here is an overview of the present status and future prospects of scanning tunneling microscopy. Topics covered include the physical basis of the scanning tunneling microscope, its instrumentation aspects, and its use for structural and spectroscopic imaging—on a scale which extends to atomic dimensions. Associated experimental and theoretical studies are reviewed, including several which suggest potential applicability of this new type of microscope to a relatively broad range of biological, chemical, and technological areas.

Squeezable tunneling junctions *by P. K. Hansma, p. 370.* Squeezable tunneling junctions establish the current state of the art for resistance stability of mechanically adjustable tunneling structures at $\Delta R/R \approx 0.1\%$. This is sufficient for use

in connection with spectroscopies as subtle as phonon spectroscopy, but it is marginal and cannot at present be maintained to high enough bias voltage to permit molecular vibrational spectroscopy. Squeezable junctions have been used for characterizing bulk samples and for differential capacitance-voltage analyses of semiconductors.

Electronic structure and tunneling current for chemisorbed atoms *by N. D. Lang, p. 374.* We discuss the tunneling current density in the vacuum region between two planar metal electrodes, one of which has an atom chemisorbed on its surface. The relation of this current distribution to the electronic structure of the adatom is analyzed. The study of this model problem leads to a better understanding of important aspects of the current flow in the scanning tunneling microscope. The emphasis of this work is not so much on the question of resolution discussed in other theoretical studies as on the characteristic signatures of chemically different atoms.

Application to biology and technology of the scanning tunneling microscope operated in air at ambient pressure *by A. M. Baró, R. Miranda, and J. L. Carrascosa, p. 380.* We have investigated the operation of the scanning tunneling microscope (STM) in air at atmospheric pressure, thus permitting the imaging of samples without the need for subjecting them to a vacuum environment. Clearly, this may be of practical importance for many types of samples having biological and technological interest. Imaging of biological samples has been found to be possible after deposition onto a flat structureless conducting substrate (highly oriented pyrolytic graphite). Three-dimensional profiles of structures derived from the virus designated as bacteriophage $\phi 29$ could thus be obtained. Other profiles have been obtained which indicate applicability to surfaces of technological interest: for example, in the measurement of the surface roughness of industrial components with increased precision, suggesting use of the STM as a new standard instrument for that purpose.

Tunneling microscopy from 300 to 4.2 K *by S. A. Elrod, A. Bryant, A. L. de Lozanne, S. Park, D. Smith, and C. F. Quate, p. 387.* A scanning tunneling microscope (STM) has been developed for operation over the full temperature range from 300 to 4.2K. At room temperature, the instrument has been used to produce topographic images of grain structure in a copper-titanium alloy foil and of atomic structure on a Pt(100) surface. At low temperatures, the instrument can be used in a new spectroscopic mode, one which combines the high spatial resolution of the STM with the existing technique of electron tunneling spectroscopy. This new capability has been demonstrated by using the microscope to probe spatial variations in the superconducting character of a niobium-tin alloy film.

A scanning tunneling microscope for surface science studies *by J. E. Demuth, R. J. Hamers, R. M. Tromp, and M. E. Welland, p. 396.* A new design is described for a scanning tunneling microscope intended for surface science studies. The performance of the microscope is evaluated from tunneling images obtained of the Si(111) 7×7 surface. Periodic structures, point defects, and grain boundary structures are observed with atomic-scale resolution and are discussed. Illustrations of various types of image processing and data display are presented.

Defects on the Pt(100) surface and their influence on surface reactions—A scanning tunneling microscopy study *by W. Hösler, R. J. Behm, and E. Ritter, p. 403.* Structural differences between a clean, reconstructed Pt(100) surface and one exhibiting chemical or structural irregularities have been identified by means of a scanning tunneling microscope. The (temperature-dependent) defect structure of a surface which had undergone a phase transition involving mass transport was characterized and compared to results obtained using other techniques. The catalytic activity of surface step sites was probed by the thermal decomposition of ethylene. The resulting surface roughening and buildup of carbon, which could be resolved in STM images, clearly showed that the decomposition proceeds from terrace edges.

Spectroscopy of electronic states of metals with a scanning tunneling microscope *by W. J. Kaiser and R. C. Jaklevic, p. 411.* We have constructed a scanning tunneling microscope (STM) and have obtained current-voltage derivative spectra from metal surfaces. For Au(111) we have observed an electronic surface state 0.4 eV below the Fermi energy. This state has previously been observed with photoemission and oxide tunneling experiments. We have also observed strong peaks in spectra obtained from Pd(111) which we identify with surface states and effects derived from bulk energy bands. In the voltage range investigated here, tunneling takes place through the entire vacuum gap between the metallic tip of the microscope and the surface of the sample being examined. The STM spectroscopy results reported here are compared with previous experimental work and theory. These intrinsic surface states are the first which have been observed with the STM and demonstrate its unique applicability to the investigation of surface electronic structure.

Some design criteria in scanning tunneling microscopy *by D. W. Pohl, p. 417.* Optimum functioning of a scanning tunneling microscope (STM) requires tip-to-sample position control with picometer precision, a rough and fine positioning capability in three dimensions, a scanning range of at least 100 times the lateral resolution, a scanning speed as high as possible, and also, preferably, simplicity of operation. These requirements have to be satisfied in the presence of building vibrations with up to micrometer-size amplitudes, temperature drift, and other perturbations. They result in design rules, presented and discussed here, for the optimization of damping, stiffness, electrical control circuitry, and the performance of the piezoelectric actuators usually employed in STMs.

Possible mechanisms of atom transfer in scanning tunneling microscopy *by R. Gomer, p. 428.* Various mechanisms for the sudden transfer of an atom from or to the tip of a scanning tunneling microscope are considered. It is concluded that thermal desorption could be responsible and also that quasi-contact in which the adsorbed atom is in effect "touching" both surfaces, which would still be separated from each other by 2–4 Å, can lead to unactivated transfer via tunneling. For barrier widths as small as 0.5 Å, however, tunneling becomes negligible.

Volume 30, Number 5, 1986

Wide-range, low-operating-voltage, bimorph STM: Application as potentiometer *by P. Muralt, D. W. Pohl, and W. Denk, p. 443.* An STM is described which operates at voltages ≤120 V. Its 3D scanner offers a wide range of displacements, has low drift and hysteresis, and exhibits good resolution. These features make it interesting for technical applications. In addition, its use as a potentiometer is described.

Traversal time for tunneling *by M. Büttiker and R. Landauer, p. 451.* Tunneling of carriers through a barrier is characterized not only by a transmission and reflection probability, but also by the time it takes a carrier to traverse the barrier. Recent work which discusses the traversal time is summarized, and its relevance is highlighted by discussing several tunneling phenomena.

Properties of vacuum tunneling currents: Anomalous barrier heights *by J. H. Coombs and J. B. Pethica, p. 455.* A design of STM which does not use vacuum internal vibration isolation and may be cooled to liquid helium temperatures is described. Tunneling current characteristics underlying STM operations are discussed. A model is presented to explain the anomalously low (<1 eV) tunneling barrier heights often observed. On the basis of this model we suggest criteria for obtaining reliable STM images.

Mono-atomic tips for scanning tunneling microscopy *by H.-W. Fink, p. 460.* By field-ion microscopy techniques, we have been able to create stable tips whose very ends are made up of just one individual atom, deposited from the gas phase onto an upper terrace of a pyramidal (111)-oriented tungsten tip. The first three layers of the tip consist of one, three, and seven atoms, respectively. Consequences of these observations for the understanding of energy transfer between gas-phase atoms and a solid surface are also discussed.

Scanning tunneling microscopy of cleaved semiconductor surfaces *by R. M. Feenstra and A. P. Fein, p. 466.* Scanning tunneling microscopy is used to study the surface topography of cleaved GaAs(110) and Si(111) surfaces. For GaAs we observe 1×1 periodicity, with an [001] corrugation amplitude of typically 0.2 Å and a [110] corrugation amplitude of ~ 0.05 Å. Surface point defects are observed, consisting typically of ~ 0.7-Å-deep depressions extending along the rows. For Si(111), we find a periodicity of two unit cells, indicating the presence of the 2×1 reconstruction. We observe a maximum [211] corrugation amplitude of 0.5 Å and a [011] corrugation amplitude of < 0.02 Å, consistent with the π-bonded chain model for this surface. Structural disorder on the surface most commonly appears as "protrusions" along or crossing over between the chains. Orientational disorder is observed in the tilt of the chains.

Scanning tunneling microscopy of surface microstructure on rough surfaces *by J. K. Gimzewski and A. Humbert, p. 472.* An important feature of the scanning tunneling microscope (STM) is its ability to image nonperiodic or disordered surfaces with atomic or near-atomic lateral and vertical resolution. Many physical and chemical properties of surfaces and interfaces are sensitive to, and in some cases determined by, random roughness or surface disorder, although our understanding is hampered by the lack of suitable techniques for investigating atomic-scale features. We present STM results for microcrystalline Ag and nanocrystalline silicon surfaces which demonstrate the unique topographic data obtainable using the STM. For comparison, the STM studies are complemented by data obtained using conventional techniques. Furthermore, the ability of the STM to investigate the influence of growth conditions on surface morphology, such as ion bombardment, adsorption, and condensation temperature, is discussed.

Near-field optical scanning microscopy with tunnel-distance regulation *by U. Dürig, D. Pohl, and F. Rohner, p. 478.* Unprecedented optical image resolution (20 nm to 30 nm) has been obtained with a near-field optical scanner using light with a wavelength of half a micrometer. The key element is an extremely small aperture (~ 10 nm) placed at the very top of a pyramidal screen. The aperture is scanned in the immediate proximity of the surface to be investigated, using vacuum tunneling to sense the distance. The amount of light transmitted by both the aperture and the sample depends sensitively on the optical properties of the sample in the immediate vicinity of the aperture.

Chemical applications of scanning tunneling microscopy *by P. West, J. Kramar, D. V. Baxter, R. J. Cave, and J. D. Baldeschwieler, p. 484.* The development of a scanning tunneling microscope at the California Institute of Technology is well under way. Electron tunneling has been demonstrated, and preliminary surface images of gold films have been obtained. Additional instrumental development is required to achieve the atomic resolution which is required for the study of chemical processes on surfaces. A theoretical model is also being developed for the study of tunneling of electrons from the probe to surfaces with molecular species adsorbed, and with atomic and molecular species intervening between the probe and the surface. These experimental tools and theoretical models, which are being developed concurrently, will be applied to the study of chemical systems and processes

on surfaces. Some of the first molecular species for study will include benzene and pyridine on metal surfaces, and porphyrins and phthalocyanines on pyrolytic graphite. The applications that can be imagined for STM in surface chemistry are very broad, and the choice of additional systems for study will be based on the results of these initial experiments.

Surface modification with the scanning tunneling microscope *by D. W. Abraham, H. J. Mamin, E. Ganz, and J. Clarke, p. 492.* We describe the design and operation of a scanning tunneling microscope (STM) intended for studying surfaces. We are able to prepare samples with ion bombardment and heating, and to characterize them with LEED and Auger analysis *in situ* before scanning with the STM. Data acquisition and analysis are computer-controlled, with a wide variety of options for presentation. In the near future, we will be able to load-lock samples without venting the UHV chamber. Our STM has very low thermal drift, typically of the order of 1 Å/min. In addition to topographical measurements of the surface, we have obtained spatially resolved maps related to the height of the tunnel barrier. We have investigated the effects of scratching the surface with the tip, and have also succeeded in depositing material from the tip onto the surface. Surface diffusion of material is found to play an important role in both of these processes.

STM activity at the University of Basel *by M. Ringger, B. W. Corb, H.-R. Hidber, R. Schlögl, R. Wiesendanger, A. Stemmer, L. Rosenthaler, A. J. Brunner, P. C. Oelhafen, and H.-J. Güntherodt, p. 500.* We have built a scanning tunneling microscope (STM) with a design similar to that published by Binnig and Rohrer. An electromagnetic device called the "maggot" has been developed for nanometer-scale movement over a wide temperature range. So far we have studied the surface topography of a Pd(100) single crystal, of a glassy $Pd_{81}Si_{19}$ alloy, and of graphite. Finally, nanometer-scale structures have been produced with the STM, indicating its potential for high-resolution microfabrication.

Applications of a high-stability scanning tunneling microscope *by H. van Kempen and G. F. A. van de Walle, p. 509.* We have constructed a scanning tunneling microscope which is quite insensitive to vibrations and has a low thermal drift. Low thermal drift is obtained by using a compensating structure for the z-axis (perpendicular to the sample surface) of the scan unit and by utilizing symmetry in the x- and y-directions. To get a low sensitivity to vibrations, we made the scanning unit compact and rigid. A very light tip holder construction allows a high scan speed. We studied different surfaces of Ag single crystals and compared the results with a study of the same surfaces from the inside by a method based on the use of focused conduction electrons. A terrace surface model, introduced to explain the focusing results, has been confirmed for the Ag(001) surface. A description of tip preparation and tip shape is given.

Construction of a UHV scanning tunneling microscope *by S. Chiang and R. J. Wilson, p. 515.* In our laboratory, we have built an ultrahigh-vacuum scanning tunneling microscope (STM). The STM is mounted onto one flange in an ultrahigh-vacuum chamber which is connected by a transfer chamber to a surface-analysis system equipped with 500-Å-resolution SAM/SEM, XPS, LEED, and sample-heating and sample-cleaning facilities. Samples and tips can be moved throughout the combined vacuum system. We describe the design and performance of the instrument and show some preliminary data on Si(111). We also show some recent results of experiments on tip preparation.

A scanning tunneling microscope for the investigation of the growth of metal films on semiconductor surfaces *by T. Berghaus, H. Neddermeyer, and S. Tosch, p. 520.* We describe a scanning tunneling microscope which is part of an apparatus designed for the investigation of metal-semiconductor surfaces. The main parts of the tunneling unit are the piezoelectric walker ("louse") carrying the sample for the coarse approach and a piezoelectric xyz system for movement of the tip. The xyz system is self-compensating with regard to uniform thermal expansion. Our first measurements have been obtained on a Au(110) surface. The spatial resolution allows the observation of monoatomic steps. Corrugation perpendicular to the rows and troughs of the (110) surface with an amplitude around 1 Å is also visible. The noise in the z direction under optimum conditions is smaller than 0.2 Å rms.

Scanning tunneling microscope automation *by M. Aguilar, P. J. Pascual, and A. Santisteban, p. 525.* A computerized system for scanning tunneling microscope control, data acquisition, and display is presented. It is based on the IBM Personal Computer Model XT or AT. An IBM Data Acquisition and Control Adapter card is used for the PC to control the electronic equipment and to measure the voltages applied to the three STM piezoelectric elements and the tunneling current. An IBM Professional Graphics Controller card is used for real-time display of the STM "images." The software has been designed in a modular way to allow the replacement of these cards and other improvements to the equipment. A discussion of performance (scanning speed and size of images) is included. Some image analysis and display tools are included for *a posteriori* image processing.

Theory of scanning tunneling microscopy and spectroscopy: Resolution, image and field states, and thin oxide layers *by N. Garcia, p. 533.* Theoretical studies on the scanning tunneling microscope and its spectroscopic version are reviewed. This research has shown that the conductance of the tunneling electrons is strongly influenced by the classical image potential. The introduction of this potential increases the conductance, although the slope of the logarithm of the conductance versus electrode separation remains practically constant. The image force also has focusing effects on the tunneling electrons and produces a minimum in the

resolution for ~ 5 Å electrode separation. Spectroscopic levels have been calculated for the image states held by the tunneling potential. The results of this work agree very well with experimental data and indicate that the evolution of the observed tunneling spectroscopic levels with applied field is very sensitive to the image potential, and, moreover, that the whole series of image states can be obtained by extrapolation to zero applied field. Work is also presented on the theoretical aspects of tunneling spectroscopy of thin oxide layers grown on a metal substrate—NiO on Ni(100). From theory and experimental data, information can be obtained about the electronic band structure and the number of layers of the oxide. All results of the above research are in accord with the experimental data of Binnig, Rohrer, et al.

Computer automation for scanning tunneling microscopy *by P. H. Schroer and J. Becker, p. 543.* Computer automation has become a necessary part of the scanning tunneling microscope (STM). We have designed a comprehensive data acquisition and image processing IBM PC/AT workstation to complement the STM. The computer is used to control the raster scans, to acquire multichannel analog data, and to store the data for later analysis both at the workstation and on a host computer. New features include the use of PC/AT extended memory, large backup storage, and image processing at the workstation. Real-time gray-scale imaging provides quick and comprehensive information to scientists and engineers. The system is flexible enough to interface with many microscope designs and with other laboratory automation projects. A mainframe host computer system is loosely coupled with the workstation to provide off-line processing, hard copy, and mass storage.

The behavior and calibration of some piezoelectric ceramics used in the STM *by S. Vieira, p. 553.* The high resolution and displacement measurement in the scanning tunneling microscope are dependent upon the behavior of the piezoelectric ceramics used for moving the tip. In this paper certain characteristics and features of piezoelectric ceramics relevant to the desired precision are discussed. These characteristics are the relaxation aftereffects that follow the change of the applied electric field, and the temperature dependence of the piezoelectric response. The above effects have been analyzed in four commercial piezoelectric ceramics by the three-terminal capacitance measurement technique.

Volume 30, Number 6, 1986

Microtasking on IBM multiprocessors *by P. Carnevali, P. Sguazzero, and V. Zecca, p. 574.* The demand for very high computing performance has become increasingly common in many scientific and engineering environments. In addition to vector processing, parallel computing is now considered a useful way to enhance performance. However, parallel computing tends to be unpopular among users because, with presently available technology and software, it requires explicit programmer intervention to exploit architectural parallelism. This intervention can be minor in some cases, but it often requires a nonnegligible amount of program restructuring, or even a reformulation of some of the algorithms used. In addition, it makes program debugging considerably more difficult. Tools for interprocedural program analysis, able to analyze at high levels large FORTRAN programs composed of many subroutines and to perform an automatic high-level parallel decomposition, are being developed at IBM. As an additional possibility, low-level parallelism could also be exploited automatically. This would greatly simplify the program analysis needed in a compiler in order to automatically insert parallel constructs in a program. However, this could only be efficient in an environment in which the cost for scheduling and synchronizing parallel activities is extremely small. Such a "microtasking" environment has been realized, at a prototype level, for IBM multiprocessors and is described in this paper. Its applicability and potential are illustrated with the help of some simple examples. These examples show that, with microtasking, good processor utilization and useful speed-ups are achieved even for fine-grain problems. The user interface of the prototype is described, both at the FORTRAN and at the Assembler level, and the possibility of incorporating this environment into a FORTRAN compiler is discussed.

Compiling APL: The Yorktown APL Translator *by G. C. Driscoll, Jr. and D. L. Orth, p. 583.* The Yorktown APL Translator (YAT) permits functions written in APL to be compiled using an existing compiler for part of the process. It also creates tables that allow the APL2 Release 2 interpreter to call the compiled code. The code can also be called from a Fortran main routine. This paper outlines the history of APL compilation, the motivation for producing YAT, the design choices that were made, and the manner of implementation. Sample APL functions and their translations are shown, and the time required to interpret these functions is compared with the time required to execute the compiled code. Possible further work is discussed.

Program analysis and code generation in an APL/370 compiler *by W.-M. Ching, p. 594.* We have implemented an APL/370 compiler which accepts a subset of APL that includes most language features and a majority of APL primitive functions. It produces System/370 assembly code directly to be run independently of an interpreter. The compiler does not require variable declarations. Its front end, which is machine-independent, employs extensive type-shape analysis based on a type-shape calculus of the primitive functions and global dataflow analysis. Its back end does storage mapping, code generation for various primitive functions, and register management. For one-line functions, compiled code executes at 2–10 times the speed of the interpreter. On programs with many iterations, the execution time of the compiled code not only is dramatically better than that of interpretation, but is actually fairly close to that of FORTRAN. This removes the performance penalty of APL in computation-intensive applications.

An automatic overlay generator *by R. Cytron and P. G. Loewner, p. 603.* We present an algorithm for automatically generating an overlay structure for a program, with the goal of reducing the primary storage requirements of that program. Subject to the constraints of intermodule dependences, the algorithm can either find a maximal overlay structure or find an overlay structure that, where possible, restricts the program to a specified amount of primary storage. Results are presented from applying this algorithm to three substantial programs.

An execution architecture for FP *by T. Huynh, B. Hailpern, and L. W. Hoevel, p. 609.* FP is a functional programming language proposed by John Backus to liberate programming from the "von Neumann Style." We define here an architecture (FP machine model) that is intended to allow easy and efficient implementation of FP on a conventional von Neumann machine. We present a new execution architecture for FP based on Johnston's contour model and on Hoevel and Flynn's notion of DEL/DCA architectures. We call our architecture DELfp, a Directly Executed Language for FP.

Early error detection in syntax-driven parsers *by A. V. Moura, p. 617.* The early error detection capabilities of syntax-driven parsers are studied. The classes of weak precedence and simple mixed-strategy precedence parsers are chosen as the object of study. Very similar techniques could be used to obtain related results for other classes of syntax-driven parsers. We investigate whether the correct-prefix and the viable-prefix properties can be enforced within these classes: A negative result is obtained for the first class and a positive one for the second. Moreover, for the simple mixed-strategy class the relationship between early error detection and parser size is studied. Some lower bounds on the parser size are proven for simple mixed-strategy precedence parsers that have the viable-prefix property.

A computerized autoradiographic technique for the simultaneous high-resolution mapping of myocardial blood flow and metabolism *by A. L'Abbate, C. Paoli, R. Porinelli, and M. G. Trivella, p. 627.* By using both radioactive particulate and diffusible tracers, combined with computerized processing and analysis of generated autoradiographs, a new technique for the high-resolution mapping of organ blood flow and metabolism has been developed. In this paper the technique has been evaluated for the myocardium.

Mechanical design of the cartridge and transport for the IBM 3480 Magnetic Tape Subsystem *by D. J. Winarski, W. W.-C. Chow, J. G. Bullock, F. B. Froehlich, and T. G. Osterday, p. 635.* The IBM 3480 Magnetic Tape Subsystem has achieved significant improvements over its predecessor, the 3420, in speed, data density, and floor-space requirements. The 3-megabyte-per-second data rate of the 3480, which is 2.4 times that of the 3420, was accomplished through the use of chromium-dioxide tape stored in a compact, single-reel cartridge threaded to a take-up reel. A sixfold increase in data density over the 3420 allowed reduction of the size of the 3480 reel, which was essential for rapid acceleration of the tape in the 3480's reel-to-reel transport. The increased data density also demanded substantial improvements in tape guiding and tape-motion control. This paper describes how mechanical analysis and design contributed to the achievement of these advances and aided in overcoming the ensuing problems.