Preface

Recent advances in semiconductor and interconnection technologies and microprocessor design have resulted in a choice of computing options that has redefined the role of traditional large-scale computer systems. These systems are increasingly used as servers for databases, files, communications, and centralized operations. In addition to their ability to handle large volumes of data, transactions, and computations at high speeds, the systems are distinguished by high levels of data integrity and availability, which are required to reliably process critical applications.

The characteristics of large-scale data processing, data integrity, and high availability were developed for System/370 during the past two decades by numerous improvements in architecture, hardware, and software such as multiprocessing, large-scale I/O processing, the ability to address gigabytes of real storage and terabytes of virtual storage, and special-purpose processing. Several of these are described in IBM journals: System/370 channel subsystem design (*IBM Journal of Research and Development* **27**, No. 3, 1983); 3090 system multiprocessor with the Vector Facility—the first special-purpose processor (*IBM Systems Journal* **25**, No. 1, 1986); and virtual storage addressing extensions via ESA/370 architecture (*IBM Systems Journal* **28**, No. 1, 1989).

IBM recently introduced a family of large-scale processors, known as Enterprise System/9000 (ES/9000), with many new or enhanced functions and the associated system software. The functional enhancements constitute a new generation of architecture known as Enterprise Systems Architecture/390 (or System/390 architecture) and make available to System/370 users significant new functions in an evolutionary manner. The technical advances in System/390 architecture and ES/9000 systems design described below extend the System/370 advantages of data integrity and availability, as well as I/O handling and the processing of large volumes of data.

System/390 architecture lays the foundation for IBM large-scale system structure in the coming decade. The System/390 architecture includes

- An Enterprise Systems Connection (ESCON) architecture to improve the distance, data rate, connectivity, and availability for connecting ES/9000 systems to I/O devices, local area networks (LANs), workstations, and other systems with fiber optic technology.
- A system structure and architecture for connecting system complexes (sysplexes) to achieve increases in overall capacity through the addition of systems, and to increase system availability.
- An architecture and system design for a high-performance coprocessor for cryptographic functions.
- A significantly improved platform for technical computing which includes new functions in hardware, compiler, and operating systems.
- A VM/ESA architecture for support of virtual addressing extensions in the VM/ESA operating system (*IBM Systems Journal* **30**, No. 1, 1991).

System/390 architecture is implemented in the entire family of ES/9000 processors, which span a large range in performance: This range is achieved through different processor design points, and through the use of low-power CMOS chips at the low end and high-performance bipolar semiconductor chips packaged on thermal conduction modules (TCMs) at the midrange and high end. The high-end models (Models 820, 860, and 900) achieve high performance and extend the state-of-the-art design by including 1) sophisticated branch prediction, superscalar design with multiple execution elements, out-of-sequence completion of instructions, and shared second-level cache; 2) enhanced power/thermal/cooling systems, including redundant power systems which automatically compensate for power supply variations and improve processor availability; and 3) advanced semiconductor devices, circuits, and packaging (*IBM Journal of Research and Development* **36**, No. 5, 1992).

Many of the technical advances made in System/390 architecture, the ES/9000 processors, and the power/thermal system design of the high-end, water-cooled processors are described in this topical issue of the *IBM Journal of Research and Development*. The papers focus on the advances embodied in System/390 architecture included in the ES/9000 processor family, on key technical advances in high-end processor design, and on topics unique to low-end processors. (The ES/9000 Type 9121 family of single-frame, air-cooled, midrange processors was discussed in the *IBM Journal of Research and Development* **35**, No. 3, 1991).

This issue is organized into three sections. The first seven papers deal with descriptions of advances introduced through the Enterprise Systems Connection Architecture (ESCON). The objectives of the ESCON architecture and the system design concepts are discussed by Calta et al. Aulet et al. describe the fiber optic technology used in the ESCON systems, while Elliott and Sachs discuss the functional components of the architecture. The design of the ESCON Director, which is a nonblocking I/O switch capable of dynamically connecting any input port to any output port, is described by Georgiou et al., while the hardware design issues of a fiber optic channel that was designed as a common building block for the entire family of ES/9000 processors are presented by Flanagan et al. Cwiakala et al. discuss the system structure for improving system availability through dynamic reconfiguration of the I/O subsystem, and Coleman et al. describe the

coexistence of the ESCON architecture in the open-system environment that includes fiber-distributed data interfaces (FDDI) for workstation attachment.

The next two papers discuss the system complex (sysplex) structure. Dhondy et al. describe the structure and the high-availability design of its timer, which provides a common timing source for time-of-day clocks in attached processors within a sysplex system. Swanson et al. discuss the structure and multiple virtual storage (MVS) services for coupling multiple systems into a sysplex.

In the next paper, Smith and Yeh describe the system design and design issues of the cryptographic coprocessor that is available on the high-end models of the ES/9000 processors.

Gibson and Rao describe the ES/9000 processors and their programming system environment from an engineering/scientific perspective.

The four papers in the second section describe the design of high- and low-end ES/9000 processors. Liptay discusses the design of the high-end processors. A description of the floating-point processor used in the low-end ES/9000 processors is presented by Dao-Trong and Helwig. Ackerman et al. describe the techniques used for verifying correctness of the high-end processors, including the use of a special-purpose parallel processor. Chen et al. discuss the design techniques for concurrent error detection at low cost as well as fault isolation and recovery in the high-end processors.

The power and thermal system improvements to the high-end, water-cooled ES/9000 mainframe processors are discussed in the next three papers. The design and analysis of a fault-tolerant power system that immediately compensates for a power supply failure are described by Covi. The improvements made in component- and system-level cooling of water-cooled mainframe processors are discussed by Delia et al. Goth et al. describe the enhanced heat transfer from logic chip surfaces that is due to a modification in piston design and to conductivity improvements in the hardware of the thermal conduction module (TCM).

The papers in this issue describe a number of the key advances and developments incorporated into System/390 architecture and supporting system software and into ES/9000 water-cooled processors. Many additional functions, such as the ability to share channel subsystems across logical Processor Resource/Systems Manager (PR/SM) partitions, subsystem storage isolation to address improvements in availability, and functions introduced only through system software, were introduced after these papers were written.

The numerous advances described in this issue of the *Journal* are the results of the efforts of many technical professionals in IBM whose goal was to develop a new generation of large-scale systems. Their work is presented by authors who are among the leaders in their fields. It is through the cumulative efforts of all these people that this special issue is possible.

N. M. Donofrio

General Manager,
IBM Enterprise Systems