Preface

In the 1980s the industry's transition from bipolar to CMOS technology was driven by the increased power demand for high-performance circuits. Although the reduced power consumption of CMOS technology initially came at the cost of performance, the density advantage of CMOS was soon turned into a performance benefit. In the early years of CMOS technology, it was inconceivable that the power limit would be reached again. Simple scaling rules dictated the progress of the technology, doubling the transistor count every two years. However, a significant transition appears in the 130-nm to 65-nm-node regime, in which passive power density moves from forming a minor part of the total to becoming dominant. This has effectively halted traditional scaling in CMOS and forced technologists and designers to deal once more with a power-constrained device and chip design. At the core of this dilemma is the MOSFET transistor, which must deliver a maximum drive current at the lowest possible off-leakage current. The undesired leakage from the device subthreshold regime is a manifestation of degraded short-channel control. Another source of leakage is the increase in gate current that is due to thinner gate dielectrics. The latter is responsible for the fact that scaling of traditional oxide-based gate dielectrics comes to a halt at the 90-nm node with a physical dielectric thickness of about 1 nm. Of course, this has a direct impact on gate-length scaling and short-channel control.

Since the total tolerable power dissipation of a chip is limited by the cooling capability of the package, a tradeoff must be made between active and passive power. The goal is to minimize passive power at the benefit of active power and switching speed. The concept of *active power* is best characterized by switched charge multiplied by voltage swing. Switching speed is measured by switched charge divided by switching current. The objective is therefore to minimize capacitances and supply voltages and to increase transistor drive current while simultaneously keeping the leakage current at a reasonable level. This task requires a device with the best possible electrostatic integrity (short-channel behavior) and drive current; the first property is related to the device architecture, and the second to the transport properties of the channel materials.

The advanced silicon technology papers in this issue of the *IBM Journal of Research and Development* discuss device and material options to improve short-channel effects and current drive. Short-channel improvement is obtained through the proper choice of device architecture and the introduction of new gate dielectrics with higher electrical permittivity than conventional oxide-based gate dielectrics. Device drive current benefits from the introduction of metal gates, which eliminates polysilicon depletion. Drive current is also increased by mobility-enhancing techniques such as process-induced strain, the proper choice of substrate materials, or combinations thereof.

In an overview paper, Haensch et al. discuss the possible benefits of new device architectures and materials, with a particular focus on the switching behavior of simple ring oscillators. Next, Antoniadis et al. develop a simple model for the intrinsic switching delay of a device that is based on a source-injection model, drawing conclusions with respect to the scaling behavior of the device. They also review the status of research efforts to improve electron and hole mobility using heterostructures of strained Si and strained SiGe on both bulk and insulator substrates. The next paper, by Shang et al., focuses on germanium as a high-mobility alternative channel material. In the paper that follows, recent progress in gate-stack engineering is reviewed by Gusev et al. Looking toward the future, Hiramoto et al. present current research on silicon nanoscale devices which utilize physical phenomena observed in nanostructures that may prolong the extendability of CMOS devices.

Since a competitive CMOS technology delivers the best power and performance tradeoff for the product, we must understand how best to design transistors in the presence of power constraints. A primary objective is to obtain as much performance as possible for a fixed amount of power. In the end, chip performance, not device performance, is most important to consider for achieving the performance objective. Two contributions investigate this in detail. Frank et al. present a simplified model containing the basic elements for determining chip performance, including intrinsic transistor characteristics, circuit delay, tolerance issues, basic microprocessor composition, power dissipation, and heat removal. This permits the impact of technology options to be studied over a range from single-device to chip-level performance. This model also addresses the impact of manufacturing-related process tolerances on the power–performance balance at the chip level. Bernstein et al. take a more detailed look at process-, device-, and circuit-level responses to systematic and random components of tolerance. Their paper discusses the ways in which exploratory, novel structures emerging as evolutionary CMOS replacements are likely to change the nature of variability in coming generations. The next paper, by Ketchen and Bhushan, completes this excursion with an assessment of relevant higher-level product technology. They present a newly devised concept that relies on product-representative test structures for measuring and characterizing CMOS circuit performance, power, and variability at speeds characteristic of present-day microprocessors. This diagnostic tool provides early

learning in process development, monitors mature processes in manufacturing, and enables model-to-hardware correlation and tracking of product performance.

Energy efficiency has become a ubiquitous design requirement for digital circuits, and aggressive supply-voltage scaling has emerged as the most effective way to reduce energy use. The paper by Hanson et al. reviews circuit behavior at low voltages, specifically in the subthreshold regime, and suggests new strategies for energy-efficient design.

While the contributions mentioned above are anchored in traditional microprocessor design with the MOSFET device as its core, Topol et al. discuss three-dimensional integrated circuits that contain multiple layers of active devices. This technology has the potential to dramatically enhance chip performance, functionality, and device packing density. It is a technology in its infancy, and the presentation touches on the technology components that must be in place in order to take full advantage of the 3D opportunity.

Power, performance, and variability determine the space in which modern silicon technology exists. The contributions in this issue of the *IBM Journal of Research and Development* explore the opportunities in this space and discuss possible paths for the development of future technology nodes.

Wilfried Haensch
Manager, Device and Integration Technology
IBM Research Division

Meikei Ieong
Senior Manager
IBM Research Division