

# Adaptive Body Bias for Reducing Impacts of Die-to-Die and Within-Die Parameter Variations on Microprocessor Frequency and Leakage

James W. Tschanz, *Member, IEEE*, James T. Kao, *Member, IEEE*, Siva G. Narendra, *Member, IEEE*, Raj Nair, *Member, IEEE*, Dimitri A. Antoniadis, *Fellow, IEEE*, Anantha P. Chandrakasan, *Senior Member, IEEE*, and Vivek De, *Member, IEEE*

**Abstract**—Bidirectional adaptive body bias (ABB) is used to compensate for die-to-die parameter variations by applying an optimum pMOS and nMOS body bias voltage to each die which maximizes the die frequency subject to a power constraint. Measurements on a 150-nm CMOS testchip which incorporates on-chip ABB, show that ABB reduces variation in die frequency by a factor of seven, while improving the die acceptance rate. An enhancement of this technique, that compensates for within-die parameter variations as well, increases the number of dies accepted in the highest frequency bin. ABB is therefore shown to provide bin split improvement in the presence of increasing process parameter variations.

**Index Terms**—Body bias, CMOS digital integrated circuits, forward bias, low-power circuits, microprocessors, parameter variations, substrate bias, within-die variation.

## I. INTRODUCTION

POWER density has become a significant concern in microprocessor design due to the large numbers of transistors integrated in a single die and the increasing clock frequencies. The power density limit of a processor, which is dictated by the thermal design of the system, impacts system cost and maximum operating frequency. Power constraints are even more stringent in mobile processor designs in which long battery life is desirable. At the same time, processors must achieve high frequencies under this power constraint. The goal of a processor design, therefore, is to achieve the maximum operating frequency while meeting the power density constraint.

Process parameter variations result in fabricated dies with variations in maximum operating frequency and power consumption [1], [2]. Within-die variations cause differences in transistor characteristics across a single die. In addition, distribution of device parameters change from die to die within a wafer and across multiple wafers and lots. Together, these variations result in a distribution of die frequencies and leakages. Normally the distribution is divided into several regions, or “bins” and microprocessors are placed into the highest-possible frequency bin that meets the power specification. Some dies




Fig. 1. Measured leakage power and frequency for 62 dies.

cannot be accepted because of either low operating frequency or excessive power consumption.

Fig. 1 shows the measured frequency and leakage for 62 testchip dies on a single wafer. Due to both die-to-die and within-die parameter variations, there is significant variation in both frequency and in leakage power. In order to be accepted, each die must meet a minimum frequency requirement (represented by the vertical dashed line). At the same time, each acceptable die must meet the maximum power consumption requirement. Assuming a worst-case power density of  $20 \text{ W/cm}^2$ , the maximum allowed leakage power for each die can be calculated using the die frequency, switched capacitance and a worst-case activity factor  $\alpha$ . Any die with leakage power exceeding the maximum allowed at its frequency of operation violates the total power constraint. This die must be either accepted at a lower operating frequency, or discarded if the standby leakage power ( $I_{SB}$ ) limit is exceeded. A significant number of dies become unacceptable since they fail to meet one of these two constraints.

The leakage and frequency for a single die can be controlled to some extent using body bias, which employs a non-zero body-to-source bias to modulate the threshold voltage of a transistor. Reverse body bias (RBB)[3]–[5] has been employed in recent years as an effective technique for reducing the leakage power of a design in standby mode by raising the voltage of the pMOS  $N$ -wells with respect to the supply voltage, or by lowering the voltage of the substrate relative to ground. Similarly, forward body bias (FBB) [6]–[8] has been used in active mode to increase the operating frequency of a design, although this increases the leakage power as well. In addition, FBB has the desirable result of improving the short-channel effects of a transistor, thus reducing sensitivity to critical-dimension variation [9].

Manuscript received March 15, 2002; revised June 3, 2002.

J. W. Tschanz, S. G. Narendra, R. Nair, and V. De are with Microprocessor Research, Intel Laboratories, Intel Corporation, Hillsboro, OR 97124 USA.

J. T. Kao, D. A. Antoniadis, and A. P. Chandrakasan are with the Massachusetts Institute of Technology, Cambridge, MA 02139 USA.

Digital Object Identifier 10.1109/JSSC.2002.803949

By controlling the transistor threshold voltages ( $V_T$ ) through process adjustment or body bias, the distribution can be moved to the left (by raising the  $V_T$ ) or to the right (by lowering the  $V_T$ ). Ideally, the  $V_T$  should be lowered for dies that are too slow and raised for dies that are too leaky, maximizing the number of dies in the middle of the distribution that meet both frequency and leakage constraints. This can be accomplished through the use of bidirectional adaptive body bias (ABB) [9]–[12]. This bidirectional ABB technique is described in Section II and a specific testchip implementation is detailed. Measurement results from this ABB test chip are used throughout this paper to demonstrate the effectiveness of ABB. Section III describes a method for correcting die-to-die variations using ABB and measurements show the resulting effect on die acceptance percentage, frequency and leakage power. The effects of within-die parameter variation are explored in Section IV and an enhancement of the ABB technique is presented which compensates for these variations as well. Finally, Section V concludes the paper by summarizing the improvements in die acceptance rate and processor frequency made possible by ABB.

## II. ABB IMPLEMENTATION

Bidirectional ABB allows each die on a wafer to have the optimum threshold voltage which maximizes the die frequency subject to the power constraint. The threshold voltage of each die is controlled not only by process but also by the application of the appropriate amount of FBB or RBB. Dies which are too slow receive FBB, increasing the die frequency as well as the die leakage. Dies that violate the leakage constraint receive RBB, reducing the leakage as well as the frequency. In this way, the combined effect of parameter variations is compensated by changing the  $V_T$  of the devices. In leakage-sensitive circuit topologies, such as dynamic circuits, it is recommended that the amount of forward bias applied be limited by noise constraints, if the die is too slow but the transistors in those circuits already have a low threshold voltage.

The pMOS and nMOS body bias voltages which result in this optimum threshold voltage may be applied by an external source or by an on-chip body bias generator. Similarly, the control circuitry that determines the optimum body bias voltage may be implemented off-chip or integrated in the die. We have implemented a testchip that incorporates the body bias generator and control circuitry for pMOS devices on die and allows off-chip control and biasing for the nMOS devices. This testchip has been fabricated in 150-nm CMOS technology and demonstrates the effectiveness of bidirectional ABB.

### A. Testchip Components

Each testchip die [Fig. 2(a)] contains 21 “subsites” distributed over a  $4.5 \times 5.3 \text{ mm}^2$  area in two orthogonal orientations. Each of these subsites [Fig. 2(b)] represents a circuit block of a microprocessor design and contains a complete ABB generator and control circuit in addition to critical path circuit blocks. The organization of a testchip subsite is shown in Fig. 3. The circuit block, or CUT, contains key circuit elements of a microprocessor critical path, meant to model the effect of body bias on the frequency and leakage of a real microprocessor design.

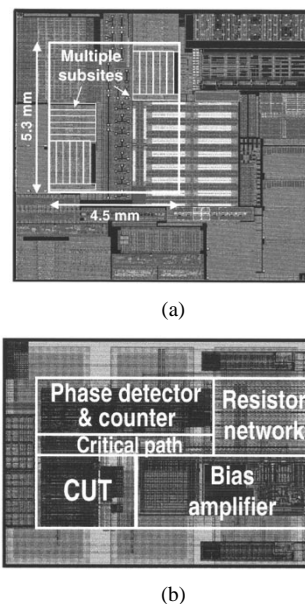


Fig. 2. (a) Die micrograph, noting locations of testchip subsites. (b) Micrograph of a single testchip subsite.

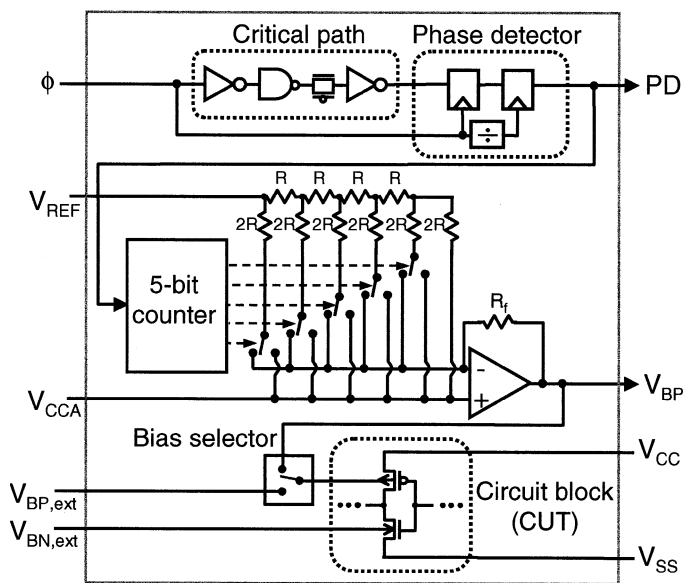


Fig. 3. Block diagram of ABB testchip.

This circuit block can be configured in a ring-oscillator structure for frequency measurements, or disabled to allow leakage power measurements. One critical path from this circuit block is replicated and a target clock frequency  $\phi$  is applied externally. This represents the desired frequency of operation for the circuit block. The output of the critical path is sampled by a phase detector structure consisting of two flip-flops which compare the critical path delay with the target clock period. The output of the first phase detector flip-flop is sampled by a second flip-flop clocked with a divided clock signal. This allows sufficient time for the body bias generator to stabilize and the critical path frequency to adapt to the new body bias before the phase detector is updated again. The output from this second flip-flop, denoted “PD,” is used to clock a 5-b digital counter whose value represents the desired body bias to apply. Finally, it is necessary to

Bias Mode	Condition	Range
NBB → FBB	$V_{CCA} = V_{CC}$ $V_{REF} > V_{CCA}$	FBB: $0 \rightarrow V_{REF} - V_{CCA}$
NBB → RBB	$V_{CCA} = V_{CC}$ $V_{REF} < V_{CCA}$	RBB: $0 \rightarrow V_{CCA} - V_{REF}$
FBB → RBB	$V_{CCA} < V_{CC}$ $V_{REF} < V_{CCA}$	FBB: $V_{CC} - V_{CCA} \rightarrow$ RBB: $2V_{CCA} - V_{REF} - V_{CC}$

Fig. 4. Bias modes provided by on-chip body bias generator.




Fig. 5. Adaptive bias operation.

convert this 5-b digital code to an analog body voltage which is applied to the pMOS transistors in both the CUT and in the critical path. This is accomplished through the use of a D/A converter consisting of an R-2R resistor network and an op-amp driver. The output voltage of this body bias generator, which biases the pMOS transistors in the CUT and in the critical path, is a function of the 5-b counter value as well as the supply voltages  $V_{REF}$  and  $V_{CCA}$ . Different ranges of unidirectional—FBB or reverse—or bidirectional body bias values (Fig. 4) can be selected by using appropriate values of  $V_{REF}$  and  $V_{CCA}$  and by setting a counter control bit. For the desired body-bias range of 500-mV RBB to 500-mV FBB, (a range of 1 V in 32 steps) this bias generator can achieve 32-mV bias resolution. ABB can also be accomplished by using the phase detector output to continually adjust off-chip bias generators through software control, instead of using the on-chip circuitry, until the frequency target is met.

### B. Testchip Operation

For a specific, externally applied nMOS body bias, this on-chip circuitry automatically generates the pMOS body bias that minimizes leakage power of the CUT while meeting a target clock frequency. An example of measured testchip operation is shown in Fig. 5. Initially, the operating frequency of the circuit block is lower than the desired target frequency. The body voltage therefore reduces, forward-biasing the pMOS transistors and increasing the frequency. This continues until the circuit frequency matches the desired frequency target, at which point the phase detector output switches to a one, disabling the counter. This adaptation process is done at high temperature since this represents the worst-case frequency of the part as well as the largest leakage power expected. Once the optimum bias voltages are determined, they may be permanently programmed into the chip in the case of an on-chip bias gen-

erator, or programmed into the external voltage regulators supplying the off-chip bias voltage.

## III. DIE-TO-DIE ABB

Using only a single pMOS/nMOS body bias combination per die, ABB can be used to compensate for die-to-die variations. The goal of the ABB procedure is to find the optimum pMOS/nMOS bias combination that maximizes the die frequency while meeting the leakage constraint. Because each subsite represents a critical path of a microprocessor design, the die clock frequency is the minimum of the CUT frequencies and active leakage power is the sum of the CUT leakages. We have evaluated several ABB techniques through measurements on 21 CUTs per die at 1.1 V  $V_{CC}$  and 110 °C, for 62 dies on a wafer.

### A. Simple ABB (S-ABB)

In the simplest ABB scheme, the optimum bias voltages are determined through measurements on a single circuit block on the die. Consider, for example, a microprocessor consisting of many large circuit blocks, each of which is equally critical. One of these circuit blocks contains a replica critical path and a phase detector which communicates with a central body bias generator. Based on the frequency of this critical path, the central bias generator determines the body bias which must be applied to meet a target frequency and this body bias is applied to all circuit blocks on the die. Because the optimum body bias is determined through measurement on a single circuit block, this method therefore ignores any within-die variation that results in differences in frequency and leakage among circuit blocks fabricated on the same die.

The central bias generator, along with the required body bias buffers and additional bias routing, incurs  $\sim 2\%$  total die area overhead [6]. Because the current in the transistor body is at least two orders of magnitude smaller than the supply current, the body bias grid is significantly less dense than the supply grid. The body bias buffers, which are distributed throughout the die, are designed so that the body bias voltage tracks variations in the local supply voltage of each circuit block. In addition, these buffers ensure that any noise coupled to the body bias grid does not exceed the expected noise for a design without body bias. No noise-related effects have been observed through measurements of body-biased circuits when compared to traditional implementations without body bias [6].

The optimum bias combination is determined by applying a target clock frequency to the die. This represents the highest-possible operating frequency for the die under the given power constraint, achievable only when leakage power is zero. This maximum clock frequency is shown on Fig. 1 as the intercept of the slanted leakage line with the frequency axis. For that target frequency, an nMOS bias is applied from an off-chip source and the on-chip control circuitry is enabled to automatically adapt the pMOS body bias to meet the frequency target. The leakage power for this combination is measured and the process repeated for all nMOS bias levels. The combination which results in lowest leakage is chosen as the best choice

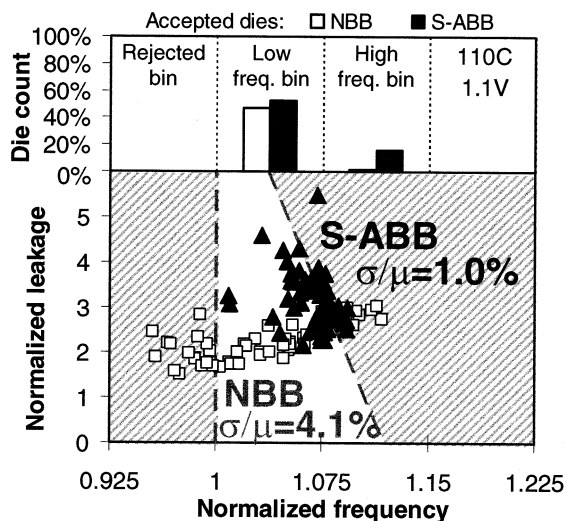


Fig. 6. Leakage versus frequency distribution for NBB and S-ABB.

for the target frequency and the leakage is compared to the maximum allowed leakage for that circuit block. If the leakage is higher than the constraint, the frequency target must be lowered and the process repeated. Otherwise, the optimum bias combination has been found.

The effect of S-ABB is shown in Fig. 6. The leakage versus frequency distributions for NBB and S-ABB are shown on the bottom plot, while the histogram shows the number of acceptable dies in each frequency bin for the two techniques. Here the frequency range is divided into two acceptable bins and the sum of the two bins gives the total number of accepted dies. When no body bias (NBB) is used, only 50% of the dies are acceptable, mainly in the lowest frequency bin. ABB applied to the same 62 dies reduces the frequency variation  $\sigma/\mu$  from 4.1% to 1.0%, resulting in a larger number of accepted dies and more dies in the highest frequency bin. However, there are still a significant number of dies which fail to meet the leakage constraint. These arise from the fact that a single circuit block is used to determine the body bias for the entire die, while within-die variations cause differences in circuit block frequency and leakage.

**B. Die-to-Die ABB**

An improved ABB technique takes these within-die variations into account when determining the best pMOS/nMOS bias combination per die. The technique is similar to S-ABB except that each circuit block requires its own phase detector structure and the central bias generator takes into account each phase detector output when determining the appropriate body bias voltage. This can be accomplished by combining the individual phase detector signals with an “OR” structure so that the bias generator counter is updated if any circuit block does not meet the target frequency. Area overhead is slightly higher than for S-ABB due to the additional phase detectors required, but still only 2%–3% of the total die area. This technique therefore works in the presence of within-die variations, as demonstrated in Fig. 7. When compared to NBB, ABB reduces the frequency variation  $\sigma/\mu$  to 0.69%, resulting in 100% acceptable dies. Furthermore, 32% of the dies now fall




Fig. 7. Leakage versus frequency distribution for NBB and ABB.

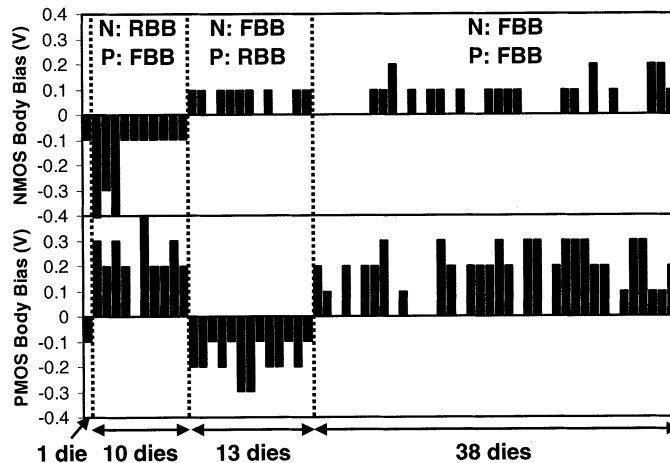


Fig. 8. Histogram of applied bias voltages for ABB.

in the highest frequency bin. Therefore die-to-die ABB results in an increase in the number of acceptable dies, as well as a net frequency increase. It is evident, however, that in order to be most effective, within-die variations must be considered.

The distribution of body bias voltages resulting from ABB is shown in Fig. 8. The 62 dies are grouped into four categories based on the pMOS and nMOS bias required. A majority of the dies (38) receive forward or zero bias for both pMOS and nMOS transistors, while a smaller number require a combination of forward and reverse bias. Only a single die uses reverse bias for both pMOS and nMOS transistors. In addition, a modest bias range is required –300 mV of forward or reverse bias is sufficient for most dies.

**IV. WITHIN-DIE VARIATION COMPENSATION**

By comparing the results of S-ABB and ABB, it is evident that within-die variations influence the effectiveness of the ABB technique. In this section, we examine the impact of within-die variations as technology scales and describe a method of compensating these variations as well.




Fig. 9. Die frequency distribution versus number of critical paths ( $N_{CP}$ ) per die.




Fig. 10. Device  $I_{ON}$  variation versus critical path delay variation.

#### A. Effects of Within-Die Variation

The relative impacts of die-to-die and within-die parameter variations depends on the number of critical paths ( $N_{CP}$ ) per die. In [2] it is shown through statistical simulations that as the number of critical paths per die increases, the within-die delay variation causes the mean frequency  $\mu$  to reduce while the variation  $\sigma$  in critical path delays reduces as well. This is confirmed through testchip measurements as shown in Fig. 9. As the number of critical paths per die increases from 1 to 14, the mean frequency reduces by  $\sim 8\%$  and  $\sigma$  reduces as well. As  $N_{CP}$  exceeds 14, however, there is no significant change in the frequency distribution. Because the effect of  $N_{CP}$  on die frequency saturates above  $N_{CP} = 14$ , using measurements of 21 critical paths per testchip die to determine die frequency is sufficiently accurate for obtaining frequency distributions of microprocessors which contain  $100$ 's of critical paths.

The impact of within-die variations is also a function of the number of gate stages in the critical paths. Previous measurements [10] on 49-stage ring oscillators showed that  $\sigma$  of the WID frequency distribution is  $4\times$  smaller than  $\sigma$  of the device saturation current ( $I_{ON}$ ) distribution. However, measurements on the testchip containing 16-stage critical paths (Fig. 10) show that  $\sigma$ 's of WID critical path delay distributions and nMOS/pMOS  $I_{ON}$  distributions are comparable. Since typical microprocessor critical paths contain 10–15 stages and this number is reducing by 25% per generation [13], impact of within-die variations on frequency is becoming more pronounced.

#### B. Within-Die ABB (WID-ABB) Technique

While the improved ABB scheme considers within-die variations in determining the optimum bias combination per die, it is not possible to compensate for these variations using only a single bias combination per die. The WID-ABB technique al-

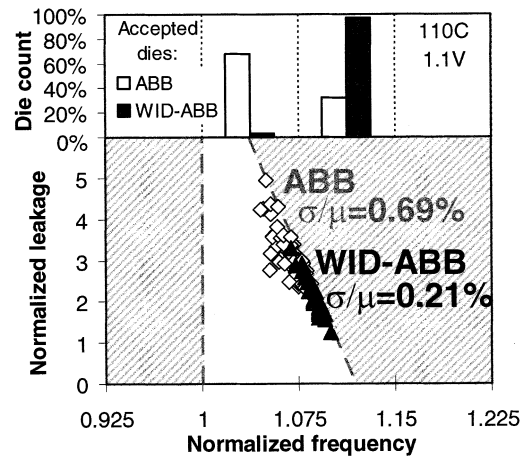


Fig. 11. Leakage versus frequency distribution for die-to-die ABB and WID-ABB.

lows each large circuit block in the design to have its own unique bias combination which controls the frequency and leakage of only that circuit block. Therefore, in this technique, each circuit block requires its own phase detector as in ABB, while the central bias generator must be capable of generating several bias voltages and selecting the appropriate bias for each circuit block on the die. Alternately, the adaptation circuitry from the central bias generator may be distributed as well, although a central bias generator is still required to generate a PVT-invariant reference voltage. In the WID-ABB scheme, each subsite independently adjusts its own body bias to meet the applied frequency target and therefore the total die leakage is minimized for the target frequency. For nMOS implementation, a triple-well process is necessary. While this technique requires a more complex central bias generator, the area of this block is amortized over the entire die and thus the area overhead is similar to ABB.

Fig. 11 shows the effect of WID-ABB when compared to the ABB scheme. The frequency variation has been further reduced by a factor of 3 and again all dies are accepted. In addition, 99% of the dies are now accepted in the highest frequency bin, compared to 32% for ABB. By using multiple bias voltages per die, precise control over the die frequency and leakage is possible. The result is an improvement in die acceptance rate and in frequency bin split. A histogram of optimum body bias voltages for circuit blocks on several adjacent dies (Fig. 12) shows that a significant number of blocks receive RBB for both pMOS and nMOS transistors, enabling large leakage savings. In addition, FBB is mainly required for the pMOS devices. This suggests that the skew for this wafer favors fast nMOS and slower pMOS devices, which is corrected using body bias. The bias distribution changes from wafer-to-wafer and from lot-to-lot as the processing conditions change.

The complexity and overhead of the body bias generator and control circuitry depends on the required resolution in body bias voltage. Fig. 13 shows the effect of bias resolution on the effectiveness of the ABB and WID-ABB techniques. It is evident that 300-mV bias resolution is sufficient for die-to-die ABB while 100-mV resolution is required for WID-ABB. The 32-mV resolution provided by the on-chip bias generator is therefore sufficient for either application.




Fig. 12. Histogram of applied body biases for circuit blocks using WID-ABB.

Bias resolution	ABB		WID-ABB	
	dies, $F > 1$	$\sigma/\mu$	dies, $F > 1.075$	$\sigma/\mu$
500mV	79 %	2.87 %	2 %	1.89 %
300mV	100 %	1.47 %	66 %	0.50 %
100mV	100 %	0.69 %	97 %	0.21 %

Fig. 13. Effect of bias resolution on effectiveness of ABB and WID-ABB.

V. CONCLUSION

The bidirectional ABB technique has been shown to reduce the effects of process parameter variations on the frequency and leakage of microprocessor dies. Using a single body bias value per die, ABB reduces  $\sigma$  of the die frequency distribution by 7  $\times$ , compared to NBB, allowing all dies to meet the given minimum frequency and maximum power constraints. The number of high-frequency parts can be further increased by the application of WID ABB, reducing  $\sigma$  of the die frequency distribution by an additional 3  $\times$  and allowing virtually 100% of the dies to be accepted in the highest frequency bin. The ABB technique therefore allows die acceptance rate and bin split improvement in the presence of increasing process parameter variations.

ACKNOWLEDGMENT

The authors wish to thank S. Menon, B. Bloechel, M. Borkar, and G. Dermer for measurements and lab support; J. Rattner, S. Borkar, R. Hofsheier, F. Pollack for encouragement and support.

REFERENCES

[1] K. A. Bowman, S. G. Duvall, and J. D. Meindl, "Impact of die-to-die and within-die parameter fluctuations on the maximum clock frequency distribution," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2001, pp. 278–279.

[2] —, "Impact of die-to-die and within-die parameter fluctuations on the maximum clock frequency distribution for gigascale integration," *IEEE J. Solid-State Circuits*, vol. 37, pp. 183–190, Feb. 2002.

[3] A. Keshavarzi, S. Ma, S. Narendra, B. Bloechel, K. Mistry, T. Ghani, S. Borkar, and V. De, "Effectiveness of reverse body bias for leakage control in scaled dual  $V_t$  CMOS ICs," in *Proc. ISLPED*, Aug. 2001, pp. 207–212.

[4] T. Kuroda, T. Fujita, S. Mita, T. Nagamatsu, S. Yoshioka, K. Suzuki, F. Sano, M. Norishima, M. Murota, M. Kako, M. Kinugawa, M. Kakumu, and T. Sakurai, "A 0.9-V, 150-MHz, 10-mW, 4 mm<sup>2</sup>, 2-D discrete cosine transform core processor with variable threshold-voltage (VT) scheme," *IEEE J. Solid-State Circuits*, vol. 31, pp. 1770–1779, Nov. 1996.

[5] S. Thompson, I. Young, J. Greason, and M. Bohr, "Dual threshold voltages and substrate bias: Keys to high performance, low-power, 0.1  $\mu$ m logic designs," in *VLSI Technology Dig. Tech. Papers Symp.*, 1997, pp. 69–70.

[6] S. Narendra, M. Haycock, V. Govindarajulu, V. Erraguntla, H. Wilson, S. Vangal, A. Pangal, E. Seligman, R. Nair, A. Keshavarzi, B. Bloechel, G. Dermer, R. Mooney, N. Borkar, S. Borkar, and V. De, "1.1 V 1 GHz communications router with on-chip body bias in 150 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2002, pp. 270–271.

[7] S. Vangal, N. Borkar, E. Seligman, V. Govindarajulu, V. Erraguntla, H. Wilson, A. Pangal, V. Veeramachaneni, M. Anders, J. Tschanz, Y. Ye, D. Somasekhar, B. Bloechel, G. Dermer, R. Krishnamurthy, S. Narendra, M. Stan, S. Thompson, V. De, and S. Borkar, "A 2.5 GHz 32b integer-execution core in 130 nm dual-Vt CMOS," in *IEEE ISSCC'02 Dig. Tech. Papers*, Feb. 2002, pp. 412–413.

[8] Y. Oowaki, M. Noguchi, S. Takagi, D. Takashima, M. Ono, Y. Matsunaga, K. Sunouchi, H. Kawaguchiya, S. Matsuoka, M. Kamoshida, T. Fuse, S. Watanabe, A. Toriumi, S. Manabe, and A. Hojo, "A sub-0.1  $\mu$ m circuit design with substrate-over-biasing," in *IEEE ISSCC Dig. Tech. Papers*, 1998, pp. 88–89.

[9] M. Miyazaki, G. Ono, T. Hattori, K. Shiozawa, K. Uchiyama, and K. Ishibashi, "A 1000-MIPS/W microprocessor using speed-adaptive threshold-voltage CMOS with forward bias," in *IEEE ISSCC Dig. Tech. Papers*, 2000, pp. 420–421.

[10] M. Miyazaki, G. Ono, T. Hattori, and K. Ishibashi, "A 1.2-GIPS/W microprocessor using speed-adaptive threshold-voltage CMOS with forward bias," *IEEE J. Solid-State Circuits*, vol. 37, pp. 210–216, Feb. 2002.

[11] J. Tschanz, J. Kao, S. Narendra, R. Nair, D. Antoniadis, A. Chandrakasan, and V. De, "Adaptive body bias for reducing impacts of die-to-die and within-die parameter variations on microprocessor frequency and leakage," in *IEEE ISSCC Dig. Tech. Papers*, 2002, pp. 422–423.

[12] J. Kao, "Subthreshold Leakage Control Techniques for Low Power Digital Circuits," Doctoral dissertation, MIT, Cambridge, 2001.

[13] V. De and S. Borkar, "Technology and design challenges for low power and high performance," in *Proc. ISLPED*, Aug. 1999, pp. 163–168.



**James W. Tschanz** (S'94–M'99) received the B.S. degree in computer engineering in 1997, and the M.S. degree in electrical engineering in 1999, both from the University of Illinois at Urbana-Champaign.

Since 1999, he has been a Circuits Researcher at Intel Laboratories, Hillsboro, OR. His research interests include low-power digital circuits, design techniques, and methods for tolerating parameter variations. He is an adjunct faculty member at the Oregon Graduate Institute, Beaverton, OR, and has authored several papers and patents pending.



**James T. Kao** (M'01) received the B.S. degree in electrical engineering and computer science from the University of California at Berkeley in 1993, and the M.S. and Ph.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, in 1995 and 2001, respectively, with a Ph.D. dissertation entitled "Subthreshold Leakage Control Techniques for Low-Power Digital Circuits."

He is currently with Silicon Laboratories, Austin, TX, designing phase-locked loops. His interests include low-power and mixed signal circuits.



**Siva G. Narendra** (S'91–M'99) received the B.E. degree from Government College of Technology, Coimbatore, India, in 1992, the M.S. degree from Syracuse University, Syracuse, NY, in 1994, and the Ph.D. degree from Massachusetts Institute of Technology, Cambridge, in 2002.

He has been with Intel Laboratories, Hillsboro, OR, since 1997, where his research areas include low-voltage MOS analog and digital circuits and impact of MOS parameter variation on circuit design. He is an Adjunct Faculty with the Department of

Electrical and Computer Engineering, Oregon State University, Corvallis. He has authored or co-authored over 16 papers and has 15 issued and 27 pending patents in these areas.


Dr. Narendra is an Associate Editor for the IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION SYSTEMS and a Member of the Technical Program Committee of the 2002 International Symposium on Low-Power Electronics and Design.



**Raj Nair** (M'98) obtained the Bachelor of Engineering degree from the University of Mysore, Mysore, India, in 1986, and the M.S.E.E. degree from Louisiana State University, Baton Rouge, in 1994.

He spent about eight years developing and implementing machine automation systems, test and measurement instrumentation, and signal conditioning and data acquisition systems. He joined Intel, Hillsboro, OR, in 1995. He was the architect, designer, and implementer of Intel's first

on-chip distributed voltage regulation system. He was also an architect and designer of a CMOS image sensor chip for the Intel 971 Digital Camera Kit and researched CPU clocking, power delivery, packaging, and signaling at Intel's Microprocessor Research Laboratories before joining Intel's Assembly Technology Development (ATD) division of the Technology and Manufacturing Group (TMG), where he is currently responsible for project management in ATD's strategic initiatives enabling processor power delivery and IO. His interests also include analog circuit design, optical interconnects and micro-electromechanical devices, and actuation techniques. He holds ten U.S. patents and has numerous publications pertaining to voltage regulation, digital image sensing, digital clock distribution, analog circuit design and high-speed signaling.



**Dimitri A. Antoniadis** (M'79–SM'83–F'90) received the B.S. degree in physics from the National University of Athens Athens, Greece, in 1970, and the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, in 1976.

In 1978, Dr. Antoniadis joined the faculty at Massachusetts Institute of Technology (MIT), Cambridge, where he currently holds the Ray and Maria Stata Chair in Electrical Engineering. He was co-founder and first Director of the MIT Microsystems Technology Laboratories, and from 1993 to

2000, he was Director of the SRC MIT Center of Excellence for Microsystems Technology. Currently, he is Director of the National, Multi-University Focus Research Center for Materials, Structures and Devices, centered at MIT. His initial research activities covered the area of measurement and modeling of the earth's ionosphere and thermosphere ranging from instrument design to computer simulation. He led the development of the first two generations of the SUPREM process simulator and since then, his technical activity has been in the area of semiconductor devices and integrated circuit technology. He has worked on the physics of diffusion in silicon, thin-film technology and devices and quantum-effect semiconductor devices. His current research focuses on the physics and technology of extreme-submicron Si, SOI and Si/SiGe MOSFETs.

Dr. Antoniadis is the recipient of the Solid State Science and Technology Young Author Award of the Electrochemical Society in 1979, the Paul Rappaport Award of the IEEE in 1998, and the 2002 Andrew Grove Award of the IEEE.



**Anantha P. Chandrakasan** (S'87–M'95–SM'01) received the B.S., M.S., and Ph.D. degrees in electrical engineering and computer sciences from the University of California, Berkeley, in 1989, 1990, and 1994, respectively.

Since September 1994, he has been at the Massachusetts Institute of Technology, Cambridge, and is currently an Associate Professor of Electrical Engineering and Computer Science. His research interests include the ultra low power implementation of custom and programmable digital signal processors,

distributed wireless sensors, multimedia devices, emerging technologies and CAD tools for VLSI. He is the co-author of *Low Power Digital CMOS Design* (Norwell, MA: Kluwer Academic, 1995) and the co-editor of *Low Power CMOS Design* (Piscataway, NJ: IEEE Press, 1997) and *Design of High-Performance Microprocessor Circuits* (Piscataway, NJ: IEEE Press, 2000).

Dr. Chandrakasan received the Analog Devices Career Development Chair from 1994 to 1997, the NSF Career Development award in 1995, the IBM Faculty Development award in 1995, and the National Semiconductor Faculty Development Award in 1996 and 1997. He has received several best paper awards including the 1993 IEEE Communications Society's Best Tutorial Paper Award, the IEEE Electron Devices Society's 1997 Paul Rappaport Award for the Best Paper in an EDS publication during 1997 and the 1999 Design Automation Conference Design Contest Award. He has served on technical program committees for various conferences including ISSCC, VLSI Circuits Symposium, DAC, and ISLPED, as a Technical Program Co-chair for the 1997 International Symposium on Low-Power Electronics and Design (ISLPED), VLSI Design 1998, and the 1998 IEEE Workshop on Signal Processing Systems, as a General Co-Chair of the 1998 ISLPED, as an Elected Member of the Design and Implementation of Signal Processing Systems (DISPS) Technical Committee of the Signal Processing Society, the Signal Processing Sub-Committee chair for ISSCC 1999 through 2001, and the Program Vice-Chair for ISSCC 2002, and is the Technical Program Chair for ISSCC 2003. He was also an Associate Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS from 1998 to 2001.



**Vivek De** (S'86–M'92) received the Ph.D. degree in Electrical Engineering from Rensselaer Polytechnic Institute, Troy, New York in 1992.

He is a Principal Engineer and Manager of Low Power Circuit Technology at Microprocessor Research of Intel Labs, Hillsboro, OR. He has authored 82 technical papers in refereed international conferences and journals and two book chapters on low power design. He has 23 issued patents and 45 more patents filed (pending).

Dr. De served as Technical Program Chair of 2001 ACM/IEEE International Symposium on Low Power Electronics and Design (ISLPED'01), General Chair of ISLPED'02 and Technical Program Chair of 2002 ACM Great Lakes Symposium on VLSI. He served on technical program committees of ARVLSI and ISQED conferences. He is the guest editor of a special issue on low power electronics for IEEE TRANSACTIONS ON VLSI SYSTEMS and an adjunct faculty at the Department of Elec and Computer Engineering at Oregon State University. He is the recipient of a best paper award at the 1996 IEEE International ASIC Conference in Portland, OR.