



Intel[®] Technology Journal

Intel's 45nm CMOS Technology

45nm SRAM Technology Development and Technology Lead Vehicle

45nm SRAM Technology Development and Technology Lead Vehicle

Uddalak Bhattacharya, Technology and Manufacturing Group, Intel Corporation
Yih Wang, Technology and Manufacturing Group, Intel Corporation
Fatih Hamzaoglu, Technology and Manufacturing Group, Intel Corporation
Yong-Gee Ng, Technology and Manufacturing Group, Intel Corporation
Liqiong Wei, Technology and Manufacturing Group, Intel Corporation
Zhanping Chen, Technology and Manufacturing Group, Intel Corporation
Joe Rohlman, Technology and Manufacturing Group, Intel Corporation
Ian Young, Technology and Manufacturing Group, Intel Corporation
Kevin Zhang, Technology and Manufacturing Group, Intel Corporation

Index words: SRAM, sleep transistor, forward body bias, ECC, raster, GTL, PLL, DTS, DLL, VCO, QPI

ABSTRACT

Every new generation of process technology at Intel is developed and certified using an SRAM-based “X-chip.” X6 is the technology lead vehicle used for the 45nm technology serving as a platform for the co-optimization of circuit design and process technology for SRAMs as well as critical design collaterals for products.

As the workhorse of the embedded memory, SRAMs play an essential role in all Intel products in achieving power-performance goals. SRAMs are also ideally suited for process-defect sensitivity and detection. The SRAMs on X6 had featured several different SRAM designs and register files that were individually optimized to take advantage of the Hi-K metal gate process for various product applications. Intel’s revolutionary 45nm technology was instrumental for aggressive SRAM scaling. The tileable SRAM array in X6 was architected to directly support product applications. The X6 also served as the vehicle for several critical memory circuit technology developments, including second-generation dynamic sleep control and dynamic Forward Body Bias (FBB). To support process and design learning, X6 includes an infrastructure of advanced test features: for example, an Error Correction Code (ECC) emulator is designed to quantify the benefit of error corrections and a Programmable Built-in Self Test (PBIST) for high-speed testing with raster capability. Fine-granularity In-Die-Variation (IDV) oscillators track process variation. Critical circuits such as electrically programmable Fuse, Phase Lock Loop (PLL), Digital Thermal Sensor (DTS),

and advanced I/Os allowed technologists and designers to work closely to optimize the process and circuits earlier.

The X6 testchip successfully met the goals of process and critical collateral certification to support both process and product development needs and played an essential role in Intel’s rapid product ramp at the 45nm technology node.

INTRODUCTION

“X-chips” have been serving the needs of process technology development at Intel for many generations. They have always been based on SRAM building blocks. SRAM has the unique advantages of stressing the most critical design rules of a process technology, providing large area coverage that increases defect sensitivity, and of providing fine addressable granularity to enable fault isolation and analysis.

In this paper we focus on X6, the testchip that was used to certify the Hi-K metal gate process technology at the 45nm node. Over the generations, the X-chips have evolved beyond the original goal of providing test data needed to identify process defects, to a platform for the co-optimization of process technology and critical design collaterals for the products. In the following section, we provide an overview of the X-chip.

Testability is a key concern with X-chips, and X6 is no exception. The testchip must be sensitive to process defects and must exercise the critical design rules. A process should not get certified on the testchip and encounter yield or performance issues when a product is

subsequently ramped on it. The complexity of modern semiconductor processes results in significant lead time from wafer starts to end of line, while the total time available to certify a process remains, at best, the same. For the process engineer this means only a limited number of “info-turns” during the development phase. To reduce this information turnaround time, X6 includes a test infrastructure that allows the collection of relevant data with reduced test time and features that aid fault isolation. We discuss this aspect further in the test features section.

We now provide descriptions of the critical design collateral content of X6 SRAMs and some key mixed-signal circuits that have high process sensitivities in the context of product design. Therefore, these classes of circuits are obvious choices to be included in the testchip. We discuss a product-ready and tileable common SRAM subarray design along with performance results. In addition to 6T SRAM arrays, X6 also included multi-port SRAM, a.k.a., register file (RF) memories, for the first time to broaden the process and product co-learning. New fuse technology is developed with X6 to provide for expanded product requirements. A significant number of analog elements such as different kinds of Phase Lock Loop (PLL), Digital Thermal Sensors (DTS), and advanced I/O circuits were also part of the X6 testchip. In the following sections these collaterals are described along with the key learnings for technology development and circuit optimization.

X-CHIP OVERVIEW

The X-chip series started many generations ago with the X1 testchip that was the certification vehicle for the 180nm node process technology. From X1 to X6 the role and content of the testchip has grown significantly. However, in the midst of these changes some elements and challenges remain the same.

The memory cell still remains the most critical collateral, the design of which requires careful consideration. The footprint of the memory cell often sets the design rule limit at a given technology node, and it has a large impact on the product die-size, given the tens of megabyte-sized caches of modern CPUs. The stability of the cell at lower voltages often determines the low voltage operation of the products and hence the power consumption. Any systematic process defect mechanism will have a huge impact on product yield due to the tiling of millions of memory bits in product caches. Due to these reasons it has been a product design requirement for several generations to copy exactly the memory bits developed on the X-chips.

The SRAM bits are organized into addressable units called subarrays that have a rectangular matrix-like structure with rows and columns. During a read or write operation from the memory subarray, a specific row and

specific columns are activated depending on the address, and a group of bits called a Word are read or written. A subarray is designed to be very compact since it is tiled many times to form an on-die cache in a real product. The compactness can be achieved by exercising the tightest design rules. Once a subarray design is complete, large portions of the X-chip area can be tiled with these subarrays with minimal additional effort. Thus, the SRAM on the X-chips has always proved to be sensitive to many different kinds of process defects. The fine granular addressability of the compact subarrays allows rapid isolation of the defect location.

On the test infrastructure side X-chips have always carried an I/O interface that can communicate at the desired speeds with all the different tester platforms. A Programmable Built-in Self Test (PBIST) has also been an integral part of X-chips to allow high-speed, on-die testing and burn-in tests. Among the mixed-signal collaterals, X-chips have always included the PLL. This also serves as a clock multiplier for on-die, high-speed SRAM testing.

X-chip design faces some unique challenges compared to product design. All the key electrical parameters of the transistors and interconnects will evolve significantly from the beginning to the final goal of the technology, when the process reaches its maturity. Yet, if the design is not robust enough to comprehend this evolution and still be largely functional, it will not be able to provide robust silicon data for continuous process learning. Testchip circuits are designed largely based on engineering judgment with limited simulation data available to the designers when the design is finalized. The concept of “correct-by-construction” is essential in ensuring the robust design. Layout design rules get defined almost simultaneously with the testchip design process. This creates uncertainties and the inevitable last-minute changes for testchip designers. The impact of design rule uncertainties is mitigated by working closely with the rule-definition teams. Another reason for tight coupling with the design-rule-development activity is that the testchip often throws up scenarios not considered at the time of design rule definition. The testchip provides the first reality check for the process design rule set and validation flows. The design time allowed for the testchip is rather short. The testchip design schedule must be synchronized with the process integration schedule so that it does not become a bottleneck for technology development. Process learning accelerates significantly with the arrival of the first testchip silicon.




Figure 1: X6 reticle

Figure 1 shows a diagram of the X6 reticle. There are two different flavors of die on this reticle, and each die is instanced twice. Thus we have four SRAM-based dice on the reticle, and two dice are allocated for discrete test structures. We focus on the SRAM-based X6 testchip. The two flavors of the X6-SRAM testchip have similar architecture and content. However, having two different flavors allows designers to test different flavors of the same circuits and compare directly the merits of each. While the same can be done on a single testchip design, having two die flavors makes it convenient by adding parallelism to the design process, which then shortens the design time.

X6 Test Features

The main SRAM chip-level tileable unit is called the Raster Unit (RU). An RU is a portion of SRAM convenient for generating raster maps. It is significantly bigger than a small subarray but small enough to be tested in a reasonable time. All RUs can be tested in parallel, which saves test time. The RU-based architecture of the testchip allows easy integration of different flavors of

memory bits and subarray designs while ensuring that each flavor is statistically significant in size.

The chip-level inputs are collected from the I/Os located at two edges of the die. The major busses are located along a center spine of the die connecting the RUs to the I/O and other control blocks. Another center spine in the orthogonal direction contains the major non-SRAM blocks such as the PBIST, Test Access Port (TAP), Fuse unit, and PLL. The synchronous circuits are all located within the RU with the exception of PBIST. There are two clock domains in each RU. One of them is a high-frequency clock and the other a low-frequency clock. The high-frequency clock is contained entirely within the RU, and no synchronous paths communicate between the RUs. This feature makes the chip design very tolerant of clock skew.

The PBIST on X6 can be programmed to generate the test patterns required for SRAM testing. This is useful in cases where the tester cannot support complex test patterns, as in the burn-in test platform and also for high-frequency testing where expensive high-speed testers and complex I/O circuitry can be avoided by on-die testing. The scheme for high-frequency testing is illustrated in Figure 2. The PBIST is part of the low-frequency clock domain. The advantage is that the PBIST design is unlikely to become the bottleneck for frequency and hence is tolerant to the process being off-target in terms of performance, during the development stage. A Parallel-to-Serial (PS) converter converts the low-frequency instructions received as parallel instructions to high-speed serial instructions [1]. An on-die compare circuit compares the “read” data from the SRAM, and the “expect” data from the PBIST pattern to determine pass/fail.

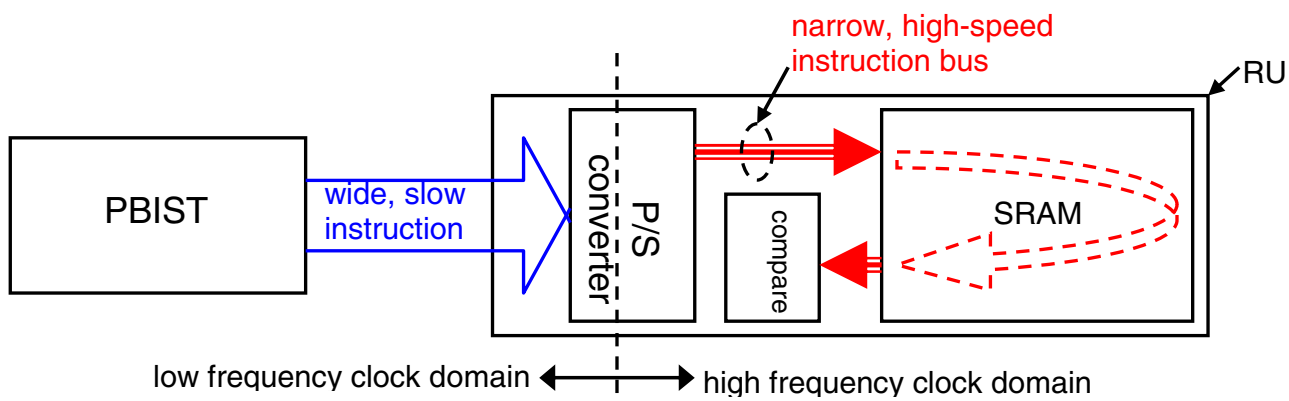


Figure 2: High-frequency testing scheme

The high-frequency testing scheme described above is adequate for determining the SRAM maximum frequency of operation (F_{max}). However, for SRAM design validation and process learnings, it is also quite useful to have raster data at F_{max} to determine what limits F_{max} —especially for products. The high-frequency raster feature will stall the PBIST when a failure occurs. The faulty addresses will be recorded in on-die storage elements and later can be scanned out to identify the location of the faulty bits.

Error Correction Codes (ECC) are becoming very important in modern SRAM design. ECC is an essential part of the large array for data integrity under all product conditions. Since X6 is used to optimize the SRAM memory bit and subarray design, it is important to quantify the impacts of having ECC. X6 has a simple ECC emulator that can quantify the benefit of ECC.

X6 has a TAP that controls the scan function. There are many scan chains dedicated to a specific purpose. Two major roles of the scan chains for process learning are to physically isolate a defect and to support process variation tracking through the use of In-Die-Variation (IDV) oscillators. When a process defect occurs between two stages in pipelined implementation of synchronous circuits, it is possible to feed the logic cone with the defect with various input combinations by using a scan chain and also by recording the output of the logic cone the same way. By analyzing the logic cone output for different input combinations you can, in a large number of cases, pinpoint the failing device. The majority of the circuitry above the subarray level in the RUs is pipelined, and scan provides a useful method of failure isolation and process learning. IDV oscillators are well known for their

ability to track process variation of different process parameters ranging from transistor performance to various kinds of leakage, interconnect delay, and device mismatch [2]. The easiest way to control and extract data from IDV oscillators is through scan chains. X6 provides a robust set of scan chains for this purpose that allow a fine spatial granularity of IDV oscillators on-chip, which otherwise would be impossible to manage using I/O bumps alone.

X6 Memory Design: SRAM and Register File

CMOS technology scaling continues to drive the increase of on-die memory density to meet performance needs in applications such as microprocessors. Meanwhile, the device variation and leakage are increasing as the miniaturization of the transistor continues. As a result, it has become increasingly challenging to design SRAM with an adequate stability margin for low-voltage operation while at the same time keeping the power consumption low enough to meet system-level power requirements. The X6 testchip is a 153Mb SRAM design that is optimized for a 45nm Hi-K Metal-Gate technology [3].

Figure 3 shows the SEM top-down view of the $0.346 \mu\text{m}^2$ 6T-SRAM cell fabricated in the 45nm Hi-K Metal-Gate CMOS technology [3]. The cell design takes full advantage of the new technology features such as trench-contact in achieving high-density and low-voltage operation. The Hi-K Metal Gate transistor technology essentially eliminates the gate leakage in the bitcell. The overall measured cell leakage is reduced by 10x, as shown in Figure 3. This leakage reduction provides enormous benefits in power-constrained applications.

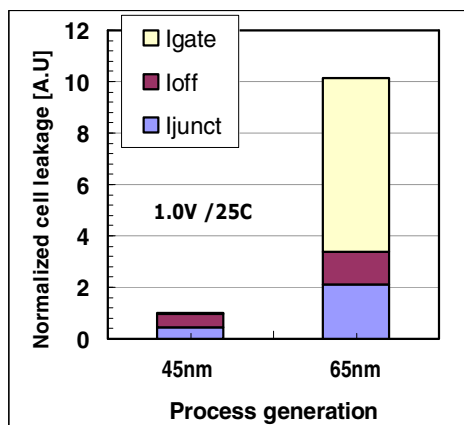



Figure 3: SEM top-down view of 0.346 mm² SRAM cell in 45nm technology (left); SRAM cell leakage comparison between 65nm and 45nm technologies (right)




Figure 4: Measured voltage-frequency schmoos

Figure 5 describes the array architecture and configuration along with critical timing of the array. A 16KB subarray with a 256-row bitline and a 128-column wordline is first

constructed to achieve optimal array efficiency while meeting an aggressive frequency target. The subarray has both built-in column and row redundancies for yield improvement. The terminology in Figure 5 for rows and columns is (M+N) where M is the base number and N is the redundant value. To meet today’s microprocessor bandwidth requirements, the subarray is designed to support 64-bit wide Read and Write. Write is performed with a 64-bit data stream within a single cycle. Read data comes out of the subarray in two consecutive cycles with two 32-bit “chunks,” in order to minimize the global routing congestion. Both Sleep and Forward Body Bias (FBB) have independent control circuits along the 256x256 sector boundary, which provides the fine granularity to achieve balanced design between switching power and area efficiency. The design has achieved up to 3.8GHz operating frequency at 1.1V power supply as shown in Figure 4.




Figure 5: 16-KB subarray configuration and critical signals' timing diagram

The PMOS strength in the 6T SRAM cell is essential to maintain the cell stability during the active mode for low-voltage operation. Using lower threshold voltage PMOS is often prohibitive due to excessive transistor leakage. In X6, a dynamic FBB for the PMOS in the SRAM cell is developed to improve the robustness of low-voltage operation while meeting stringent product and manufacturing requirements at minimum design overhead. Figure 6 shows the critical circuits for the dynamic FBB design.




Figure 6: Dynamic SRAM PMOS FBB circuit

The amount of FBB is determined by the ratio of two PMOS devices (PL and PD), and these devices have a built-in programming control. To meet fast and dynamic requirements, an NMOS pull-down path, formed by

transistors ND and NS and controlled by a pulse signal, is employed to achieve the fast voltage transition at the N-well. A feedback or shutoff mechanism is also used to prevent the N-well voltage from dropping too low and causing excessive junction leakage. The trip points of the inverters, SI1 and SI2, are optimized to meet this need. The pull-down signal pulse is programmable and generated off the wakeup signal. It starts the discharge of the N-well voltage one cycle before the WL is turned on in order to ensure that the N-well voltage level has reached the intended static level. By applying the FBB selectively to the activated portions of the large array, the overall power impact from FBB is kept to a minimum. Test results demonstrate that a wide range of FBB strength can be achieved under high-frequency operation. The stronger PMOS under FBB can improve the minimum operating voltage up to 75mV without increasing the overall SRAM leakage power dissipation.

The use of dynamic sleep design to lower SRAM power supply has been proven to be effective in reducing the static power consumption by reducing leakage [4]. The SRAM voltage control during the sleep state is critical to maintain the integrity of the data stored in the array. High-volume manufacturing requires programming capability [5]. A scheme with active feedback control on SRAM VCC has been proposed to improve Process-Voltage-

Temperature (PVT) variation [6] requiring off-chip voltage reference. In this design, an on-die programmable voltage generator is designed with N-well-based precision resistors. It provides low design overhead as well as insensitivity to different PVT conditions. A simplified Op-Amp further reduces the overall area overhead down to less than 2%. The integrated new design is shown in Figure 7. SRAM VCC is more sensitive to temperature, as

it is dependent on subthreshold leakage, which is the dominant leakage source in Hi-K metal gate technology where gate leakage has been eliminated. The use of the active control along with the temperature-insensitive on-die reference voltage generator provides a much tighter SRAM VCC distribution compared to passive control. This translates into about 100mV lower standby voltage during sleep state and better leakage reduction.




Figure 7: Active SRAM VCC control with integrated on-die programmable reference voltage generation




Figure 8: 45nm Intel® Core™2 Processor with 6MB L2 cache that is formed with the common 16KB subarray

The modular architecture of the X6 SRAM design described above has enabled the 16KB subarray to be used directly as the building block for a 6MB L2 cache in the next-generation Intel® Core™2 processor-based CPU [7]. The die photo of this product with the L2 cache highlighted is shown in Figure 8. This silicon-verified and process-optimized subarray was one of the key contributors to the successful production ramp of this product. Thus, the utility of using a common SRAM design for similar applications across different products to reduce manufacturing risk is clearly established. A

detailed description of the SRAM design can be found in [8].

Multi-port register file memories are very common and important in CPU and other logic products. Due to unique circuit topologies used in this kind of memory, the sensitivity to process variation, e.g., NMOS and PMOS strength ratio, has increased as technology scaling continues. The X6 testchip contained several important RF arrays that are directly used by the lead CPU products at Intel. They have proved to be very effective in optimizing the transistor parameters as well as providing better circuit-modeling enhancement for various designs.

X6 FUSE

Fuse-based electrically programmable read-only memories (PROM) have been widely used in microprocessors for a variety of important applications. These include improving microprocessor yield by repairing defective SRAM bits using row/column redundancy, permanently storing die identification, and effectively trimming devices used in analog circuits such as thermal sensors and I/O. Fuse designs in X6 were optimized for the Hi-K metal gate process. In this generation we developed an array-based design for high-density fuses that significantly improved the area efficiency of the fuse block for increased product applications.

X6 Mixed Signal and Analog Circuits

The X6 technology lead vehicle includes a group of mixed signal and analog collateral circuits. CPU products require increasing amounts of analog circuit content. The analog technology dependencies differ from digital dependencies by device and component type, and they have a stronger focus on gain and operating range metrics as opposed to digital drive current and capacitive load. A technology that produces high-yielding, high-performing SRAM and digital logic does not ensure analog circuit functionality and performance. Including analog circuits on the X6 technology lead vehicle enabled evaluation of unique analog technology requirements.

A set of circuits targeted at the lead 45nm product were incorporated in the design to allow greater technology flexibility. They include PLL, Gunning Transceiver Logic (GTL) I/O buffers, and DTS. The PLL design also includes multiple types of Voltage Controlled Oscillators (VCO) for design optimization. The GTL I/O buffer contains programmable termination resistors for design tuning. All these circuits have been directly used in the lead product design. Figure 9 shows a waveform captured from the product I/O.




Figure 9: GTL I/O waveform (X-axis 1ns/div)

As CPUs demands more data bandwidth, the future generation of Intel CPUs will feature new high-speed serial I/Os, also called Quick Path Interconnect (QPI) circuits. Many circuits used in the physical implementation of QPI circuits have strong process technology dependency. Several critical components of QPI circuits, including differential transceiver, Delay Locked Loop (DLL), duty cycle correction circuits, and low-jitter LC-PLL are incorporated in X6 and provide significant early silicon-based learning. A capability of 6.4GT/s has been demonstrated as shown by the eye-diagram in Figure 10.

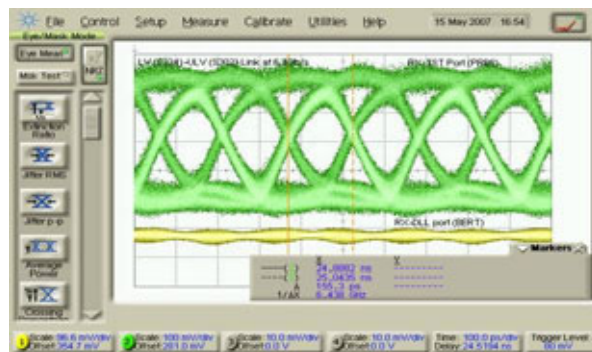


Figure 10: 6.4GT/s operation shows clean eye-diagram and error-free with BIST

SUMMARY

The role of the X6 testchip in the 45nm Hi-K metal gate process development is discussed in this paper. X6 played a critical role in supporting process development by providing a platform for process yield and performance enhancement. In parallel it also served as a test bed for co-optimization of critical design collaterals with the process technology. There exists a significant advantage in having the process technology and the critical design collaterals ready simultaneously for product use. Design considerations and performance of some of the collaterals such as the SRAM memory bit and subarray, PLL, DTS, and high-speed I/Os has been outlined. The early effort on X6 was a key contributor to the success of the lead product ramp in the 45nm Hi-K metal gate technology.

ACKNOWLEDGMENTS

We are grateful to many members of the PTD technical staff including members of advanced memory design, mixed signal circuit collateral development, advanced circuits and technology integration, and process technology teams for their contributions to this work.

REFERENCES

- [1] W. Chen and U. Bhattacharya, "Method for bi-directional data synchronization between different clock frequencies." *US Patent 6,956,918*.
- [2] M. Hatizalambrou, A. Neureuther, and C. Spanos, "Ring oscillator sensitivity to spatial process variation." *1st International workshop on Statistical Metrology (IWSM)*, June 1996.
- [3] K. Mistry, et al., "A 45nm Logic Technology with high-k+Metal Gate Transistors, Strained Silicon, 9 Cu Interconnect Layers, 193nm Dry Patterning, and 100% Pb-free Packaging." *IEDM Technical Digest*, December 2007, pp. 247–250.
- [4] K.-S. Min, K. Kanda, and T. Sakurai, "Row-by-row dynamic source-line voltage control (RRDSV) scheme for two orders of magnitude leakage current reduction of sub-1-V-VDD SRAM's." In *Proceedings IEEE*

International Symposium Low Power Electronics and Design (ISLPED), August 2003, pp. 66–71.

- [5] K. Zhang, et al., “SRAM design on 65-nm CMOS technology with dynamic sleep transistor for leakage reduction.” *IEEE J. Solid-State Circuits*, vol. 40, April 2005, pp. 895–901.
- [6] M. Khellah, et al., “A 256-Kb Dual-VCC SRAM Building Block in 65-nm CMOS Process with Actively Clamped Sleep Transistor.” *IEEE J. Solid-State Circuits*, vol.42, January 2007, pp. 233–242.
- [7] V. George, et al., “Penryn: 45-nm Next Generation Intel® Core™2 Processor.” *ASSCC Technical Digest*, November 2007, pp. 14–17.
- [8] F. Hamzaoglu, et al., “A 153Mb-SRAM design with dynamic stability enhancement and leakage reduction in 45nm High-K Metal-Gate CMOS technology.” *ISSCC Technical Digest*, February 2008, pp. 376–377.

AUTHORS’ BIOGRAPHIES

Uddalak Bhattacharya is a Principal Engineer in the Logic Technology Development Organization at Intel Corporation. His interests include SRAM test circuit design, testing, measurement, and yield analysis. He received his Ph.D. degree in Electrical Engineering from the University of California, Santa Barbara in 1996. His e-mail is uddalak.bhattacharya at intel.com.

Yih Wang is a Senior Design Engineer in Intel’s Logic Technology Development Organization. He joined Intel in 2001 after receiving his Ph.D. degree from the University of Florida. He has worked on SRAM developments for a variety of Intel logic process technologies. Currently, he is responsible for developing the low-power and high-speed SRAM designs for future logic processes. His e-mail is yih.wang at intel.com.

Fatih Hamzaoglu received his Ph.D. degree from the University of Virginia, Charlottesville, in 2002, in Electrical Engineering. After his Ph.D., he joined the Logic Technology Development Group at Intel Corporation. Since then, he has been working on low-power, high-performance SRAM cache design. He is the author/co-author of 19 papers and holds four issued and 6 pending patents. His e-mail is fatih.hamzaoglu at intel.com.

Yong-Gee Ng is a Senior Design Engineer at the Portland Technology Development Group. His interests are in low-power, high-performance SRAM cache designs. He received B.A. and M.S. degrees in Electrical Engineering from Arizona State University, Tempe, in 1989 and 1991, respectively. His e-mail is yong-gee.ng at intel.com.

Liqiong Wei is a Component Design Engineer in the Logic Technology Development Group at Intel. Her technical interests include low-power, high-performance

memory design and validation. She received her B.S. and M.S. degrees in Electrical Engineering from Peking University, Beijing, and a Ph.D. degree in Electrical and Computer Engineering from Purdue University. She is a member of the IEEE. Her e-mail is liqiong.wei at intel.com.

Zhanping Chen received a B.S. degree in Computer Science and Technology from Peking University, Beijing, China, and a Ph.D. degree in Electrical and Computer Engineering from Purdue University, West Lafayette, IN, in 1991 and 1999, respectively. He joined Intel Corporation in June 1999 and currently leads the development of programmable read-only memory (PROM). His research interests include CAD tool development and circuit design/optimization for low power and high performance. Dr. Chen received the Best Paper Award at the 2000 IEEE International Symposium on Quality Electronic Design, San Jose, CA. His e-mail is zhanping.chen at intel.com.

Joe Rohlman is an Engineering Manager in the Advanced Design Group in the Logic Technology Development Organization. He leads a team of engineers focused on mixed signal circuit design on Intel’s next-generation process technology. Prior to joining the Logic Technology Development Organization, Joe held management and design roles on Microprocessor and Ethernet product teams. Joe received an MSEE degree from the University of Michigan. His e-mail is joseph.f.rohlman at intel.com.

Ian Young is a Senior Fellow and Director of Advanced Circuits and Technology Integration in the Technology and Manufacturing Group at Intel. He received his B.A. and M.S. degrees in Electrical Engineering from the University of Melbourne, Australia. He received his Ph.D. in Electrical Engineering from the University of California, Berkeley. He is currently directing the development of high-speed serial I/O and mixed-signal analog circuits in 32nm logic and SOC technologies and doing research for optical I/O. His e-mail is ian.young at intel.com.

Kevin Zhang is an Intel Fellow and Director of Advanced Memory Circuits and Technology Integration in the Logic Technology Development Group at Intel Corporation. He is responsible for embedded memory technology development and has led the design and validation of technology vehicles from 90nm to 45nm technology generations, which played a key role in both process technology development and product design. He holds a Ph.D. degree in Electrical Engineering from Duke University. His e-mail is kevin.zhang at intel.com.

BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino logo, Core Inside, FlashFile, i960, InstantIP, Intel, Intel logo, Intel386, Intel486, Intel740, IntelDX2, IntelDX4, IntelSX2, Intel Core, Intel Inside, Intel Inside logo, Intel Leap ahead., Intel Leap ahead. logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel StrataFlash, Intel Viiv, Intel vPro, Intel XScale, IPLink, Itanium, Itanium Inside, MCS, MMX, Oplus, OverDrive, PDCharm, Pentium, Pentium Inside, skool, Sound Mark, The Journey Inside, VTune, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Intel's trademarks may be used publicly with permission only from Intel. Fair use of Intel's trademarks in advertising and promotion of Intel products requires proper acknowledgement.

*Other names and brands may be claimed as the property of others.

SPEC®, SPECint® and SPECfp® are registered trademarks of the Standard Performance Evaluation Corporation. For more information on SPEC benchmarks, please see <http://www.spec.org>

Microsoft, Windows, and the Windows logo are trademarks, or registered trademarks of Microsoft Corporation in the United States and/or other countries.

Bluetooth is a trademark owned by its proprietor and used by Intel Corporation under license.

Intel Corporation uses the Palm OS® Ready mark under license from Palm, Inc.

Copyright © 2008 Intel Corporation. All rights reserved.

This publication was downloaded from
<http://www.intel.com>.

Additional legal notices at:
<http://www.intel.com/sites/corporate/tradmarx.htm>

For further information visit:

developer.intel.com/technology/itj/index.htm