

## Preface

On June 9, 1997, IBM announced a new generation of microprocessor-based S/390<sup>®</sup> mainframe-class servers that deliver performance comparable to that of previous systems built with older, more costly bipolar technology. With this announcement, the S/390 Parallel Enterprise Server Generation 4 (G4) marks the completion of the first stage of the transition of S/390 to CMOS (complementary metal oxide semiconductor) technology, and the beginning of a program to extend this technology to the year 2000 and beyond.

Prior to 1994, IBM utilized CMOS technology in the 9370 and 9221 processor families. These early models paved the way for the major transformation to the 9672 processor family introduced in 1994. The objective of the 9672 processor family, from a hardware point of view, was to produce a CMOS processor within four years that would be equivalent to the largest bipolar processor. This new family had to maintain support for all of the traditional attributes of high reliability and availability associated with the existing S/390 mainframe family. It also had to incorporate the new requirements of the growing client/server computing paradigm. In the first paper in this issue, Rao et al. present an overview of this transition.

In September 1996, IBM introduced a CMOS S/390 processor (G3) whose performance is approximately 75% of that of the largest S/390 bipolar processor. The G3 processor system was designed with sufficient bandwidth to support a higher-performance CMOS microprocessor. In June 1997, G4 was announced at the performance level of the largest S/390 bipolar mainframe, thus marking the successful introduction of the G4 microprocessor into the G3 structure. Further, the new CMOS servers also include support for enhanced software and I/O, targeted at open, client/server, network-centric computing.

I had the pleasure of announcing fourteen new 9672 Enterprise Server RX5 models in Japan. This announcement included the IBM S/390 Parallel Enterprise Server Generation 4, which is powered by up to ten CMOS microprocessors and delivers an average performance improvement of up to 33 percent over the previous generation (G3) of CMOS technology-based S/390 systems announced merely nine months before. Of the many important aspects of successful large-systems development, not all could not be addressed in this edition of the *IBM Journal of Research and Development*. We have covered only a subset of the many hardware advances achieved during the development of the G3 and G4 processor products.

Our thanks to the many authors from the IBM S/390 Hardware Development laboratories in Poughkeepsie, New York, and Boeblingen, Germany, the IBM Microelectronics Division laboratory in Endicott, New

York, and the IBM Thomas J. Watson Research Center in Yorktown Heights, New York, who have taken time to document this achievement. Thanks and congratulations to all members of the S/390 team worldwide, and to the other IBM teams whose efforts have contributed to the success of this transformation. We would like also to extend special thanks to the IBM Microelectronics Division for the timely delivery of technology in support of the S/390 server products.

Papers by Doettling et al. and Mak et al. describe the system design and the cache structure for the G3 and G4 servers, respectively. The G3/G4 servers were jointly defined so that the G4 processor and level-2 cache could replace the G3 processor and level-2 cache while maintaining the rest of the G3 structure (nest). This strategy was important to the reduction in overall server development schedule and expense. In the next paper, Gregg describes how G3 and G4 addressed the changing I/O environment, allowing for support of new network adapters and a new internal system link. The paper by Webb and Liptay presents a new processor architecture (using a unique hardware fault detection and recovery design) that gives the G4 processor a significant performance enhancement over previous S/390 processors. In fact, the G4 processor achieves a cycle time that is equal to that of the best CISC processors on the market today. The G4 floating-point unit is described in the paper by Schwarz et al. Many enhancements to the S/390 floating-point logic implementation, circuit design, and chip physical design are presented.

Successfully designing high-performance CMOS circuits was critical to producing a G4 CMOS processor that could replace the most powerful IBM bipolar processors. Sigal et al. describe the technology and circuit techniques used to support design points up to 400 MHz. Crucial clock circuits, PLL design, and array circuit designs are discussed.

Such complexity of custom circuit design and chip development was not required for the G3 processor and its support chips. Kick et al. describe the rationale and the methodology used to develop the support chips for G3 and G4 and the processor and level-2 cache for G3. The methodology used to develop the G4 high-performance microprocessor is then described by Shepard et al. This paper addresses in detail key concerns of chip physical design. The complexities are significant, highlighting the tremendous achievement of the G4 design. It should also be clear from the contrast between the Kick and Shepard papers that there is a need for both methodologies, given schedule and cost constraints.

The next group of papers highlight the simulation/verification and testing of the G3 and G4 designs. Certainly it is of the utmost importance to have 100% functional components to implement the design as

specified. Both G3 and G4 products were developed on tight schedules as IBM S/390 pushed to complete the transformation from bipolar to CMOS as quickly as possible. This made it imperative to achieve a high functional design prior to the first hardware implementation, as described in the papers covering the verification of the G3 and G4 products at all levels of design. Wile et al. describe the verification of the G4 processor and level-2 cache chip. Schlipf et al. present a formal verification technique used on the G3 MBA chip to enhance the verification process. The paper by Koerner and Licker covers the verification process that ensures the correct interaction of service element code and S/390 system hardware. Wile describes a simulation tool set used by the G4 designers to improve the quality of their unit designs. Van Huben discusses the use of two-cycle simulation in the G4 verification process, which helped to flush out problems that typically arise from the use of "black boxes." Next, Hallock et al. describe a common programming interface for cycle simulators that made it easy to interface with multiple different simulators. This allowed the G4 verification team to use different simulators as necessary to address the different requirements of unit versus system simulation. The last paper in this issue, by Huott et al., addresses the chip test strategy and methodology used on the G4 microprocessor and level-2 cache chip.

Through continued innovation in microprocessor design, system structure, and technological improvement, the S/390 mainframe-class server has evolved successfully into a platform that delivers the cost/performance needed by our customers. We will continue this evolution of S/390 as a balanced platform that will help support our customers' success into the next century.

Ross A. Mauri  
Vice President, Global Hardware Development  
IBM System/390 Division

*Guest Editor*

*Editor's note:* The editors wish to acknowledge the many contributions of Guest Editor Cyril Price of the IBM S/390 Hardware Development group, who organized the subject matter, recruited the authors, negotiated the commitment of peer reviewers, and served as advisor and arbitrator on the many technical and logistical questions that arose during the publication process. Without his assistance this topical issue could not have been done.