

Preface

Within the past decade CMOS has become the technology of choice for a broad range of semiconductor products. High-density DRAMs, high-speed processors, and low-power devices for mobile applications are key examples. Underlying this widespread appeal are the distinguishing advantages that CMOS provides: an exceptionally low power-delay product, the ability to accommodate millions of devices on a single chip, and flexible, custom design methodologies which permit optimization, as required, for lowest cost, lowest power, or highest speed.

Accompanying the versatility of CMOS is the economy of providing support for a single technology, rather than for a collection of several disparate ones.

This issue of the *IBM Journal of Research and Development* is a collection of papers on IBM CMOS technology which illustrates this breadth of application and describes the increasing sophistication that is present in the underlying design tools and fabrication techniques.

The opening papers, by Sechler and Grohoski, describe the impact of CMOS on the system design of high-performance processors such as those used in IBM's RISC System/6000 workstation family.

Following these are several papers describing chip designs which together demonstrate the flexibility of CMOS in providing superior solutions to applications that traditionally required much more disparate supporting device technology. Included are contributions by Bernstein et al., describing the trade-off of power for performance in the IBM 601 PowerPC chip, and two papers which show the versatility of CMOS DRAM cores surrounded by application-specific circuitry. The first, by Sunaga et al., is on a wide-I/O (64Kb $\times$ 32) chip for graphics applications; the second, by Ellis et al., demonstrates the flexibility in DRAM applications that can be offered to system designers. The use of CMOS in applications at the interface between technically conflicting signal regimes, such as photonics and analog, is illustrated in the next three papers, by Kuchta et al. on CMOS photodetector/preamplifiers, by Ewen et al. on CMOS-based custom circuits for high-speed data communication, and by Shin et al. on the design of a custom low-power, high-performance CMOS-based digital signal-processing circuit. Closing the section on chip design are a paper by Bechade and Houle, which presents a technique for clock multiplication, and a paper by Dhong et al., which describes a circuit for providing a low-noise CMOS TTL-compatible off-chip driver.

Design automation tools have had to keep pace with the increasing complexity and diversity of CMOS applications, as described in the next three papers. Bose and Surya present techniques for a structured approach to timing verification in high-speed RISC processors. Bergamaschi et al. describe a now commercially available method for

synthesizing gate-level networks from a high-level functional description. Kuehlmann et al. describe a formal verification program for confirming the equivalence of CMOS transistor-level designs and their high-level functional specification.

Today's CMOS technology is the result of several decades of evolution, illustrated by the papers in the final section. Adler et al. describe the influence provided by DRAMs in this evolutionary process. Chesebro et al. discuss the complexity of channel-length control across large chips, an effect which can seriously limit chip performance. The next two papers speak to the current state of the production art: Leonovich et al. describe the process by which a modern silicon production facility learns, and Koburger et al. describe how the base technology provided by DRAMs is extended to provide the features needed for high-performance, low-power logic applications at 2.5 volts. Finally, the extendibility of CMOS is examined in two papers which explore the fringes of today's capabilities: Shahidi et al. take a look at the implications of providing a CMOS technology which operates at 1 volt, and Taur et al. show results recently realized in fabricating CMOS transistors at channel lengths of 100 nm and less.

This issue was brought together by the diligent effort of the following individuals: Peter Cottrell, James Dunn, J. Michael Hancock, Tak Ning, Rosemary Previti-Kelly, W. David Pricer, Robert Sechler, and Lewis Terman, who identified and solicited the papers as well as many of the peer reviewers. We are grateful as well to Karen Papo, who provided invaluable and responsive administrative support. Finally, we acknowledge all those who make up the IBM CMOS technology team. Their contributions are, of necessity, only partially represented by the papers in this issue.

Daniel J. Fleming

Guest Editor

Editor's note: The editors gratefully acknowledge the leadership of guest editor Daniel Fleming in the planning and development of this issue. Formerly Director of External Development for the IBM Microelectronics Division and now retired from IBM, Dan identified and recruited the steering committee, organized the subject matter as he has described it above, obtained the commitment and support of executive management, and provided direction, motivation, and encouragement to all of us.