

Preface

The IBM System z9™, introduced on September 16, 2005, extends mainframe capabilities with advances in silicon technology, electronic packaging, virtualization, and reliability, availability, and serviceability (RAS), to name a few. Just over a year after its introduction, more than 1,200 Enterprise Class (EC, formerly z9-109) and 600 Business Class (BC) systems have been shipped, evidence that the new technologies are resonating within the mainframe computing arena.

From the outset, the design strategy was to build upon and enhance the framework introduced by the predecessor of the z9™, the IBM z990. This has resulted in significant enhancements to overall system performance and system availability, as well as new and improved functionality with no increase in floorspace requirements and only modest increases in system power consumption. The z9 retains the modular processor book structure and system organization of the z990, but extends performance and availability through a combination of innovative physical packaging and system firmware. The z9 provides nearly twice the total computing capacity measured in MIPS, doubles the available memory capacity to 512 gigabytes, provides up to 172.8 gigabytes per second of internal bandwidth capability, and increases the number of available central processing units, or CPUs, to 54. The z9 offers users the ability to configure up to 60 logical partitions, which is twice the capability of the z990. Further, the z9 has introduced additional specialty processor capabilities which have also been made available for the currently deployed z990 systems.

The System z9 also provides enhancements to the mainframe I/O environment. Perhaps most notable are the performance enhancements made to IBM FICON® channels. The MIDAW (Modified Indirect Addressing Words) facility provides a performance boost for certain types of data transfers by eliminating data chaining overhead. Eight months after the initial z9 offering, a four-port FICON adapter supporting up to 4 Gb/s data bandwidth (an increase from 2 Gb/s in the previous generation) became available. FICON connectivity was extended by eliminating the card-plugging constraints found in previous mainframe systems. Another important system-level constraint, the number of attachable, addressable I/O devices, was relieved with the introduction of multiple subchannel sets, an exclusive z9 capability. Many other new I/O functions such as virtualization for small computer system interface (SCSI) Fibre Channels, IPv6 support for HiperSockets™, and IBM Communications Controller for Linux® further complement z9 connectivity.

The System z9 has delivered several new advances in system availability through the development of the Concurrent Driver Upgrade (CDU) and Concurrent

Book Replacement/Upgrade (CBR) functions, which enable system maintenance and functional upgrades without disruption to the customer workload. They constitute a major step toward the elimination of planned system outages for performing these tasks. Redundant I/O Interconnect helps enable CBR by providing a dynamic means for switching I/O interconnections “on-the-fly” while also providing advanced recovery and concurrent maintenance capabilities. System fault tolerance and recovery are further improved in the z9 by the Enhanced System Recovery framework and Dynamic Oscillator Switchover.

The System z9 also includes enhanced security features working in conjunction with corresponding enhancements to the latest levels of the z/OS®, z/VM®, and z/Linux operating systems.

This double issue of the *IBM Journal of Research and Development* features 18 papers describing many of the new technologies which make the IBM System z9 a unique mainframe computing platform. The topics covered range from silicon and packaging technologies, system design, and advances in higher-availability computing to some of the novel design methodology advances which made these new technologies possible.

The emphasis on security and availability in the System z9 necessitated a high-performance yet stable hardware platform upon which to build these new functions. The paper by Poindexter et al. describes the development of the 90-nm silicon technology used in the z9. The paper addresses general aspects of silicon process development as well as some of the unique issues presented by the z9 microprocessor. The paper by Mayer et al. elaborates further on this theme, describing additional chip- and system-level techniques used to optimize the yield and performance of the custom microprocessor and its supporting chipset.

As in past zSeries® mainframe designs, the System z9 family features the industry’s densest high-performance processor subsystem packaging and interconnection. The ability to package up to 16 processing cores along with large amounts of L2 cache on a single ceramic MCM substrate, then to interconnect up to four of these MCMs with high-speed buses, is unique in the industry. The paper by Harrer et al. addresses the design and verification of this complex system organization, while the following paper, by Berger et al., discusses in greater detail the high-performance circuits and circuit topologies which provide such high-performance chip-to-chip interfaces on the MCM as well as the connections to the memory cards and the system interconnect bus.

Improving chip yields through SRAM wordline redundancy is not a new concept, nor is the concept of storing critical chip manufacturing information within the chip itself. However, the z9 employs a novel eFUSE

circuit technology to accomplish these and other seemingly mundane though critical functions. The paper by Rizzolo et al. describes this eFUSE technology and the methodologies developed for its deployment.

The System z9 has introduced significant new functions to mainframe computing. A major new technology creates new types of processor capability from the basic microprocessor building blocks deployed in the z9. The System z™ Application Assist Processors (zAAPs) and Integrated Information Processors (zIIPs) are described in detail in the paper by Wyman et al. These new specialty processors provide function and performance advantages that were not previously available, although these capabilities have now been made available for currently deployed z990 systems.

Enhancing overall system and enterprise-wide security was a major initiative in the development of the System z platform. Arnold et al. describe a number of the hardware- and firmware-based design enhancements which combine to give the System z unparalleled leadership in security and encryption. IBM mainframes have also long been recognized for leading-edge virtualization technology. A major extension to this virtualization leadership is provided by N_Port Identifier Virtualization for SCSI FCP SAN (Fibre Channel Protocol storage area network) environments. The paper by Srikrishnan et al. describes the System z9 solution to this complex issue. The last paper describing system functional enhancements, by Zee et al., addresses a novel approach to a communications controller design, including enhanced virtualization capability.

The System z9 continues the unparalleled leadership of IBM mainframes in system availability and serviceability. The z9 provides significant fault tolerance improvements via Enhanced System Recovery and Dynamic Oscillator Switchover functions. The paper by Oakes et al. discusses the innovative recovery infrastructure designed into the z9, while the next paper, by Mueller et al., describes a fully redundant, completely transparent system oscillator switchover capability. These functions have contributed to the excellent availability characteristics demonstrated by the z9 thus far.

Reduction of planned system outages to further improve system availability was also targeted by z9 developers. Using innovative hardware packaging and system firmware, it is possible to remove a processor book for repair or hardware upgrade without interrupting customer operation. This feature includes not only book removal and replacement, but also dynamic switching of attached I/O to other, operating books in the system. The paper by Conklin et al. describes the Concurrent Book Repair/Upgrade function, and the paper by Helmich et al. provides a detailed discussion of the Redundant I/O Interconnect function. In addition to book replacement,

another major contributor to planned system outages in prior systems was the installation of new firmware drivers for driver maintenance or delivery of new function. IBM mainframes have long supported concurrent firmware patch capability. This capability has now been extended to complete new firmware drivers as well through the Concurrent Driver Upgrade. Muehlbach et al. describe the challenges of this extremely complex design point and the novel solution to them.

The innovative hardware and firmware technologies introduced in the IBM System z9 would not have been possible without corresponding advances in design methodology and design and verification tools. The next series of papers describe some of these advances. Axnix et al. describe an open-standard-based set of firmware development tools which will continue to provide the mainframe system firmware development framework for some time. Theurich et al. then discuss extensions to the mainframe firmware emulation capabilities which were utilized to improve code quality effectively. Finally, Duale et al. address the challenges associated with the design and verification of an early implementation of the decimal floating-point format in the System z9. Although full decimal floating-point capability is not available in current mainframe systems, this early implementation provides an advance application and operation system test vehicle.

The final paper in this issue, by Webb et al., presents a novel approach to software reuse across platforms which was successfully implemented in the System z9 to quickly provide AIX®-based SCSI device support to the z/VM operating system.

The design of the IBM System z9 was a major development effort requiring the skills of engineers and scientists from many Systems and Technology Group development laboratories. Thanks are due to the many authors from the Systems and Technology Group and the IBM Thomas J. Watson Research Center who have taken the time to document these outstanding accomplishments.

Paul R. Turgeon
Program Manager
System z Hardware Development
IBM Systems and Technology Group

Guest Editor