Volume 36, Number 1, 1992

Preface *by J. C. McGroddy, p. 2.*

Thin-film-transistor/liquid crystal display technology—An introduction *by W. E. Howard, p. 3.* Liquid crystals are simple and very efficient electro-optic transducers, or light valves. Thin-film transistors are simple electronic control devices which can be fabricated on large transparent substrates. These two technologies, when combined, allow the fabrication of electronic displays which challenge the dominance of the cathode ray tube (CRT). This paper reviews the history of this important development, presents the current status in comparison to the color CRT, and describes the remaining challenges to be overcome if the color CRT is truly to be displaced.

A gray-scale addressing technique for thin-film-transistor/liquid crystal displays *by P. M. Alt, C. G. Powell, B. L. Owens, Jr., and H. Ifill, p. 11.* An addressing technique allowing continuous-tone color images to be rendered on thin-film-transistor/liquid crystal displays having bilevel drivers is described. The technique uses multiple subfields per frame, with driver voltages changed synchronously with the field data. By using N bits of data per pixel, excitation is applied to the display one bit-plane per field for N consecutive fields. The technique is analyzed, its benefits and limitations discussed, and experimental results presented. Up to 16 gray levels have been demonstrated with good image quality.

Cell design of gray-scale thin-film-transistor-driven liquid crystal displays *by H. Takano, S. Suzuki, and H. Hatoh, p. 23.* The desirable liquid crystal (LC) panel design (here called "cell design") for a gray-scale thin-film-transistor (TFT) -driven twisted nematic (TN) liquid crystal display (LCD) is discussed in terms of display legibility and ease of fabrication. To optimize cell design for gray-scale application, some key display factors such as contrast ratio, color change, and viewing cone are evaluated for various cell geometries and cell thicknesses. The cell geometries discussed are combinations of two display modes (a normally white mode and a normally black mode, in which the optical axes of the exit polarizers are placed perpendicular and parallel to those of the entrance polarizers, respectively) and two optical eigenmodes (an extraordinary-ray mode and an ordinary-ray mode, in which the transmission axes of the entrance polarizers are parallel and perpendicular to the entrance rubbing directions, respectively). A new driving scheme of threshold-voltage bias application to the LC cell is proposed to overcome the TN LCD shortcoming of a narrow viewing cone. We have adopted a cell design for a 512-color TFT LCD: 1) a first minimum normally white (NW) mode as polarizer arrangement for ease of fabrication, 2) an extraordinary-ray mode (e-mode) as optical eigenmode with a novel driving scheme (threshold-voltage biased) for gray-scale improvement in eliminating brightness reversals, and 3) a retardation ($d\Delta n$) value of 0.47 μm for further optimization of proper gray-scale order and color change. We have called this mode "threshold-voltage-biased e-mode NW," or "biased e-mode NW."

Color filter for 10.4-in.-diagonal 4096-color thin-film-transistor liquid crystal displays *by T. Koseki, T. Fukunaga, H. Yamanaka, and T. Ueki, p. 43.* Color filters with a wide color reproduction gamut and high transmittance were developed for 10.4-in.-diagonal 4096-color thin-film-transistor liquid crystal displays using pigment-dispersed photosensitive polymers. The transmission spectrum of each color pixel was designed in conjunction with other components such as backlight and polarizers in order to meet front-of-screen quality requirements. To improve screen quality, a low-resistivity common electrode was used, eliminating the top coating. A repair technique utilizing back-exposure was also developed to improve production yield. This pigment-dispersed-type color filter has the merits of a simple process, low fabrication cost, good uniformity, high reliability, and applicability to high-resolution displays. There is a problem involving deterioration of contrast ratio caused by the depolarization effect of the color filter. We measured depolarization factors for several pigments and showed that the yellow pigment was the major contributor. This depolarization effect has been minimized.

Lateral field effect in twisted nematic cells *by A. Lien and R. A. John, p. 51.* The lateral field existing in the ON state of a liquid crystal display (LCD) can result in an undesired reverse tilt domain in each pixel, thereby leading to poor contrast ratio of the display. As pixel size becomes smaller to meet the requirements of a high-information-content display, the problem becomes worse. In this paper, we present experimental results and theoretical analysis of the lateral field effect in the twisted nematic (TN) cell, which is most commonly used in the thin-film-transistor liquid crystal displays (TFT/LCD). The effects of various cell parameters, such as pretilt angle, bus-line-to-pixel spacing, and cell gap thickness, on the pixel reverse tilt domain and on the corresponding optical performance have been studied in detail. The results are useful for TFT/LCD design and cell fabrication.

Functional testing of TFT/LCD arrays *by L. C. Jenkins, R. J. Polastre, R. R. Troutman, and R. L. Wisnieff, p. 59.* The thin-film-transistor liquid crystal display (TFT/LCD) is emerging as the leading flat-panel display in computer applications. TFT array characterization is important to the research, development, and manufacturing of TFT/LCDs. This paper describes a new Dynamic Array Tester developed for that purpose and describes some examples of its use.

Reactive ion etching technology in thin-film-transistor processing *by Y. Kuo, p. 69.* This paper discusses reactive ion etching (RIE) process issues in preparing thin-film transistors (TFTs) for liquid crystal displays (LCDs). Three areas were

examined in detail: gate metal etch, dielectric etch, and a-Si:H etch, both intrinsic and n^+ doped. Although there are different requirements for each step, the basic principles for the etching process are similar. For example, each process includes three major mechanisms: plasma-phase chemistry, particle transport phenomena, and surface reactions. All data on the etching results were interpreted according to these principles. Finally, a TFT characteristic curve based on RIE of some of the most critical process steps is presented.

Study of the V_{th} shift of the thin-film transistor by the bias temperature stress test *by Y. Fujimoto, p. 76.* Amorphous silicon thin-film transistors (a-Si:H TFTs) are now widely used as the switching device in the active-matrix addressing of liquid crystal displays. One concern is the potential instability problems associated with the threshold voltage (V_{th}) shifts to higher values after prolonged operating times. The reason for this V_{th} shift has been widely discussed, and two models accounting for it have been suggested. One model explains the shifts by the trapping of electrons in the insulator, the other model by the creation of the metastable states at the a-Si:H/SiN_x interface. Our TFT insulator has the rather complicated structure of an anodic oxide film, SiO_x, SiN_x sequentially stacked over the gate electrode, which makes it difficult to separate the contribution of each layer. To confirm the V_{th} shift mechanism and the contribution of each layer of insulator to the V_{th} shift, we have prepared samples with a series of different insulators and have measured the bias dependence of their V_{th} shifts. Our results show that the anodic oxide film makes no contribution to the V_{th} shift, and it makes little difference to the V_{th} shift whether the next insulator is SiN_x or SiO_x. The latter fact may be explained in two ways. One is that both SiN_x and SiO_x make the same contribution to the V_{th} shift. Alternatively, neither SiN_x nor SiO_x makes any contribution to the V_{th} shift, but the a-Si:H/gate insulator interface has some contribution. From these experiments, it cannot be determined which of the proposed mechanisms is consistent with the behavior of the V_{th} shift of our a-Si:H TFT.

TDI charge-coupled devices: Design and applications *by H.-S. Wong, Y. L. Yao, and E. S. Schlig, p. 83.* The design and applications of charge-coupled devices (CCDs) operated in the time-delay-and-integration (TDI) mode are reviewed. Design issues regarding the use of the TDI-CCD imager for visible imaging applications are discussed. Aspects pertaining to its parallel array, serial-to-parallel interface, serial register, modulation transfer function (MTF), discrete charge motion, motion synchronization, clocking, number of integrating stages, noise, dynamic range, sensitivity, output uniformity, device yield, pixel size, and spectral response are highlighted in the context of their effect on system performance. Its imaging characteristics are compared to those of the photodiode linear array imager, and design studies and experimental results for a family of TDI-CCD imagers for scanning documents and museum art objects are described.

Volume 36, Number 2, 1992

Preface *by N. O. Lipari, p. 138.*

Plasma-based dry etching techniques in the silicon integrated circuit technology *by G. S. Oehrlein and J. F. Rembetski, p. 140.* Plasma-based dry etching techniques play a major role in the formation of silicon-based integrated circuits. The first part of this paper reviews our understanding of the means for achieving etching directionality and selectivity in reactive etching using glow discharges. Relevant trends in magnetically enhanced rf diode systems, microwave-excited electron cyclotron resonance plasmas, process clustering, real-time process monitoring and control, and computer modeling of glow discharges are discussed in the second part of the paper.

Stress-induced dislocations in silicon integrated circuits *by P. M. Fahey, S. R. Mader, S. R. Stiffler, R. L. Mohler, J. D. Mis, and J. A. Slinkman, p. 158.* Many of the processes used in the fabrication of silicon integrated circuits lead to the development of stress in the silicon substrate. Given enough stress, the substrate will yield by generating dislocations. We examine the formation of stress-induced dislocations in integrated circuit structures. Examples are presented from bipolar and MOS-based integrated circuit structures that were created during developmental studies. The underlying causes of oxidation-induced stress and the effect on such stress of varying oxidation conditions are discussed. The knowledge thus gained is used to explain dislocation generation during the formation of a shallow-trench isolation structure. The importance of ion-implantation processes in nucleating dislocations is illustrated using structures formed by a deep-trench isolation process and a process used to form a trench capacitor in a DRAM cell. The effect of device layout geometry on dislocation generation is also examined. We show how TEM observations can be used to provide more information than solely identifying those process conditions under which dislocations are generated. By combining TEM observations with stress analysis, we show how the sources of stress responsible for dislocation movement can be identified.

Application of electron and ion beam analysis techniques to microelectronics *by T.-S. Kuan, P. E. Batson, R. M. Feenstra, A. J. Slavin, and R. M. Tromp, p. 183.* The application of electron microscopy, scanning tunneling microscopy, and medium-energy ion scattering to microelectronics is reviewed. These analysis techniques are playing an important role in advancing the technology. Their use in the study of relevant phenomena regarding surfaces, interfaces, and defects is discussed. Recent developments and applications are illustrated using results obtained at the IBM Thomas J. Watson Research Center. Potential advances in the techniques are also discussed.

Numerical modeling of advanced semiconductor devices *by W. Lee, S. E. Laux, M. V. Fischetti, G. Baccarani, A. Gnudi, J. M. C. Stork, J. A. Mandelman, E. F. Crabbé, M. R. Wordeman, and F. Odeh, p. 208.* Numerical modeling of the electrical behavior of semiconductor devices is playing an increasingly important role in their development. Examples that pertain to advanced MOSFETs and bipolar transistors are presented to illustrate the importance of taking into account three-dimensional as well as nonequilibrium and nonlocal physical phenomena to effectively characterize the electrical behavior of such devices.

Integrated processing for microelectronics science and technology *by G. W. Rubloff and D. T. Bordonaro, p. 233.* This paper is a review of integrated processing—an approach to microelectronics fabrication in which sequential processes are linked by wafer transfer through a clean, controlled environment (e.g., high vacuum or inert gas). The approach is rapidly becoming the state of the art in microelectronics research, development, and manufacturing. In microelectronics research, it provides a means for advancing mechanistic understanding and material quality through *in situ* fabrication of test structures and extensive *in situ* diagnostics. In microelectronics development and manufacturing, it promises process simplification, improved contamination control and yield, and potentially more flexible equipment utilization. With increasing emphasis on ultraclean processing, involving control of reactive impurities as well as particles, and on real-time process monitoring and control, applications of integrated processing are moving toward a common ground in which state-of-the-art research techniques can be used to address key issues in development and manufacturing, and provide in return substantive guidelines for manufacturing design and practice.

Volume 36, Number 3, 1992

Preface *by J. H. Griesmer, p. 328.*

Arthur Samuel: Pioneer in Machine Learning *by G. Wiederhold and J. McCarthy, p. 329.*

Capturing the deep meaning of texts through deduction and inference *by F. Antonacci and C. M. Calamani, p. 333.* One of the main problems in the computer analysis of natural language is understanding sentences beyond a surface level, i.e., making inferences about likely circumstances and drawing plausible conclusions. At the first level, a natural-language-understanding system can answer simple and trivial questions; in order to extend the domain of possible questions that it can answer, the system must make presuppositions and recognize implications that depend on certain events (also called actions). The IBM Rome Scientific Center has developed a prototype system that is able to make inferences about what might be true. This system has been integrated with a text-understanding system (System N), also developed at Rome.

Explaining SLDNF resolution with non-normal defaults *by M. A. Casanova, A. S. Hemerly, and R. A. de T. Guerreiro, p. 347.* This paper defines a default logic interpretation for normal programs that has the following major characteristics. First, it directly captures the true nature of SLDNF resolution as an extension of SLD resolution. Second, it is semantically convincing, but it requires neither an elaborated nonstandard interpretation nor a radical rewriting of the program clauses that would make it difficult to understand their meaning. Last, it extends known results for stratified normal programs to programs that satisfy a weaker condition.

STORK and PENGUIN: Logic programming systems using general clauses and defaults *by R. A. de T. Guerreiro, A. S. Hemerly, and M. A. Casanova, p. 361.* This paper describes two logic programming systems with the expressive power of full clausal first-order logic and with a nonmonotonic component. They provide a direct generalization of pure Prolog and can be implemented using the same technology as Prolog processors. The inference engine of both systems is based on the weak-model elimination method which, in the case of the second system, is extended to incorporate defaults.

Logic programming with typed unification and its realization on an abstract machine *by C. Beierle, p. 375.* Logic programming can benefit from a typing concept which supports many software engineering principles such as data abstraction, modularization, etc. From a computational point of view, the use of types can drastically reduce the search space. Starting from these observations, this paper gives a survey of many-sorted, order-sorted, and polymorphic approaches to type concepts in logic programming. The underlying unification procedures for ordinary term unification, order-sorted unification, and in particular for polymorphic order-sorted unification are given in the style of solving a set of equations, giving a common basis for comparing them. In addition, the realization of these unification procedures on a Warren Abstract Machine-like architecture is described. Special emphasis is placed on the abstract machine developed for PROTOS-L, a logic programming language based on polymorphic order-sorted unification.

Zephyr: Toward true compiler-based programming in Prolog *by Y. Asakawa, H. Komatsu, H. Etoh, Y. Hama, and K. Maruyama, p. 391.* Prolog is widely used in prototyping, especially in artificial intelligence, but it has yet to gain widespread acceptance in application development. We think that the problems in this area result from the programming style enforced in existing Prolog systems. Zephyr is a new Prolog system refined and enhanced to help solve such problems. It allows users to do modular programming by always using a compiler instead of an interpreter. In this paper, we describe the unique features of Zephyr which make

this possible, focusing especially on *package, metafunctions,* and *tables,* and the implementation of the system on OS/2.®

An experiment in constructing an open expert system using a knowledge substrate *by C. V. Apté, R. A. Dionne, J. H. Griesmer, M. Karnaugh, J. K. Kastner, M. M. Laker, and E. K. Mays, p. 409.* This paper discusses an experiment in the use of an object-centered knowledge representation service to provide a common conceptual model for the construction of a large knowledge-intensive decision support tool. A core knowledge substrate forms a common resource for a variety of problem-solving activities and a basis for the rapid construction of new capabilities. FAME, a substantial expert system to aid in the financial marketing of IBM mainframes, has been built and extensively tested in the field to validate our tools and techniques.

Use of natural language for knowledge acquisition: Strategies to cope with semantic and pragmatic variation *by T. Wetter and R. Nüse, p. 435.* The large amount of verbal data from common knowledge-elicitation methods suggests using the data directly for knowledge acquisition by means of sophisticated natural-language analyzers (NLAs). In this paper, we analyze the feasibility of such an approach theoretically and present a number of examples. In the theoretical part of the text we first provide a detailed analysis of the *entities* involved, i.e., the domains of expertise, the qualities of knowledge about domains, the properties of generic sentences and texts in natural languages, and the conclusions to be drawn from the limited expressiveness of formal representations. Then we discuss the *processes* of transforming knowledge into natural language and of transforming natural language into formal language. Since much can go wrong in both processes, we derive desired relations or *validity criteria* among the entities and *strategies* to meet the criteria. We believe that this broad theoretical framework can be used to analyze and compare existing attempts at directly using natural language for knowledge acquisition, and thus assess the present status of the field.

Topological reasoning about dextrous grasps *by T. N. Nguyen and H. E. Stephanou, p. 469.* The need, in robot manipulation, for higher levels of dexterity and versatility than those provided by grippers and by special-purpose end-effectors has prompted much research effort during the last decade on the design and control of multifingered hands. Most work on multifingered robot hands has dealt with low-level, numeric control, commonly based on screw theory and tools drawn from line geometry, differential geometry, kinematics, and dynamics. Current numeric, contact-based schemes, however, are limited to tip prehension (intentional grasping by the fingertips). The intriguing ease with which humans perform grasping and manipulation activities has concurrently triggered new investigations to provide robots with humanlike, prehensile capability for complex tasks in unstructured environments. These investigations have resulted in numerous AI-oriented, task-directed, distributed, symbolic schemes that have been conducted essentially independently. Efforts to link symbolic and numeric schemes have been undertaken, but the results have been rather modest. This paper deals with an intelligent, integrated symbolic-numeric scheme for dextrous manipulation, using a topological approach. In this paper, we introduce a reasoning scheme called topological reasoning that is used in conjunction with a grasp-based, topological model for uniform representations of multifingered robot hands at different levels of detail (e.g., whole hand, finger, joint), and discuss its application to dextrous manipulation (grasp selection and regrasping). We show that using topological reasoning, both hand posture and hand functionality can be derived from symbolic, high-level task requirements and object attributes, and can be transformed into numeric, low-level, joint space variables. Furthermore, the reasoning scheme is applicable not only to tip prehension, but also to palm prehension and any combination of the two.

Optical recognition of hand-printed characters of any size, position, and orientation *by S. Di Zenzo, M. Del Buono, M. Meucci, and A. Spirito, p. 487.* This paper deals with the optical recognition of text data in documents such as engineering drawings, land-use and land-register maps, and utility maps. The automatic computer acquisition of these documents is performed through the basic steps of vectorization of the line-structure and recognition of the text data interspersed in the document. The latter data are usually handwritten by professional draftsmen, and may have any size, position, and orientation. We review some of the features appropriate to this particular OCR problem, and suggest a special recognition strategy. Numerous examples are given. The results obtained with a prototype system on actual land-register maps are reported.

Volume 36, Number 4, 1992

Preface *by N. M. Donofrio, p. 533.*

Enterprise Systems Connection (ESCON) Architecture—System overview *by S. A. Calta, J. A. deVeer, E. Loizides, and R. N. Strangwayes, p. 535.* This paper serves as an introduction to a wholly new IBM data processing interconnection system called Enterprise Systems Connection (ESCON™) Architecture™. Utilizing state-of-the-art fiber optic technology, the ESCON system introduces a unique concept to computer interconnection topology, the dynamic switched point-to-point connection. A comprehensive solution to the interconnection of data processing equipment and systems, the ESCON system offers superior connectivity, bandwidth, distance, and ease of installation. The ESCON architecture is directed toward the structuring of large dispersed multisystem data processing centers with campus distributed user communities, but it is equally well suited to the needs of small processing configurations. The paper first reviews the objectives of the ESCON development and then gives a more detailed discussion of the system design

alternatives and choices which were made. Topics discussed are the fiber optic technology, the interconnection topology, the ESCON architecture, and the design of the major system elements.

IBM Enterprise Systems multimode fiber optic technology *by N. R. Aulet, D. W. Boerstler, G. DeMario, F. D. Ferraiolo, C. E. Hayward, C. D. Heath, A. L. Huffman, W. R. Kelly, G. W. Peterson, and D. J. Stigliani, Jr., p. 553.* This paper describes the first implementation of optical fiber technology for the I/O channel connections of the IBM Enterprise Systems Connection (ESCON™) Architecture™. The ESCON optical link line rate is 200 megabits per second and is capable of transmission over distances of 3 km. The link is composed of a serializer, electro-optic transmitter, duplex fiber optic cable, electro-optic receiver, and deserializer. The serializer and deserializer respectively perform the conversions from parallel to serial and serial to parallel formats. The clock which is used to retime the serial data in the deserializer is extracted from the encoded serial signal using a phase-locked loop (PLL) technique. The optical link technology selected to achieve the data processing system requirements is InGaAsP/InP 1300-nm LED, InGaAsP/InP PIN photodiode, and multimode optical fiber. A duplex fiber jumper cable is designed with a rugged, low-profile, polarized connector, with a unique protective cap which recedes as it is mated. The optical link loss budget is determined by dividing the link into two major categories: available optical power and cable plant loss. The link design ensures that the minimum available power is greater than the maximum cable plant loss. The design parameters and trade-offs of the optical link are discussed in this paper. Unique measurement techniques and tools to ensure reliable and consistent link performance are described.

The IBM Enterprise Systems Connection (ESCON) Architecture *by J. C. Elliott and M. W. Sachs, p. 577.* The IBM Enterprise Systems Connection (ESCON™) Architecture™ is the architecture for the new fiber optic serial-I/O channels for the processors in the IBM System/390® family. The architecture is based on message exchanges, which replace the byte-oriented protocols of the predecessor parallel interface architecture. Its interconnection topology employs a dynamic crosspoint switch. This paper describes the major functional components of the architecture and discusses some of the technical problems that were solved during its development.

The IBM Enterprise Systems Connection (ESCON) Director: A dynamic switch for 200Mb/s fiber optic links *by C. J. Georgiou, T. A. Larsen, P. W. Oakhill, and B. Salimi, p. 593.* This paper describes the function and hardware structure of the Enterprise Systems Connection (ESCON™) Director™, an I/O switch capable of providing dynamic, nonblocking, any-to-any connectivity for up to 60 fiber optic links operating at 200 Mb/s. Optoelectronic conversion at the switch ports allows the switching of the fiber optic links to be done electronically. The establishment of paths in the switching matrix is done by means of a hard-wired, pipelined controller at a maximum rate of five million connections/disconnections per second. Routing information is provided in the header of data frames. The switch-port function, switching matrix, and matrix controller were implemented in the IBM 1-μm CMOS "standard cell" technology. The paper discusses the system interconnection philosophy, details of the data flow, the switch hardware architecture, the design methodology, and the approach to technology implementation.

The IBM Enterprise Systems Connection (ESCON) channel—A versatile building block *by J. R. Flanagan, T. A. Gregg, and D. F. Casper, p. 617.* The IBM Enterprise Systems Connection (ESCON™) environment required the design of a single channel that could be attached to the entire line of Enterprise System/9000™ processors and deliver the performance required by the top of that line. In addition to the channel, other functions were needed, such as the ESCON channel-to-channel adapter. All of these functions were required to be implemented using the same channel hardware. This paper describes the key elements of the IBM ESCON channel design.

MVS Dynamic Reconfiguration Management *by R. Cwiakala, J. D. Haggar, and H. M. Yudenfriend, p. 633.* This paper presents an overview of the Dynamic Reconfiguration Management (DRM) function of MVS/ESA™ and its support of the IBM Enterprise System/9000™ family of machines. Dynamic Reconfiguration Management is the ability to select a new I/O configuration definition without needing to perform a power-on reset (POR) of the hardware or an initial program load (IPL) of the MVS operating system. Dynamic Reconfiguration Management allows the installation to add, delete, or modify definitions for channel paths, control units, and I/O devices, in both the software and hardware I/O configurations.

Fiber Distributed Data Interface attachment to System/390 *by J. J. Coleman, C. B. Meltzer, and J. L. Weiner, p. 647.* IBM has articulated a commitment to open systems. An important element in this direction is connectivity based upon pervasive industry standards. A fundamental standard that enables System/390® to participate in heterogeneous systems environments is the Fiber Distributed Data Interface (FDDI), defined by ANSI. The initial IBM offering in support of FDDI is attachment to System/390 machines via the 3172 Interconnect Controller. FDDI provides a high-performance alternative to lower-speed local area networks (LANs) for attachment of workstations to System/390 mainframes. A key feature of the 3172 Micro Channel® (MC) controller is its internal bus structure, derived from PS/2® technology. The 3172 FDDI adapter is capable of data rates up to 80 megabytes per second (MBps). This should be sufficient to support multiple FDDI LANs at their rated speed of 10 MBps. Also, because of its MC orientation, the 3172 FDDI adapter

is potentially extendable to other platforms derived from PS/2 and RISC System/6000® technology.

Coordination of time-of-day clocks among multiple systems *by N. R. Dhondy, R. J. Schmalz, R. M. Smith, Sr., J. Thomas, and P. Yeh, p. 655.* The IBM Enterprise Systems Architecture/390™ External-Time-Reference (ETR) architecture facilitates the synchronization of time-of-day (TOD) clocks to ensure consistent time-stamp data in an installation with multiple systems. The ETR architecture also provides a means by which the TOD clocks can be set automatically, without human intervention, to an accurate standard time source. This paper reviews the design considerations involved in providing these functions—along with "clock integrity" and continuous operation—as a consistent extension of the System/370™ TOD-clock architecture. The paper also provides a functional description of the IBM 9037 Sysplex Timer®, which is an implementation of the sending unit of the ETR network.

MVS/ESA coupled-systems considerations *by M. D. Swanson and C. P. Vignola, p. 667.* One of the most important future MVS environments will be to provide on-line transaction processing and decision support through the use of large data "warehouses" in heterogeneous networks. To provide the large capacity and high availability required for this environment, it will be necessary to use multiple, cooperating MVS systems to provide those services. The design of such coupled MVS systems must accommodate a number of factors related to both the general environment and the user's needs, including availability, growth, granularity and scale, systems management, migration and coexistence, and single-system images. In this paper, we describe the system services available with MVS/ESA™ SP Version 4 that can support the coupled-systems environment, and discuss how those services can be exploited.

Integrated Cryptographic Facility of the Enterprise Systems Architecture/390: Design considerations *by R. M. Smith, Sr. and P. C. Yeh, p. 683.* This paper reviews the considerations that shaped the design of the Enterprise Systems Architecture/390™ Integrated Cryptographic Facility. It describes design issues, alternatives, and decisions, and it provides the rationale behind some of the decisions. Issues related to performance, security, usability, and availability are covered.

Design of the IBM System/390 computer family for numerically intensive applications: An overview for engineers and scientists *by D. H. Gibson and G. S. Rao, p. 695.* The IBM System/390® (S/390) computer family provides a two-order-of-magnitude performance range for numerically intensive applications. The engineer or scientist can use the same operating system, compiler, and run-time environment commonly across the family. This paper provides an overview of primary S/390 hardware and software products of interest for numerically intensive applications, including MVS/ESA™, VM/ESA®, AIX/ESA™, and the extension of FORTRAN for very large applications and parallel applications. The primary portion of the paper is focused on details of design interest in three specific hardware products within the S/390 family, with emphasis on the Enterprise System/9000™ (ES/9000™) Model 900. Also described is a potential parallel-computing configuration using the ESCON Director™. The paper concludes with a discussion of the generic system environments within which S/390 products can support the technical user.

Design of the IBM Enterprise System/9000 high-end processor *by J. S. Liptay, p. 713.* The "high-end" water-cooled processors in the IBM Enterprise System/9000™ product family use a CPU organization and cache structure which depart significantly from previous designs. The CPU organization includes multiple execution elements which execute instructions out of sequence, and uses a new virtual register management algorithm to control them. It also contains a branch history table to remember recent branches and their target addresses so that instruction fetching and decoding can be directed more accurately. These models also use a two-level cache structure which provides a level 1 cache associated with each processor and a level 2 cache associated with central storage. The level 1 cache uses a store-through organization, and is split into two separate caches, one used for instruction fetching and the other for operand references. The level 2 cache uses a store-in method to handle stores.

A single-chip IBM System/390 floating-point processor in CMOS *by S. Dao-Trong and K. Helwig, p. 733.* A floating-point processor with the IBM System/390® architecture is implemented in one CMOS VLSI chip containing over 70,000 cells (equivalent inverters), using a transistor channel length of 0.5 μm. All floating-point instructions are hard-wired, including the binary integer multiplications. The chip is implemented in a 1-μm technology with three layers of metal. All circuits are realized in standard cells except for a floating-point register and a multiplier array macro, which are custom designed to save chip area. Instructions are performed in a five-stage pipeline with a maximum operating frequency of 37 MHz. The chip measures 12.7 mm × 12.7 mm, and dissipates 2 W. It is part of the chip set which forms the core of the IBM Enterprise System/9000™ Type 9221 entry-level models.

Simulation of IBM Enterprise System/9000 Models 820 and 900 *by D. F. Ackerman, M. H. Decker, J. J. Gosselin, K. M. Lasko, M. P. Mullen, R. E. Rosa, E. V. Valera, and B. Wile, p. 751.* The discovery and removal of logic design errors early in the development cycle is critical to timely availability of market-driven processor products. This paper describes the part played by simulation in the verification of the high-end models of the IBM Enterprise System/9000™ (ES/9000™) processor family, and how that effort advanced the state of the art of logic design simulation. The increased complexity of

the ES/9000 design over that of the IBM Enterprise System/3090™ (ES/3090™) necessitated a larger simulation effort. New tools and methods were developed. Two simulation missions were established. Element simulation addressed ES/9000 functional elements (e.g., the storage controller) individually using the Compiled Enhanced Functional Simulator (CEFS), a software tool. System simulation tested two or more functional elements together using the Engineering Verification Engine (EVE), a special-purpose hardware parallel processor, and an attached IBM 3092 Processor Controller (PCE). The results achieved by simulation are discussed, together with the methods used and the impact these results had on the overall verification of the ES/9000 Models 820 and 900.

Fault-tolerance design of the IBM Enterprise System/9000 Type 9021 processors *by C. L. Chen, N. N. Tendolkar, A. J. Sutton, M. Y. Hsiao, and D. C. Bossen, p. 765.* The 9021-type processors offer the highest performance of the IBM Enterprise System/9000™ (ES/9000™) series. They also have the highest levels of concurrent error detection, fault isolation, recovery, and availability of any IBM general-purpose processor. High availability is achieved by minimizing component failure rates through improvements of the base technology, and design techniques that permit hard and soft failure detection, recovery and isolation, and component replacement concurrent with system operation. In this paper, we discuss fault-tolerant design techniques for array, logic, and storage subsystems. We also present diagnostic strategy, fault isolation, and recovery techniques. New features such as the redundant power system and Processor Availability Facility are described. The overall recovery design is described, as well as specific implementation schemes. The design process to verify the error detection, fault isolation, and recovery is also described.

Three-loop feedback control of fault-tolerant power supplies in IBM Enterprise System/9000 processors *by K. R. Covi, p. 781.* In an Enterprise System/9000™ (ES/9000™) processor, a fault-tolerant power system composed of multiple power supplies connected in parallel provides thousands of amperes of current to low-voltage (1–2 V) logic circuit boards, monitors the voltage at each board, and immediately responds to compensate for failure of a supply. If a supply fails, the very fast closed-loop response redistributes the current uniformity among the remaining supplies and allows the normal functioning of the processor logic to continue uninterrupted. This rapid response is not obtained from a conventional two-loop (current-mode) feedback power supply because the loop bandwidth is restricted by a resonance that develops in the power distribution. A third feedback loop that is added to each supply controls this power distribution resonance and makes possible the wide loop bandwidth necessary to achieve the required power system control. Analysis is presented of a three-loop control system, and a simulation of its application to a typical ES/9000 power system is described.

System cooling design for the water-cooled IBM Enterprise System/9000 processors *by D. J. Delia, T. C. Gilgert, N. H. Graham, U. Hwang, P. W. Ing, J. C. Kan, R. G. Kemink, G. C. Maling, R. F. Martin, K. P. Moran, J. R. Reyes, R. R. Schmidt, and R. A. Steinbrecher, p. 791.* The high operating speed and corresponding high chip heat fluxes in the IBM Enterprise System/9000™ water-cooled mainframe processors are made possible by improvements in component- and system-level cooling. The heart of the closed-loop water-cooling system is a coolant distribution frame (CDF) common to all water-cooled processors. The CDF provides a controlled water temperature of 21.7°C to the central electronic complex (CEC) at water flow rates up to 245 liters per minute (lpm) and rejects heat loads of up to 63 kW for the largest processor. The water flow provides cooling to multichip thermal conduction modules (TCMs), to power supplies, and to air-to-water heat exchangers that provide preconditioned air to channel and memory cards. As many as 121 chips are mounted on a TCM glass-ceramic substrate, with chip powers reaching 27 W or a heat flux of 64 W/cm^2. A separate cold plate was developed to cool these modules. The power supplies with high heat densities are primarily cooled by water which flows through a unique separable cold plate designed for ease of serviceability of the power supply. Although water cooling is utilized for components with high heat fluxes, air cooling is employed for elements of the system with lower power densities. For cards cooled by forced air, careful trade-off studies among acoustical power, chip reliability, and high availability were required. The acoustic noise emissions of all the fans and blowers were determined, and a system model was constructed to measure the noise radiated from each frame in the system. The data were used to design top covers and other components to ensure that the system could meet its thermal/acoustical requirements. A closed-loop frame in which all the heat was rejected to water was developed to meet these requirements.

Dual-tapered-piston (DTP) module cooling for IBM Enterprise System/9000 systems *by G. F. Goth, M. L. Zumbrunnen, and K. P. Moran, p. 805.* The water-cooled thermal conduction modules (TCMs) in the IBM Enterprise System/9000™ (ES/9000™) systems require a fourfold thermal improvement over TCMs in the 3090™ system. An examination of the thermal/mechanical tolerance relationships among the chips, substrate, and cooling hardware showed that a cylindrical piston would not meet this requirement. The piston was redesigned with a cylindrical center section and a taper on each end. This shape minimizes the gap between the piston and "hat" while retaining intimate contact between the piston face and chip surface during all assembly conditions. Numerical and analytical models demonstrate that this new piston shape, coupled with improved conductivity of the cooling hardware materials, exceeds ES/9000 system needs. These models were verified by tests conducted on single-site and full-scale modules in the laboratory and by tests on actual ES/9000 systems.

Volume 36, Number 5, 1992

Preface *by M. J. Attardo, p. 819.*

Advancing the state of the art in high-performance logic and array technology *by K. H. Brown, D. A. Grose, R. C. Lange, T. H. Ning, and P. A. Totta, p. 821.* High-speed silicon bipolar technology continues to meet the demands of integrated circuits for mainframe computers. IBM has developed an advanced bipolar logic and high-speed array technology for its Enterprise System/9000™ systems. This technology, code-named ATX-4, is composed of trench-isolated, double-polysilicon self-aligned bipolar devices, and has four fully planarized wiring levels with interlevel connecting studs. Chip fabrication has been implemented in 1-μm ground rules and is in full-scale manufacturing. ATX-4 represents a significant advance in providing higher-speed and lower-power logic at increased levels of integration compared with that of the ATX-1 technology used in previous generations. An overview of the design and integration of ATX-4 is discussed.

Improved performance of IBM Enterprise System/9000 bipolar logic chips *by A. E. Barish, J. P. Eckhardt, M. D. Mayo, W. A. Svarczkopf, and S. P. Gaur, p. 829.* The performance required for logic gate arrays by the IBM Enterprise System/9000™ (ES/9000™) family of water-cooled processors was obtained by redesigning chips that previously consisted of emitter-coupled logic (ECL) circuits. Multiple bipolar logic circuit families were implemented for the first time on a single IBM chip by using a modular cell approach. In 60% of the ECL circuits, ac coupling in ECL gates reduced the maximum operating power per ECL circuit on ES/9000 chips by 50% and decreased the signal delay per loaded gate by 30%, to 150 ps. About 10–20% of the remaining ECL circuits were replaced by differential current switches (DCS) which dissipated less power and improved the overall chip performance. Circuits to communicate between ECL and DCS circuit families and to improve DCS circuit reliability were included on the ES/9000 chips without affecting logic function density.

A statistical approach to quality control of non-normal lithographical overlay distributions *by R. M. Booth, Jr., K. A. Tallman, T. J. Wiltshire, and P. L. Yee, p. 835.* To achieve the high reliability and performance required by integrated circuit (IC) chips in IBM Enterprise System/9000™ processors, lithography tool centerline overlay variations between masking levels were specified at ± 0.3 μm, and circuit design images were transferred with 5× step-and-repeat photolithography tools. In contrast to data obtained from 1× lithography tools, the level-to-level overlay data which characterize deviations from circuit design rules did not fit a normal distribution, and quality control was not achieved with traditional statistical procedures. A methodology was empirically developed which transformed measured data into worst-case overlay points and approximated the data by a gamma distribution. More than 80% of the worst-case distributions were fit by the gamma distribution. The transformation of chip worst-case overlay data and the quality control testing applicable to 5× step-and-repeat lithography tool processes are described in this paper.

A four-level VLSI bipolar metallization design with chemical–mechanical planarization *by W. L. Guthrie, W. J. Patrick, E. Levine, H. C. Jones, E. A. Mehter, T. F. Houghton, G. T. Chiu, and M. A. Fury, p. 845.* A high-performance four-level semiconductor device wiring fabrication process has been developed for bipolar devices in Enterprise System/9000™ (ES/9000™) processors. The reliable interconnection of large numbers of devices on a single integrated circuit chip has been enhanced by planarizing insulators and metals using chemical–mechanical polishing processes, by a novel contact stud structure, and by a Ti-clad Al–Cu metallurgy. This paper describes the structure of the four-level wiring and elements of the process, including the silicon contacts, techniques for depositing metal and oxide to cover features with high aspect ratios, high-temperature fine-line lift-off stencils, and high-density, area array solder terminals.

Directory and Trace memory chip with active discharge cell *by P. Bunce, W. Chin, L. Clark, and B. Krumm, p. 859.* The Directory and Trace memory chip is a 7.1Kb static random-access memory with 28-bit field simultaneous compare function and independent read and write 28-bit field addressing. The array is organized as four 64 by 28 subarrays. It incorporates a unique Schottky barrier diode (SBD)-coupled cell with active discharge. As memory cells are reduced in size with each new generation, soft errors become a major concern. One method of providing high soft-error immunity is to operate the memory cell transistors in the saturation region. However, in order to write data into such memory cells, the saturation capacitance in the memory cell transistors must be discharged. In prior art, such a capacitive transistor-saturation discharge was accompanied by increased power consumption and/or delay. Typically, the new data signals themselves are used to overcome this saturation capacitance. In this design a unique SBD-coupled active discharge cell discharges the conducting transistor saturation capacitance before writing new data into the cell. Thus, it enhances the write performance and preserves the high soft-error immunity of the cell.

IBM Enterprise System/9000 clock system: A technology and system perspective *by K. Chiu, J. J. DeFazio, and T. G. McNamara, p. 867.* The minimization of system clock skew is critical to the overall performance of high-speed computer systems. This paper discusses the statistical clock skew calculation methodology employed in the analysis of the IBM Enterprise System/9000™ (ES/9000™) computer systems. Comparisons made to a worst-case design approach show the advantages of a statistical clock skew calculation and its use in the timing analysis of ES/9000 systems. Design techniques

that aid in the minimization of system clock skew are discussed throughout this paper. While many details concerning these design techniques and the statistical clock skew calculation methodology are tutorial in nature and have been used in the design of past IBM high-end machines, it is hoped that this paper will give the reader a useful understanding of the major considerations affecting the clock system and its application to an ES/9000 system.

Physical and electrical design features of the IBM Enterprise System/9000 circuit module *by E. E. Davidson, P. W. Hardin, G. A. Katopis, M. G. Nealon, and L. L. Wu, p. 877.* The requirements of the new TCM (thermal conduction module) for the IBM Enterprise System/9000™ (ES/9000™) module generated a significant number of challenges for the physical and electrical designer. For example, the need to support more circuits meant that more signal and power conductors had to be provided. In addition, the requirement for faster performance called for materials with lower dielectric constants and the use of on-module decoupling capacitors. This paper describes these changes, the design considerations that were applied to signal transmission, and the approaches that were used to contain delta-*I* and crosstalk noise in the module. Finally, the test measurements used to qualify the module are explained. The result is a TCM that more than doubles the circuit density of the TCM used for the IBM 3090™ machines, with substantially greater speed and reliability.

High-performance glass-ceramic/copper multilayer substrate with thin-film redistribution *by R. R. Tummala, J. U. Knickerbocker, S. H. Knickerbocker, L. W. Herron, R. W. Nufer, R. N. Master, M. O. Neisser, B. M. Kellner, C. H. Perry, J. N. Humenik, and T. F. Redmond, p. 889.* IBM has pioneered the use of large-area alumina multilayer ceramic substrates using state-of-the-art greensheet, molybdenum paste, and tooling technologies for its mainframe computers since 1980. During this time, a new generation of substrate materials have been developed based on copper metallization and a unique glass that crystallizes to cordierite ($2Al_2O_3 \cdot 2MgO \cdot 5SiO_2$), which has a very low dielectric constant (5.0 compared to 9.4 in previous IBM systems). The glass-ceramic/copper system provides a factor of 3 improvement in electrical conductivity over alumina/molybdenum in previous IBM systems, and the number of metallized substrate layers has been increased from 45 to 63. The thermal expansion of the new substrate ($30 \times 10^{-7}°C^{-1}$) is matched with that of the silicon chips, thereby enhancing the reliability of the 78,500 solder-bonded chip-to-substrate connections in the System/390®–Enterprise System/9000™ computers. Each substrate is 127.5 mm square and can support up to 121 complex logic and memory chips. The IBM advanced multichip module dissipates more than twice the heat flux of previous alumina/multichip modules—to 17 W/cm² at the package level and 50 W/cm² at the chip level.

Equipment-related advances in the fabrication of glass-ceramic/copper/polyimide substrates *by J. R. Kranik, J. J. Fulton, L.-S. Su, and S. M. Zimmerman, p. 905.* This paper pertains to the equipment used to produce the multilayer glass-ceramic/copper/polyimide substrates of the thermal conduction modules (TCMs) used in the IBM Enterprise System/9000™ water-cooled processors. Discussed are a flexible equipment concept applied to the punching, inspection, and testing of the glass-ceramic/copper portion of the substrates, and laser-based equipment for ablation and repair of their polyimide/copper thin-film portion.

Electrical connections to the thermal conduction modules of the IBM Enterprise System/9000 water-cooled processors *by P. J. Brofman, S. K. Ray, and K. F. Beckham, p. 921.* In a complex multichip carrier such as the thermal conduction module (TCM) of IBM high-performance mainframe processors, the interfaces between chips and their substrate as well as between the substrate and its printed circuit board must support a large number of electrical connections. Since chip, substrate, and board typically comprise very different materials, the electrical connections between them must be able to accommodate considerable thermally induced mechanical stress during assembly and use. This paper describes the pin attachment, chip attachment, wire bonding, and laser deletion processes used for forming the electrical connections to the glass-ceramic/copper/polyimide /copper substrate of the thermal conduction modules of the IBM Enterprise System/9000™ water-cooled processors.

Low-inductance decoupling capacitor for the thermal conduction modules of the IBM Enterprise System/9000 processors *by J. N. Humenik, J. M. Oberschmidt, L. L. Wu, and S. G. Paull, p. 935.* Multilayer ceramic decoupling capacitors fabricated using a barium titanate-based dielectric are used with the glass-ceramic/copper/polyimide substrates of the thermal conduction modules (TCMs) of the IBM Enterprise System/9000™ processors to suppress the voltage noise generated by the logic circuits of the semiconductor chips used in the processors. Use is made of thick-film multilayer ceramic fabrication processes and thin-film termination processes to achieve substrate-mounted capabilities which, when combined with the low-inductance design of the capacitors, minimize the inductance of the decoupling paths to their adjacent chips. When mounted on the glass-ceramic/copper/polyimide substrate of a water-cooled TCM, the capacitors suppress approximately 50% of the voltage noise, thereby enhancing the performance of the TCMs.

Aspects of the electrical design and analyses of the printed circuit boards of the IBM Enterprise System/9000 water-cooled processors *by L. E. Boone, M. R. Brinthaupt, J. A. Malack, and J. Pavlik, p. 943.* Aspects of the electrical design and analyses of the multilayer printed circuit boards of the IBM Enterprise System/9000™ water-cooled processors are discussed. The design and analyses pertained to achieving

an increase in wirability and a decrease in voltage drops, power loss, simultaneous switching noise, variation in characteristic impedance, and reflections due to the presence of stubs.

Volume 36, Number 6, 1992

The controlled experiment in knowledge-acquisition research *by C. N. Nicholson, p. 958.* This paper is based on a review of the literature about controlled experiments in research on knowledge acquisition. The review was carried out to help the author make decisions about the design of his own experiment comparing two knowledge-acquisition methods. The paper looks critically at six experiments reported in the literature, and proposes a framework within which such empirical work can be viewed. It concludes that some of the apparent difficulties can be resolved, and that controlled experiments can be a useful way of discovering the relationships at work in a knowledge-acquisition project.

Knowledge in operation *by T. Bollinger and U. Pletat, p. 965.* The L_{LILOG} knowledge representation language and an inference engine to interpret it have been developed as part of the LILOG project, where new concepts for understanding natural-language texts were investigated. L_{LILOG} is a typed predicate logic whose type system has adopted the concepts of KL-One-like languages. Further language constructs allow the formulation of default and control knowledge. The inference engine for L_{LILOG} was designed as an experimental theorem prover, allowing us to investigate the behavior of various inference calculi as well as a number of search strategies. Processing with L_{LILOG} is not restricted to a propositional reasoner for logical formulas; we are also able to delegate special kinds of inferences to external deductive components. Currently, one such external reasoner for processing spatial information on the basis of analog representation is attached to the inference engine.

Unique design concepts in GF11 and their impact on performance *by M. Kumar, p. 990.* GF11 is a 512-way SIMD parallel computer currently used to verify quantum chromodynamics theory and to explore the SIMD approach to parallel processing. System design choices, such as network design, processing element design, and other architectural features, allow GF11 to sustain very high performance, close to the 10-gigaFLOPS peak. Several applications, such as structural analysis, seismic modeling, computational fluid dynamics, and linear algebra, have been ported to GF11. Applications execute in the range of 4 to 10 gigaFLOPS. The diversity in applications that perform well on GF11 demonstrates that the SIMD architecture is effective for a much larger set of applications than previously believed. The high network and data-memory bandwidths minimize the effort required to tune applications for optimum performance.

Architecture, design, and performance of Application System/400 (AS/400) multiprocessors *by J. E. Bahr, S. B. Levenstein, L. A. McMahon, T. J. Mullins, and A. H. Wottreng, p. 1001.* The architecture, design, and performance of multiprocessors in the Application System/400® (AS/400®) family are discussed. The paper describes how this multitasking system, originally designed as a uniprocessor system, was modified to form a multiprocessor system. The unique approach, using *relatively atomic instructions*, required a minimum of change while providing significant performance gains.

Volume 37, Number 1, 1993

Architecture and performance of the ESPER-2 hard-disk drive servo writer *by H. Ono, p. 3.* The servo writer is known as the most accurate and sensitive tool on an HDD (hard-disk drive) manufacturing line. ESPER-2 is designed as a product-independent common servo writer which incorporates major advances in equipment cost, reduction of clean-room requirements, and writing accuracy. The servo writing process consists of two stages. The first stage writes the master servo pattern on as many as ten disks at once with little need for contaminant protection facilities. One such master disk is then assembled with other, raw, disks in a file, and the second stage writes the proper servo pattern on all disk surfaces by referring to the master. The effects of off-center discursions in the mechanical disk assembly and random runout (NRRO, or nonrepeatable runout) in the product spindle components are eliminated by table-lookup servo actuator control, so that writing accuracy is improved. Neither an exceptionally clean environment nor fine mechanical parts are required for the second-stage operation.

Proof of correctness of high-performance 3–1 interlock collapsing ALUs *by J. E. Phillips and S. Vassiliadis, p. 12.* A 32-bit 3–1 interlock collapsing ALU, proposed to allow the execution of two interlocked ALU-type instructions in one machine cycle using an instruction-level parallel machine implementation, is shown to produce results equivalent to a serial execution of the instructions using a 2–1 ALU. The equivalence is shown by deriving tables which represent all possible requirements for the serial execution of the instructions followed by the generalization of the table to represent sets of instructions rather than the individual instructions themselves. Consequently, the equivalence of the 3–1 interlock collapsing ALU operations with these generalized requirements of the serial execution of the instructions is shown. The correctness of a proposed high-speed interlock collapsing ALU is thereby demonstrated.

Electrical characterization and performance limits of a flexible cable *by A. Deutsch, D. W. Dranchak, G. Arjavalingam, C. W. Surovic, J. K. Tam, G. V. Kopcsay, and J. B. Gillett, p. 22.* The electrical performance of a flexible-cable test structure is characterized from low frequencies up to 25 GHz. The experimental results are used to develop and refine models which describe the performance of such cables, with particular emphasis on the contribution of dielectric and resistive losses, including skin effect. The capability of triplate flexible cables to provide high-bandwidth connections over long lengths is investigated with the models developed. A triplate design is chosen because it offers high density, limited crosstalk, no loss through radiation, and relatively inexpensive fabrication. Bit rates of 100 Mb/s–1 Gb/s are considered for propagation over lengths up to 250 cm. The paper highlights the performance-limiting factors through realistic examples, including the contribution of interfaces to other interconnection structures.

Volume 37, Number 2, 1993

Preface *by J. Horkans, p. 83.*

***In situ* surface pH measurement during electrolysis using a rotating pH electrode** *by H. Deligianni and L. T. Romankiw, p. 85.* An *in situ* technique has been developed for measuring the surface pH adjacent to a solid electrode/liquid interface during electrolysis. Measurements of the surface pH can be used to obtain insights regarding the electrodeposition of various transition metals and to obtain a better understanding of associated *in situ* surface chemistry effects. Many transition metals and alloys deposit with simultaneous hydrogen evolution and, as a result, are accompanied by a pH rise near the cathode, thereby affecting the reactivity of the nearby metal-ion species. Measurements of the surface pH of a solution containing simple salts during hydrogen evolution from a cathode were performed. The surface pH of a cathode during Ni and NiFe electrodeposition was also measured. The experiments demonstrated that, in the absence of buffers or metal ions, the surface pH rises many pH units above the bulk value. During Ni and NiFe electrodeposition, however, the surface pH of solutions consisting of simple salts and starting from a bulk pH level of 2 does not increase more than 3 pH units from the bulk value. In the case of Ni and NiFe electrodeposition, surface buffering occurs because of the hydrolysis of the metal-ion species present. Additionally, it is found that during the anomalous codeposition of NiFe, the surface pH is much lower than that required by the Dahms–Croll hypothesis.

A rotating ring-disk stripping technique used to study electroplating of Sn–Pb from methane sulfonic acid solutions *by J. Horkans, I.-C. Hsu Chang, and P. C. Andricacos, p. 97.* A rotating ring-disk stripping technique has been used to analyze Sn–Pb alloys plated from methane sulfonic acid solutions with and without a proprietary additive and to construct associated current–potential curves. The deposition of both pure Sn and pure Pb was polarized by the additive, but the polarization was much greater for Pb. For alloys plated without the additive, the potential dependence of the partial currents i_{Sn} and i_{Pb} was essentially the same as that of the pure metals. The alloy compositions were very different from the solution ratios Sn(II):Pb(II) and could be either tin-rich or lead-rich compared to the solution. In the presence of the additive, on the other hand, the alloy compositions approximated the solution compositions of the metal ions; both Pb and Sn deposition were polarized in the alloy compared to deposition of the pure metals, but the extent of polarization caused by the codepositing metal was much greater for Sn. The electrodissolution of Sn–Pb alloys in HCl shows a complex oscillatory behavior, which is produced by the selective dissolution of Sn but which may also be sustained by the formation and redissolution of sparingly soluble surface films. The oscillatory behavior disappears at low dissolution current and low rotation rate, which favor a

higher surface concentration of the dissolving metals. Composition determinations are essentially the same under conditions with and without oscillations.

Mechanistic insights into metal-mediated electroless copper plating employing hypophosphite as a reducing agent *by J. G. Gaudiello and G. L. Ballard, p. 107.* Electroless copper plating using systems containing a small amount of Ni^{2+} or Pd^{2+} as a mediator and hypophosphite as a reducing agent was investigated using several electrochemical techniques. Isothermal and component-dependent polarization, rate, E_{mix}, split-cell, and ac impedance data suggested that the systems obey mixed potential theory and function as follows: a) the mediator is initially deposited from solution to the surface of the workpiece via hypophosphite reduction, b) oxidation of the hypophosphite at mediator sites supplies charge for Cu reduction, and c) Cu plating occurs over the entire workpiece. XPS analysis and depth profiling of the resulting deposits suggested that they are homogeneous in nature and that the mediator is uniformly distributed throughout. The chemical composition determined by the XPS analysis agrees well with the results obtained by atomic emission spectroscopy. In addition, the analysis showed no evidence of Cu oxide formation.

Electroless plating of copper at a low pH level *by R. Jagannathan and M. Krishnan, p. 117.* A new process for electroless copper plating at a pH level of ≤ 9 is described. The process uses amine borane reducing agents and ligands based on neutral tetradentate nitrogen donors. The use of a variety of buffer systems is demonstrated. Electroless bath performance over a wide range of conditions is presented. The quality of the plated copper is comparable to that obtained by currently used electroless plating processes, and has a resistivity of about 1.8–2 $\mu\Omega$-cm, depending on bath composition and process parameters. Use of the process is illustrated for forming conductors and filling via holes having submicron minimum dimensions.

Feature-scale simulation of resist-patterned electrodeposition *by J. O. Dukovic, p. 125.* A numerical simulation of resist-patterned or "through-mask" electroplating has been performed to investigate shape evolution at the scale of small lithographic features. Shape evolution and step coverage have a significant influence on the shapes of such microelectronic structures as conductor lines, vias, and magnetic pole pieces. The simulation and associated analysis are based on a model for the rate distribution of the electrodeposition reaction that includes the depletion of the depositing metal ions and the inhibiting action of leveling agents. A stagnant boundary layer is assumed to be present, and the diffusion theory of leveling with a one-parameter description of kinetic inhibition is employed. The results show that when the geometry of a feature cavity makes possible the occurrence of concentration-field effects (such as radial diffusion), an uneven metal-ion flux should cause nonuniform growth at high fractions of the limiting current, and leveling agents should exert an opposing effect, even causing a strong reverse nonuniformity in some cases.

Mass transfer of an impinging jet confined between parallel plates *by O. A. Moreno, R. H. Katyl, J. D. Jones, and P. A. Moschak, p. 143.* An understanding of the mass transfer behavior of an impinging jet can be usefully applied to wet chemical processes such as water rinsing, photoresist development, and metal etching or plating. Theoretical and experimental methods were used to study the mass transfer characteristics of an axisymmetric impinging jet confined between two parallel plates. Such a configuration was used because of its potential applicability to the fabrication of printed wiring boards. The CFD (computational fluid dynamics) method was used to model fluid flow and mass transfer. An electrochemical probe based on the ferro-ferricyanide system was used to experimentally determine the mass transfer coefficients and to evaluate the applicability of the theoretical methods used. An etching method was used to characterize the mass transfer rates in a typical cupric chloride etching solution. A new observation of the effect of jet instability on the etching rate in the central impingement zone is discussed.

The quartz resonator: Electrochemical applications *by K. K. Kanazawa and O. R. Melroy, p. 157.* Since the discovery that the oscillations of resonating quartz crystals can be sustained in a liquid environment, such crystals have quickly found use as a sensitive microbalance in electrochemistry, making possible *in situ* measurements of mass changes at the electrochemical interface. The early contributions of the IBM Almaden Research Center to this exciting field of development are sketched. The principles of operation are detailed, with emphasis on an intuitive description to permit considerations of new applications. Mass density changes of the order of 10 nanograms per square centimeter (ng/cm^2) are routinely detectable as changes in the resonant frequency of about a hertz. The mass density of a monolayer of material ranges from a few tens of ng/cm^2 for polymeric materials to a few hundreds of ng/cm^2 for metals. Detailed analysis of the electrical behavior of the resonator in liquid media shows that the resonant frequency, the quality factor of the resonance, and the admittance at resonance are all sensitive to the viscoelastic properties of the contacting liquid, having implications in the study of the behavior of non-Newtonian fluids, including polymeric films.

Corrosion and protection of thin-line conductors in VLSI structures *by V. Brusic, G. S. Frankel, C.-K. Hu, M. M. Plechaty, and G. C. Schwartz, p. 173.* Thin metallic lines in VLSI circuit structures are usually encapsulated in a dielectric in order to protect them from the atmosphere and prevent corrosion. However, during processing the lines are unprotected. Some of the steps to which they are subjected during processing are quite aggressive and can result in a significant yield loss. This paper pertains to the loss which is due to corrosion during processing. It focuses on the

corrosion behavior of the two of the most commonly used conductors, aluminum and copper. Aluminum alloyed with small amounts of copper is also considered. The corrosion-related behaviors of aluminum and copper are vastly different, as is shown by their reaction with water and several processing solutions. The challenge of minimizing corrosion during processing as well as during subsequent storage and use is discussed, using suitable examples drawn from studies of thin films of the metals exposed to chemical etching, reactive ion etching, and cleaning.

Application of X-ray spectroscopy to the study of electrochemically formed surface oxide films *by A. G. Schrott and G. S. Frankel, p. 191.* Many analytical techniques can provide information regarding the chemical state, structure, and properties of materials. This paper focuses on two; X-ray photoelectron spectroscopy (XPS) and X-ray absorption spectroscopy (XAS), and their application to the study of electrochemically formed oxide films. A brief review of the phenomena underlying these techniques is provided, along with a description of the commonly used means for implementing them. Their capabilities and limitations are discussed, with an emphasis on the study of passive film composition and oxidation state. A summary of the behavior of Cr in oxide films on Al–Cr alloys is presented as an example. The coordinated use of both XPS and XAS is shown to be useful in achieving full understanding of materials systems such as electrochemically formed oxide films.

Anodic dissolution of metals at high rates *by M. Datta, p. 207.* Electrochemical metal shaping and finishing processes involve anodic dissolution of metals at high rates. This paper presents a review of some fundamental aspects related to the understanding of such processes. Included are discussions of the phenomena of passive film breakdown that lead to the transpassive dissolution of metals, some of the available information on anodic reaction stoichiometry, and the role of convective mass transport and salt precipitation layers on metal removal rate and surface finish. The use of pulsating current permits the altering of anodic mass transport rates and transpassive dissolution behavior, thereby making it possible to obtain high dissolution efficiencies even at low average current densities.

***In situ* infrared spectroscopy of the electrode–electrolyte interface** *by H. Seki, p. 227.* In the study of electrochemical processes, it is important to have a means for characterizing the molecular and ionic species at the electrode–electrolyte interface. Various types of optical vibrational spectroscopy are being used to do this *in situ*. Of these, Fourier transform infrared reflection absorption spectroscopy (IRRAS) has seen rapid progress and is in wide use. A review of the techniques used in our laboratory and examples of recent measurements are presented. The adsorbed species discussed are CO, CN^-, SO_4^{2-}, and HSO_4^-, and the electrodes are, in most cases, polycrystalline noble metals. It is shown how the interpretation of the infrared spectra is greatly aided by comparison with *ab initio* molecular orbital computations of ions and molecules on metal clusters. Some of the difficulties in the interpretation of the infrared spectra are illustrated, and the future development of optical vibrational spectroscopy for studying electrode–electrolyte interfaces is discussed.

Conduction mechanisms in contaminant layers on printed circuit boards *by J. R. White, p. 243.* AC impedance methods have been utilized to explore surface conduction mechanisms on printed circuit boards (PCBs) containing various types of solder flux contaminants. Residues from water-soluble, rosin-based, and no-clean fluxes were analyzed and evaluated for their potential impact on reliability. Impedance data for intentionally contaminated PCBs having several circuit line geometries were obtained at different relative humidities. An equivalent circuit model is presented that fits the data obtained. It is used to evaluate and distinguish among ohmic, kinetic, and diffusion effects and to predict the environmental conditions that may be detrimental for various line geometries.

Contact charging of organic materials: Ion vs. electron transfer *by A. F. Diaz and J. Guay, p. 249.* In this paper we describe some of the recent literature on the contact charging (also known as tribocharging and contact electrification) of organic materials. Although it is a very familiar phenomenon, much remains to be understood about the mechanism of charging with organic materials. It has been proposed that the charging is due to the transfer of electrons and/or ions. In some studies, the correlation between the charging and the substituent constants for a substituted series of compounds has been used to support the electron transfer mechanism, and it has been proposed that the correlation reflects systematic changes in the energy levels of the highest and lowest occupied molecular orbitals of the derivatives. In others, the detection of the ions that are transferred during contact and the correspondence between their sign and that of the transferred charge have been used to support the ion transfer mechanism. In this paper, we discuss a selected number of papers that relate the charging behavior to electrochemistry and discuss the results reported in light of the two transfer mechanisms.

Volume 37, Number 3, 1993

Preface *by J. M. Warlaumont, p. 288.*

Early history of X-ray lithography at IBM *by E. Spiller, p. 291.* We present a reconstruction of the early work on X-ray lithography at the IBM East Fishkill facility in 1969 and 1970 and a summary of the efforts at the Thomas J. Watson Research Center in Yorktown Heights, New York, between 1973 and 1976.

X-ray lithography in IBM, 1980–1992, the development years *by A. D. Wilson, p. 299.* The IBM X-ray lithography research and development program is outlined, from a

personal perspective, covering the period from the inception of the program in 1980 through the development of IBM's own storage ring for X-ray production in 1992. The following aspects, among others, are discussed: origins of the program; acquisition of an X-ray port at Brookhaven National Laboratory; masks for X-ray lithography; development of special tooling for X-ray lithography, including a wafer stepper, a precision e-beam X-ray mask writing system, and a superconducting (dipole) electron synchrotron installed in the IBM Advanced Lithography Facility (ALF) in East Fishkill, New York. Key device programs were conducted to increase understanding of the X-ray lithography process and confirm its utility.

X-ray lithography from 500 to 30 nm: X-ray nanolithography *by H. I. Smith and M. L. Schattenburg, p. 319.* Proximity X-ray lithography (XRL), using wavelengths between 0.8 and 1.5 nm, provides a near-ideal match to the "system problem" of lithography for feature sizes from 500 to 30 nm, by virtue of "absorption without scattering" and recently developed mask technology. The effects of photoelectrons, at one time thought to be problematic, are now understood not to limit resolution. With experiments and simulations via Maxwell's equations, we show that useful resolution is not limited by diffraction until linewidths are below 50 nm. It is critically important to optimize the source spatial incoherence to eliminate the deleterious effects of high spatial frequencies. Mask architecture and patterning methods are presented which we believe are compatible with manufacturing at linewidths from 500 to 30 nm. Distortion due to mask frame flexing and absorber stress can now be eliminated. Elimination of distortion at the pattern generation stage remains the problem of greatest concern. We discuss a proposed method of spatial-phase-locked electron-beam lithography which could solve this problem. Our new interferometric alignment scheme has achieved 18-nm alignment at 3σ. We assert that projection XRL using multilayer mirrors at 13 nm can never match the present performance of proximity XRL. Applications of sub-100-nm XRL, including MOS, quantum-effect, and optoelectronic devices are discussed which illustrate the benefits of high resolution, process robustness, low distortion, low damage, and high throughput.

Modeling X-ray proximity lithography *by J. Z. Y. Guo and F. Cerrina, p. 331.* Advanced semiconductor circuits, such as DRAMs, are based on very complex fabrication processes. Because of the cost and complexity involved, it is rapidly becoming impossible to adopt a "trial-and-error" approach in the development stage of a new process. Fortunately, the advances in computer power spurred by the new semiconductor devices have made it possible to compute the response of complex systems in a reasonable time on workstations. Thus, the study of a virtual representation of the process (that is, a model) can represent a solution to the high cost of process development—of course, after verification of the model accuracy through controlled experiments. A correct physical interpretation of the process under study is necessary in order to implement a model that is both accurate and extendible. This is particularly true for new approaches, such as those involved in X-ray lithography. We have studied the process of image formation in X-ray lithography and have implemented several models to predict the intensity distribution at the wafer plane. The models can be applied to the definition of an optimal exposure system that will provide the maximum exposure latitude, and to the study of new types of X-ray masks.

The Helios 1 compact superconducting storage ring X-ray source *by M. N. Wilson, A. I. C. Smith, V. C. Kempson, M. C. Townsend, J. C. Schouten, R. J. Anderson, A. R. Jorden, V. P. Suller, and M. W. Poole, p. 351.* The basic properties of synchrotron radiation are described, the design of storage rings to produce synchrotron radiation is outlined, and the criteria for matching storage ring design to the needs of X-ray lithography are discussed. Simple scaling laws are presented showing the benefits for a storage ring of using the higher fields which superconducting magnets are able to provide. Helios 1 is a compact superconducting storage ring built by Oxford Instruments for installation at the IBM Advanced Lithography Facility (ALF). Design choices for superconducting rings are discussed, and the design and construction of Helios are described. Test results from the initial commissioning of Helios at Oxford are presented, but the main data on its performance when installed at ALF are given in another paper in this issue.

Performance of the IBM synchrotron X-ray source for lithography *by C. Archie, p. 373.* The compact superconducting synchrotron X-ray source at the IBM Advanced Lithography Facility in East Fishkill, New York has been in service to customers since the start of 1992. Its availability during scheduled time is greater than 90%, with recent months frequently surpassing 95%. Data on the long-term behavior of the X-ray source properties and subsystem performance are now available. The full system continues to meet all specifications and even to surpass them in key areas. Measured electron beam properties such as beam size, short- and long-term positional stability, and beam lifetime are presented. Lifetimes greater than 20 hours for typical stored beams have significantly simplified operations and increased availability compared to projections. This paper also describes some unique features of this X-ray source and goes beyond a discussion of downtime to describe the efforts *behind the scenes* to maintain and operate it.

Design considerations for the IBM X-ray lithography facility *by J. A. Leavey and L. G. Lesoine, p. 385.* Synchrotrons, like other large particle accelerators, have historically been the tools of universities and national laboratories for research. Moving this technology to industry presents many challenges which do not exist in an academic environment. One major challenge is to develop a facility to house and support the ring in a manufacturing-like mode

where operator, customer, and public concern for radiation and industrial safety is of extreme importance. This paper describes IBM's efforts to design and build a facility to address these safety concerns.

X-ray lithography beamlines in the IBM Advanced Lithography Facility *by J. P. Silverman, R. P. Rippstein, and J. M. Oberschmidt, p. 395.* In 1991 a storage ring designed as a source of X-rays for X-ray lithography was delivered, installed, and commissioned in the IBM Advanced Lithography Facility (ALF) in East Fishkill, New York. Beamlines of two different designs have been constructed and installed on the ring to deliver the X-rays to the exposure stations. One design is intended for use with a stepper for the fabrication of integrated circuits. The second design is for a general-purpose research and development beamline which is used for unaligned exposures as well as for characterization of beamline components. The design and performance of both are described. Special attention is given to a paraboloid mirror optical system which is used to collimate the radiation from the storage ring. Both the theoretical and the measured performance of the mirror are presented and shown to be in excellent agreement. An exposure nonuniformity of less than ±3%, including contributions from both the mirror and the beryllium exit window, has been achieved.

Electron beam lithography tool for manufacture of X-ray masks *by T. R. Groves, J. G. Hartley, H. C. Pfeiffer, D. Puisto, and D. K. Bailey, p. 411.* An electron beam lithography system suitable for manufacturing X-ray masks with critical dimensions down to 0.35 μm is described. The system features a 50-kV variable shaped spot (VSS) electron column with a variable axis immersion lens (VAIL). This column is capable of maintaining 0.035-μm edge acuity of the focused spot over a 2.1-mm deflection field. These fields are stitched together over an 84 × 84-mm active pattern area via motion of an *xy* table. The table position is measured using a laser interferometer. The measurement data are fed back to the magnetic deflection to correct small errors. Maintaining positional accuracy of the beam relative to the writing surface relies on a strategy of measuring and correcting repeatable errors. This is described in detail. Pattern placement accuracy is 0.070 μm (3σ) and image size control is 0.025 μm (3σ), achieved over the entire 84 × 84-mm pattern area. This performance is achieved with yield better than 90%, as confirmed by routine measurements. The system is currently used to manufacture product X-ray masks with 0.35-μm critical dimensions. Typical measurement results on product masks are presented.

X-ray mask repair *by P. G. Blauner and J. Mauer, p. 421.* A method for repairing X-ray lithographic masks using focused ion beam technology is described and demonstrated. The ion beam is used for mask imaging, for absorber milling for opaque repair, and for deposition of X-ray-opaque material for clear repair. Solutions to the unique problems faced in executing these tasks on the high-resolution, high-aspect-ratio patterns characteristic of X-ray masks are discussed. Several effects of material redeposition during opaque repair are explored. The significance of this same redeposition during clear repair and the resulting advantage gained in using a high-yield deposition process are illustrated. Examples of repairs and printed images of these repairs are shown for feature sizes smaller than 0.25 μm.

Resist materials and processes for X-ray lithography *by D. Seeger, p. 435.* A key component that is sometimes overlooked in X-ray lithography is the resist material. The lithographic properties of these materials are extremely important if one is to take advantage of the superior lithographic performance often observed in X-ray lithography. The properties of such materials may even be more important than in conventional optical lithography, since the feature sizes delineated by this lithographic technique are much smaller. A description of X-ray resists is presented which discusses both the chemistry and the lithographic properties of these materials. The characterization and stability of these processes are highlighted.

Volume 37, Number 4, 1993

Head actuator dynamics of an IBM 5 1/4-inch disk drive *by D. P. Fazzio, M. A. Moser, C. J. Polson, and J. N. Scheffel, p. 479.* The IBM 5¼-inch disk drive contained in the IBM 9345 DASD Module provides high track density and storage capacity, dynamic test failure rates below three parts per million, and low sensitivity to assembly variations. The design techniques used to achieve the required vibrational characteristics of the head actuator assembly are described. Dynamic stability specifications are derived from drive performance requirements and the actuator servomechanical system design. Modeshapes of the actuator are determined by encoding magnetic patterns onto a disk and using the read/write heads as position transducers in an operational drive. Structural changes in the carriage assembly that might lead to design improvements are explored with models derived using finite element analysis. Taguchi orthogonal matrix experiments are used to reduce the sensitivity of the actuator to dimensional tolerances and assembly processes. The achievement of actuator assembly design objectives is verified from production yields and statistical data obtained during dynamic tests.

Statistical modeling in manufacturing: Adapting a diagnostic tool to real-time applications *by B. E. Osborn, p. 491.* This paper describes a process for constructing a statistical model to automate the analysis of data from complex diagnostic tools. The method is demonstrated on data taken from an optical emission spectrometer (OES), one of the most powerful tools used in semiconductor manufacturing for detecting the chemical composition and impurity levels in plasma processes. The analysis of OES data currently requires hours of manual effort by an expert spectroscopist, rendering it ineffective for real-time

monitoring and control. However, through the use of statistical modeling, the analysis can be performed automatically on a personal computer in a matter of seconds. The process of model construction is examined in general, and methods are developed for demonstrating how information from an expert can be combined with information from the data in order to provide a statistical basis for analysis. The effectiveness of the model is demonstrated on data from typical plasma processes.

Flexible simulation of a complex semiconductor manufacturing line using a rule-based system *by B. R. Tibbitts, p. 507.* Rule-based systems have been used to produce fast, flexible simulation models for semiconductor manufacturing lines. This paper describes such a rule-based simulator for a semiconductor manufacturing line, and the language in which it is written. The simulator is written in a rule-based declarative style that uses a single-rule "template" to move thousands of product lots through various process steps; the rule is customized as needed with data for each step, route, lot, tool, manpower skill, etc. Since line or product changes require only reading new data from a database, without reprogramming, this provides a modeling environment that is simple, flexible, and maintainable. The model is implemented in ECLPS (Enhanced Common Lisp Production System), also known as a knowledge-based or expert systems language. It handles very large models (thousands of data elements, or more) well and is very fast. Subsequent changes improved the speed several orders of magnitude over that of an older version of the model, primarily through use of a preprocessor to eliminate duplicate and redundant data, and by enforcing data typing to take advantage of special techniques for very fast processing of extremely large matches (hashed indices). ECLPS also provides a built-in simulated time clock and other constructs to simplify simulation applications. The model runs daily at the IBM semiconductor manufacturing plant in Yasu, Japan, where it has been in use for many years, currently on three different semiconductor manufacturing lines.

Component procurement and allocation for products assembled to forecast: Risk-pooling effects *by S. J. Grotzinger, R. Srinivasan, R. Akella, and S. Bollapragada, p. 523.* This paper considers procurement and allocation policies in a manufacturing environment where common components are assembled into various products that have stochastic demands. The components are allocated to the assembly of a product at a time when product demand is still uncertain (assemble to forecast, ATF). The special case of one component shared by N different products is analyzed, and insights into the general problem are obtained for the situation in which the common component can be reallocated to different products as product demands change. An allocation policy is developed for general distributions and prices in an ATF environment. The policy first addresses anomalies in the state of the system and then, for a feasible state, minimizes the expected excess finished-goods inventory. A procurement level that is nearly optimal is obtained from a Monte Carlo simulation in which the probability of satisfying all of the random product demands simultaneously is considered relative to this allocation policy. Numerical studies indicate that the total component and finished-goods inventory is significantly reduced by an allocation policy that incorporates risk pooling while still fulfilling service-level requirements.

Modeling the cost of data communication for multi-node computer networks operating in the United States *by D. R. Irvin, p. 537.* The study reported here examines the cost of data communication for multi-node computer networks operating in the United States. We begin by defining a market basket of private-line transmission services and identifying its constituent prices. Two analytic models are then proposed. The first, which derives a theoretical relationship from microeconomic considerations, gives price movement as a function of the demand for service. The second embodies a learning curve fit to historical data, wherein the slope of this curve (0.71) equals the slope of the historical curve for the advance of integrated-circuit technology. Extrapolations from the two models agree well; moreover, both extrapolations conform to long-established historical trends. These agreements lend plausibility to the idea that the price of data communication unfolds in an orderly way over the long run, and, despite the perturbation introduced by the Bell System divestiture of 1984, future price movements may return to their traditional 11% annual decline.

A load-instruction unit for pipelined processors *by R. J. Eickemeyer and S. Vassiliadis, p. 547.* A special-purpose load unit is proposed as part of a processor design. The unit prefetches data from the cache by predicting the address of the data fetch in advance. This prefetch allows the cache access to take place early, in an otherwise unused cache cycle, eliminating one cycle from the load instruction. The prediction also allows the cache to prefetch data if they are not already in the cache. The cache-miss handling can be overlapped with other instruction execution. It is shown, using trace-driven simulations, that the proposed mechanism, when incorporated in a design, may contribute to a significant increase in processor performance. The paper also compares different prediction methods and describes a hardware implementation for the load unit.

Volume 37, Number 5, 1993

Preface *by C. K. Lim, T. Caulfield, and J. Benenati, p. 582.*

Finite element analysis for Solder Ball Connect (SBC) structural design optimization *by J. S. Corbin, p. 585.* Solder Ball Connect (SBC) is a second-level surface mount electronics packaging technology in which ceramic modules containing one or more chips are joined to a circuit card (FR-4) by means of an array of nonhomogeneous solder columns. These columns consist of a high-temperature-melting 90%Pb/10%Sn solder sphere

attached to the module and card with eutectic solder fillets. The solder structures accommodate the bulk of the strain (which is due to the thermal-expansion mismatch between FR-4 and the 9211 ceramic of the modules) generated during power cycling. If the solder structures are not properly designed, the thermal strain can be a source of premature fatigue failure. In this work, finite element analysis is used to characterize the plastic strains that develop in the SBC interconnection during thermal cycling. Since plastic strain is a dominant parameter that influences low-cycle fatigue, it is used as a basis of comparison for various structural alternatives. Designed experiment techniques are used to systematically evaluate the thermal strain sensitivity to structural variables. Results are used to identify an optimally reliable structure that is robust in terms of assembly-process variables.

Attachment of Solder Ball Connect (SBC) packages to circuit cards *by M. D. Ries, D. R. Banks, D. P. Watson, and K. G. Hoebener, p. 597.* IBM has developed an assembly process to attach a new family of Solder Ball Connect (SBC) integrated circuit packages to glass/epoxy cards using surface mount technology (SMT). The process provides nearly perfect yields for the resulting solder ball joint structures and ensures reliability by controlling wear-out due to metallurgical fatigue. The package, card, and process parameters found to most strongly influence the assembly yield and reliability are summarized, and unique test hardware and analysis techniques are discussed. Process considerations, analytical techniques, and test methods described for SBC packages should apply to other ball grid array (BGA) packages.

Thermal modeling of the infrared reflow process for Solder Ball Connect (SBC) *by H. V. Mahaney, p. 609.* A thermal model of the infrared reflow process has been developed for an FR-4 card populated with an array of Solder Ball Connect (SBC) modules. The analysis of the three-dimensional, transient, finite element model accounts for radiative exchange within the infrared oven and for the heat conduction (nonisotropic) within the modules and card. Transient temperature profiles of selected points and three-dimensional temperature distributions at selected times are presented to describe the primary heat-transport mechanisms. Numerical predictions and empirical data indicate that the SBC modules are relatively isothermal throughout the infrared reflow process. Therefore, every solder ball within the array exhibits a nearly identical thermal profile. This result is fortunate, since the inner solder ball connections cannot be visually inspected. The influence of module spacing and the ability to improve the reflow process by use of a high-emissivity cap coating are demonstrated.

Thermal–mechanical strain characterization for printed wiring boards *by T. Y. Wu, Y. Guo, and W. T. Chen, p. 621.* Multilayer printed wiring boards are widely used in electronic packaging assemblies. One critical reliability concern is the thermal–mechanical strains induced by temperature change. For example, the in-plane strain affects the thermal fatigue life of surface mount solder interconnections, while the out-of-plane strain affects the mechanical integrity of the plated-through holes of the printed wiring boards. For this paper, a systematic study of the thermal–mechanical strain of epoxy–glass printed wiring boards, below and above the glass transition temperatures of the epoxies, has been carried out. The study includes measurements of properties of basic constituent materials (epoxy, glass fabric, copper), of intermediate building blocks in the fabrication process, and of final products. The study has led to a quantitative engineering model that predicts the average in-plane thermal–mechanical strain for use in modeling surface mount components on a printed wiring board, as well as the average out-of-plane thermal–mechanical strain for determining plated-through-hole reliability in thermal shock processes. The model was verified by two experimental techniques (measurement by thermomechanical analyzer, and moiré interferometry) applied to two epoxy resins and three glass fabrics, with and without copper planes. For thermal shock below the glass transition temperatures of the epoxy resins, the in-plane and out-of-plane strains are described by a modified rule-of-mixtures theory and a biaxial plane stress model, respectively. For temperatures above the glass transition temperatures, the in-plane strains are governed by the copper and glass fabric, whereas the out-of-plane strains are dominated by the incompressible fluid behavior of the epoxy resins. The nonuniform pattern of thermal expansion in regions populated with plated-through holes was examined. The reliability of surface mount solder interconnections and plated-through holes is discussed.

Solder Ball Connect (SBC) assemblies under thermal loading: I. Deformation measurement via moiré interferometry, and its interpretation *by Y. Guo, C. K. Lim, W. T. Chen, and C. G. Woychik, p. 635.* Thermal deformations that result from mismatches of coefficients of thermal expansion (CTE) in Solder Ball Connect (SBC) assemblies were investigated. The CTE mismatches of the materials and the components in the package have both macro and micro effects on the strain distributions in the SBC interconnections. The geometry of the SBC joint also has a strong influence on the solder strains in the SBC package. An experimental technique with high sensitivity and resolution, moiré interferometry, was used to obtain whole-field displacements. Thermal strains in SBC packages, especially the strain concentrations in the SBC joints, were then determined from the displacement fields. The experimental results played an important role in failure analysis, structural design optimization, and finite element model verification in the IBM SBC program. The results also show that moiré interferometry is a very powerful and effective tool in experimental studies of electronic packaging.

Solder Ball Connect (SBC) assemblies under thermal loading: II. Strain analysis via image processing, and reliability considerations *by H.-C. Choi, Y. Guo, W.*

LaFontaine, and C. K. Lim, p. 649. A novel approach to processing interferometric moiré images, called computational Fourier transform moiré, has been developed. The essential principle of this technique is to automatically calculate a whole-field strain from digitized images of interferometric moiré fringes using digital Fourier transform procedures. With the use of this technique, a whole-field strain distribution of a Solder Ball Connect (SBC) interconnection under thermal loading was obtained. The calculated strain field was then used to understand fatigue modes of SBC observed from an accelerated thermal cycling (ATC) test.