

Preface

The IBM S/390® platform revitalization which began in 1993 saw its first stage successfully completed with the introduction of the S/390 Parallel Enterprise Server™ Generation 4 (G4) CMOS processor. In September 1998, IBM announced S/390 Parallel Enterprise Server Generation 5 (G5) and thus launched the second stage of the transition from bipolar to CMOS technology.

The G5 announcement took the industry by surprise, coming in at 1069 MIPS, more than double the performance of the G4 processor. Thus, IBM had established a true growth machine for its bipolar customers. At the time of announcement, the G5 microprocessor was the fastest CISC microprocessor in the industry. The G5 server provided all of the traditional attributes of high reliability and availability associated with existing S/390 bipolar hardware, while adding significant performance growth and maintaining the advantages of CMOS: space, power, weight, and maintenance. Along with the performance growth, the G5 server offered enhancements in Parallel Sysplex®, floating-point processing, and cryptographic processing.

The great success of G5 was followed by the S/390 Parallel Enterprise Server Generation 6 (G6) announcement in May 1999, only eight months after G5. The G6 microprocessor is currently, at this writing, the fastest production processor in the industry. It uses IBM Microelectronics copper technology and extends the number of processor chips to 14, thus giving G6 a 50% increase in single-image performance and a 35% increase in uniprocessor performance.

As in previous generations of the CMOS servers, both of these new servers support enhanced software and I/O, are targeted at open, client/server, network-centric computing, and are ready for the rapidly growing e-business market.

Guru Rao, IBM's chief engineer for S/390 hardware development and newly appointed IBM Fellow, states that over the next several years, S/390 system structure will continue to evolve, driven by the needs of e-commerce. Microprocessors will continue performance scaling through the use of advanced CMOS technology, high-frequency design extending into multigigahertz range, and advanced microarchitectures. The Parallel Sysplex will be integrated into SMP design, leading to converged architectures with flexible resource management across both. Standards-based, high-performance I/O "pipes" will continue to be integrated into the system and will participate in broader workload-balancing decisions. Architectural extensions will be introduced to improve the ability of the S/390 to attract and scale new application workloads. New application development through the Enterprise JavaBeans™-based programming model will be effectively integrated into the

system structure, supporting legacy applications with strict binary compatibility.

This issue of the *IBM Journal of Research and Development* covers a wide assortment of design topics and other advances in hardware development, but describes only a subset of the many advances made to build these products, as in the previous topical issue of the *IBM Journal of Research and Development* on S/390 G3 and G4.

Thanks are due to the many authors from the IBM S/390 Global Hardware Development Laboratory, IBM S/390 Software Development, IBM Microelectronics Division, and IBM Thomas J. Watson Research Center, who have taken time to document these outstanding achievements.

We additionally thank the IBM Microelectronics Division for continued success in delivering leading-edge technology. The significant increases in performance seen in these two generations of servers are a direct result of technology advances, design advances, and the close working relationship between the technology and design teams.

The paper by Katopis et al. describes the multichip module technology and the MCM design for G5/G6. Once again, the design point for this MCM was such that it could support the following G6 server. This "common nest" approach is key to reducing development turnaround time, development expense, and product cost. This paper and the following one, by Rizzolo et al., show that the overall MCM development is driven not only by significant technology requirements but also by aggressive project-management targets for schedule, expense, function, quality, and product cost. Rizzolo et al. describe the significant effort applied to obtaining the highest level of performance from a given design point. Close integration of the design team and the technology team allows for effective test, verification, and selection of the fastest chips manufacturable. The next paper, by Turgeon et al., describes the binodal cache implemented on G5 and G6. The system design for G5 was robust enough to allow the G6 server to handle 14 CP chips (two more than the G5 design point). The paper by Check and Slegel gives an overview of the CP functional design. The design point is fairly constant from G5 to G6. Performance gains have been obtained predominantly by improvements in technology. Also described are functional enhancements such as the branch target buffer, larger L1 cache, enhancements to the instruction fetch buffer, and performance enhancements to instructions that operate on decimal data and instructions that manipulate the program status word. The paper by Averill et al. focuses on the custom physical design methodology that is critical to achieving the performance targets for G5 and G6. This paper highlights the many physical design attributes that must be addressed in order to capitalize on the capabilities of technology and logic design. Physical design

results achieved on both the G5 and G6 microprocessors are outstanding. The next two papers document the new IEEE Floating-Point design. The paper by Schwarz and Krygowski explains the hardware development effort and floating-point algorithms. Abbott et al. then describe the architecture and software support. This functionality allows S/390 to support important open applications such as those written in the Java™ language. Next, the papers by Easter et al. and by Yeh and Smith cover the S/390 CMOS cryptographic coprocessor. This is another function that will allow S/390 to support the growing requirement for encryption on the Internet.

The next group of papers addresses enhancements in S/390 I/O. Gregg et al. describe enhancements to coupling that increase the capabilities of Parallel Sysplex communications. The bandwidth achieved by the new Integrated Cluster Bus and the intersystem coupling interface significantly improves the effectiveness of our Parallel Sysplex environment. DeCusatis et al. then describe the fiber optic interconnects for S/390. The overall strategy of fiber optic interconnects in large systems is discussed, and descriptions are given for the main applications of fiber optic data links on the IBM S/390 platform. Hoke et al. describe the self-timed interface technology, how it was designed, where it is used in the system, and how the interface effectively improves bandwidth. Jackson and Langston describe performance modeling attributes and tradeoff considerations among fully shared, partially shared, and private L2 caches. Considerations in determining the key performance drivers for binodal cache design are covered for uniprocessor and large SMP configurations of G5 and G6. The paper by Rao et al. on ICB performance looks at the key performance factors for the Integrated Cluster Bus and describes reasons for continuing to improve the bandwidth of this connection. A historical perspective is given for the performance levels of ICB since 1994, and the effective coupling capability or cost of a coupling facility access versus host processor speed is described.

One of the strengths of S/390, RAS, is covered by the next two papers. Spainhower and Gregg present a historical perspective of fault tolerance for IBM systems, including the significant advances achieved for the G5 and G6 products. Mueller et al. describe the RAS strategy for G5 and G6. Their paper covers the concept of continuous reliable operation (CRO) and describes the RAS strategy with a set of closely integrated building blocks: error prevention, error detection, error recovery, problem determination, service structure, change management, and RAS measurement and analysis.

The last three papers cover some of the excellent work done to achieve a high level of design quality and ensure that the hardware design will meet its objectives. Buechner et al. explain an event-monitoring and tracing

methodology that helps the development of complex hardware systems. Use of this methodology resulted in a reduction in time to debug problems in a heavily queued system. Song et al. cover the diagnostic capability of the G5/G6 development teams. Three diagnostic methodologies were used, depending on the type of problem. The techniques are software-based diagnostics using TestBench, traditional tester-based diagnostics, and PICA, which is an advanced diagnostic technique that detects device switching activity through the back side of a chip. Finally, Van Huben et al. describe their cycle-simulation environment, which enhanced the capability to model and verify the PLL used in G5 and G6. The clocking scheme for G5 and G6 played a significant role in achieving the overall performance of these products. The success of our clock subsystem development allowed S/390 to overachieve its performance targets while delivering the products to market quickly. Put simply, bringup and integration of hardware cannot proceed without a satisfactory clock subsystem. The clock subsystem for G5 and G6 is an outstanding design.

Kevin Carswell
Vice President, Global Hardware Development
IBM System/390 Division

Cyril Price
Guest Editor

Although I have been with the S/390 Global Development organization for only a short time, I am truly honored to be working with such a world-class team. S/390 is clearly recognized as a leadership platform for e-business. By building upon our many years of experience in enterprise computing, leveraging on leading-edge semiconductor and packaging technology, and relentlessly pursuing innovations in balanced system design, S/390 will continue to provide the scalability, availability, security, and interoperability which customers demand as they grow and evolve their own e-businesses.

Kevin Carswell