Reliable and Efficient System-on-Chip Design

The recent emergence of noise and increase in process variations raise serious questions about our ability to design reliable and efficient computing systems using nanometer processes. A communication-theoretic design paradigm has been proposed as the solution.

Naresh R. Shanbhag
University of Illinois at Urbana-Champaign

Power dissipation is a concern in both microprocessors and communication systems. High power dissipation increases the substrate temperature of integrated circuits, which increases leakage currents, reduces performance and battery life for mobile applications, and adversely impacts material reliability.

Supply noise (bounce and IR drops), leakage, and interconnect noise (coupling) impact signal phase and amplitude, while process variations result in uncertainty and create a mismatch between signal paths. Both noise and process variations impact reliability, causing logic errors that can result in system failure.

To increase processor performance, the microprocessor industry is driving the scaling of feature sizes into the deep-submicron (DSM) and sub-100-nanometer regime. Unfortunately, power/performance-enhancing design techniques only aggravate the reliability problem. For example, the popular supply voltage scaling technique reduces power, but it does so at the expense of noise immunity.

Although researchers have developed complex power management systems and expensive packaging schemes, the recent emergence of noise and the dramatic increase in process variations have raised serious questions about the capacity to use nanometer process technologies to design reliable and low-power/high-performance computing systems. These concerns put at risk the affordability of microsystems and jeopardize the semiconductor industry's ability to extend Moore's law into the nanometer realm.

The design and electronic design automation (EDA) communities must work closely with the process engineering community to address these problems. Researchers in academia and industry have taken major steps in this direction by establishing the multiple-university Gigascale Silicon Research Center and the Center for Circuit & System Solutions. Both centers, funded through the Microelectronics Advanced Research Corporation (MARCO) by the Semiconductor Industry Association and the Defense Advanced Research Projects Agency, initiated reliability research thrusts beginning in 2003.

Microprocessor designs must achieve high performance and energy efficiency in the presence of noise. A *communication-theoretic* paradigm[1] for reliable and efficient system-on-chip (SoC) design views integrated microsystems as miniature communication networks operating in the presence of noise. First proposed in 1997, this paradigm has evolved into two distinct but related areas of research:

- information-theoretic techniques for determining the lower bounds on energy efficiency in the presence of noise,[1-4] and
- circuit[5] and algorithmic noise-tolerance techniques[6] to approach these bounds.

The 2003 *International Technology Roadmap for Semiconductors* (http://public.itrs.net/Files/2003ITRS/Home2003.htm) echoes the need for a communication-centric SoC design paradigm and identifies error tolerance as a design challenge.

Published by the IEEE Computer Society
0018-9162/04/$20.00 © 2004 IEEE

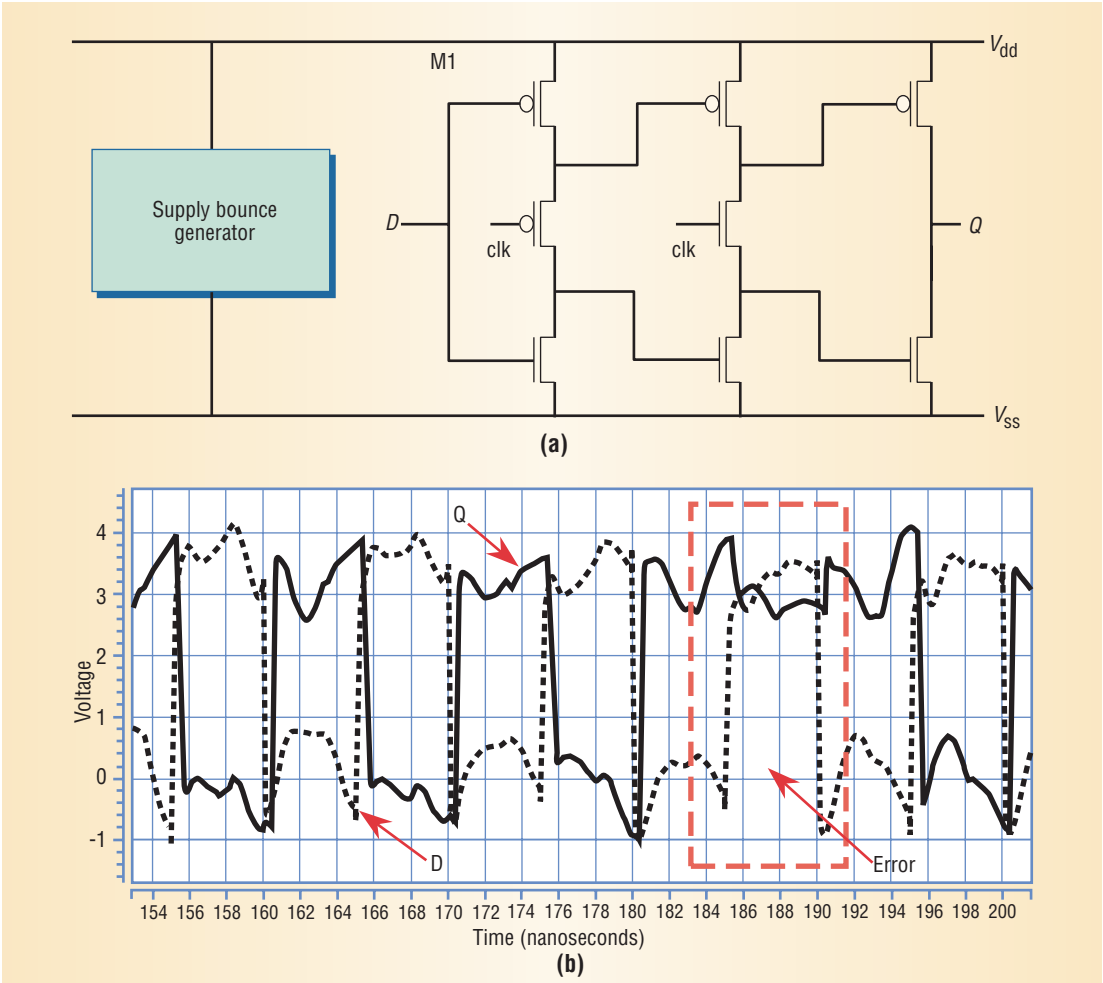

Figure 1. Noise due to supply bounce. (a) D-register circuit schematic. (b) Input and output waveforms. Q and D have the same polarity in the dotted box, indicating a logic error.

DSM NOISE

In DSM circuits, noise is any disturbance that drives node voltages or currents away from a nominal value, causing permanent as well as intermittent errors.[7] If storage elements capture these errors, the result is an observable loss in functionality. Numerous masking mechanisms can prevent an error at the output of a logic gate from propagating further. Increased delay and accidental discharge/charge of dynamic nodes are common mechanisms for such failures.

Figure 1 illustrates an error at the output of an edge-triggered register. When input D is at logic 1 and an inductive kick raises the supply node above $V_{dd} + |V_{tp}|$, where V_{dd} is the supply voltage and V_{tp} is the positive-channel metal-oxide-semiconductor (PMOS) device threshold voltage, the topmost PMOS M1 in the input stage in Figure 1a will activate and cause the logic error indicated in Figure 1b.

The probability of this error event increases with complexity and with the reduction in the threshold voltage. A typical 0.13-µm CMOS process has device threshold voltages in the 200-300 millivolt range, and it is not unusual for the power supply grid to generate supply bounce of a few hundred millivolts. Thus, the types of errors shown in Figure 1 are not unusual.

Accurately modeling the numerous noise sources is difficult. *Noise mitigation* and *noise tolerance* are two distinct ways to handle DSM process noise. The EDA industry favors noise mitigation, which involves developing noise-analysis tools that identify hot spots and then having designers mitigate the impact of noise by focusing on those areas.

Although noise mitigation is an obvious solution to the reliability problem, it is fundamentally inefficient in terms of energy conservation. In contrast, noise tolerance, which is central to the communication-theoretic SoC design paradigm, requires designers to develop circuit and system design techniques that are inherently tolerant to noise and errors. When the design must achieve both energy efficiency and reliability, noise tolerance is the preferred approach.

RELIABLE SOC DESIGN TECHNIQUES

Present-day SoC design techniques are analogous to those used in the design of communications systems more than 50 years ago.

Claude Shannon[8] first proved the feasibility of reliable data transmission over noisy communication links in 1948. Subsequently, communications system designers mastered the science of develop-

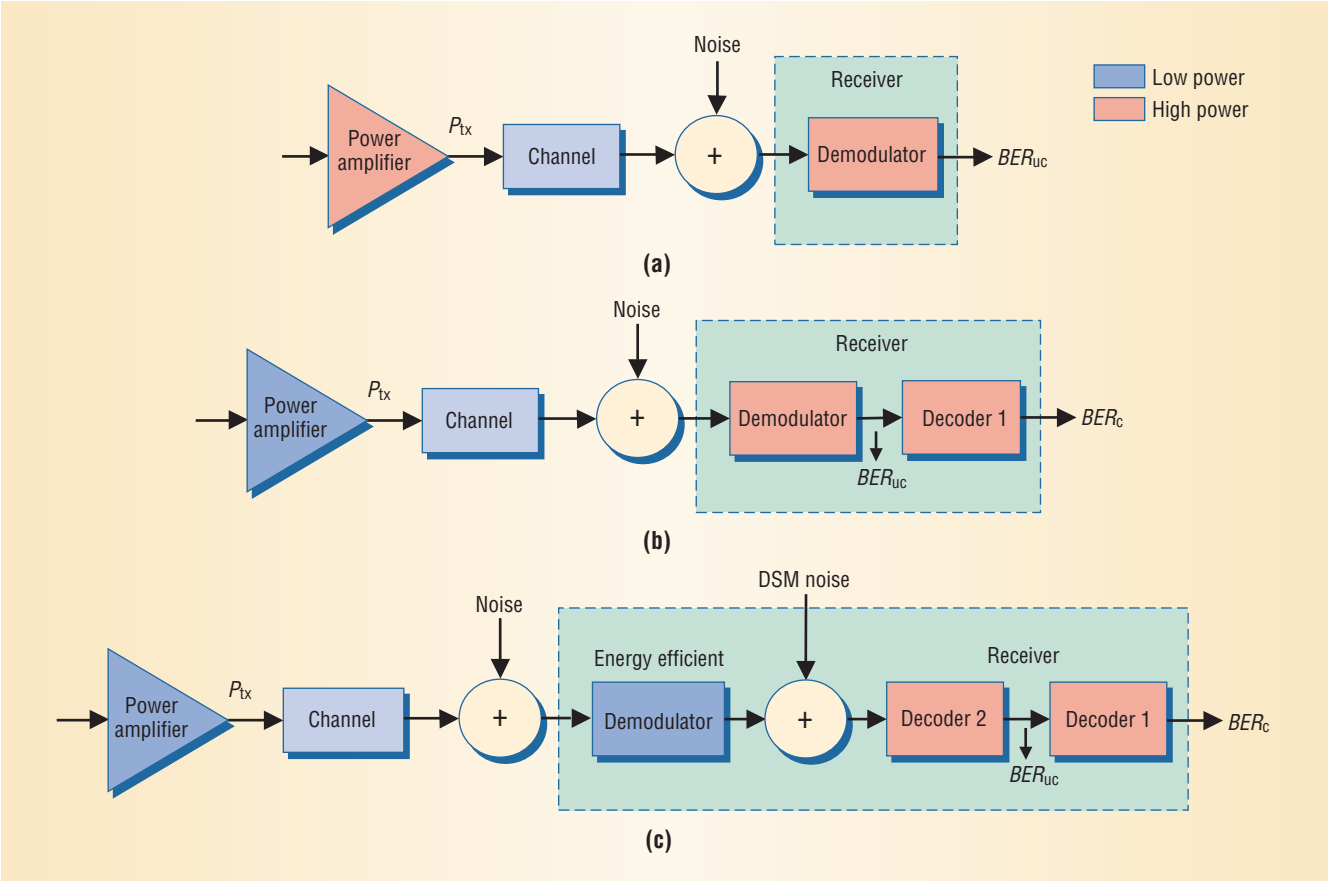

Figure 2. Communication-link design. (a) High transmit power (P_{tx}) will achieve a bit error rate (BER_{uc}) of 10^{-10} but is costly in transmit energy per bit. (b) Error-control coding makes it possible to achieve the required BER_c with low transmit power. (c) Adding error control for both channel and deep-submicron (DSM) noise provides reliability while reducing communication and computational power dissipation.

ing high-speed data transmission techniques that operate under a transmit power constraint and in the presence of noise.

SoC design must likewise be based on an integrated view of system-level reliability and energy efficiency. However, SoCs currently achieve reliability by generating signal power that exceeds noise, an approach that is extremely inefficient in terms of energy consumption.

Figure 2a illustrates a simple communication link, which consists of a power amplifier at the transmitter, the physical channel (copper, cable, optical fiber, or air), and a demodulator at the receiver; the transmitter also includes a modulator, not shown, that feeds into the power amplifier. Channel noise causes errors at the demodulator's output.

Reliable links must achieve a specific end-to-end bit error rate; a BER less than 10^{-10} is typical for data communication links. Pumping up the transmit power so that the received signal power overwhelms the channel noise power could easily achieve the required BER, but at the cost of high energy consumption as measured in terms of transmit energy per bit.

Figure 2b illustrates an alternative based on the design philosophy of correcting rather than avoiding errors. In this approach, a channel encoder at the transmitter, not shown, adds error-protection bits to the information-bearing data bits and then feeds the resulting redundant data stream into the modulator, which feeds into the power amplifier. This reduces the transmit power by many decibels.

The demodulator output is targeted to achieve a BER of roughly 10^{-3} to 10^{-4}, while the decoder output provides the required BER of 10^{-10}. In addition to meeting the BER reliability requirements, this approach is also significantly more energy efficient in terms of transmitted energy per bit. The question is whether designers can use a similar concept to reduce energy consumption in SoC designs, especially for computation.

Dramatically reducing the supply voltage would reduce speed and noise margins in the demodulator, increasing its energy efficiency. As Figure 2c shows, inserting a second decoder would correct the additional errors that this would cause. If the decoder is small relative to the demodulator, the result is a low-power link that significantly reduces both the communication and computational power dissipation, matching the demodulator's reliability with the reliability of the data it recovers.

NOISE TOLERANCE

Researchers[1-4,9] have shown that it is possible to compute reliably in a logic gate implemented in DSM process technology that is subject to random or unpredictable noise. By applying information theory to circuits, it is possible to determine the minimum energy needed to compute a task in the presence of noise.

Energy-Efficiency Bounds for DSM Circuits

Entropy, a key function in the application of information theory to circuits, is defined as

$$h(p) = -p\log_2(p) - (1-p)\log_2(1-p), \quad (1)$$

where $0 \le p \le 1$. The entropy function has the shape of an inverted parabola with minima of $h(0) = 0$ and $h(1) = 0$ and a maximum of $h(0.5) = 1$.

Assume that a two-input AND gate must process data at a rate of f_s bits per second from a data source that generates 1 and 0 with a probability of 1/2, subject to the following parameters: The nominal supply voltage V_{dd} is 1.5 volts, the noise standard deviation is σ_n, the load capacitance C_L is 50 femtofarads (fF), and the MOS devices' transconductance k_m is 200 μA/V^2.

To determine the minimum energy required to implement this gate given the data-rate requirements and noise and process parameters, it is necessary to first abstract out the impact of noise by a single parameter ε, the probability of the AND gate making an error. The relationship between ε and V_{dd} for a noise source that has a Gaussian distribution with a zero mean and standard deviation σ_n is given by

$$\varepsilon = Q\left(\frac{V_{dd}}{2\sigma_n}\right); Q(x) = \int_x^\infty \frac{1}{\sqrt{2\pi}} e^{-y^2/2}\,dy \quad (2)$$

where the noise source is assumed to appear at the gate input as noise voltage V_n. The gate makes an error if V_n exceeds the gate-decision threshold voltage $V_m = 0.5\ V_{dd}$. For example, if V_{dd} is equal to 1.5 volts and σ_n is equal to 500 millivolts, Equation 2 results in 0.067—on average, 67 errors will occur in 1,000 outputs. Note that ε is a monotonically decreasing function of V_{dd}. This trend is consistent with the notion that circuits' noise immunity increases with the supply voltage.

Employing the approximate relationship

$$f_c = \frac{k_m V_{dd}}{C_L}$$

and information-theoretic concepts, the following expressions can be used to determine dynamic power dissipation and the lower bound on supply voltage for reliable operation:

$$P_d = tC_L V_{dd}^2 f_c = h^{-1}\left[\frac{RC_L}{k_m V_{dd}} + h(\varepsilon)\right] V_{dd}^3 k_m; \quad (3)$$

$$V_{dd-min} = \frac{h(p_y) f_s C_L}{k_m\left[1 - h(\varepsilon)\right]}; \quad (4)$$

where $h^{-1}[\ ]$ is the inverse of the entropy function and $R = h(p_y)f_s$—p_y being the probability of observing a 1 at the output—is the information transfer rate. Equation 3 unveils a tradeoff between the transition activity t and the supply voltage: As the supply voltage decreases, the inverse entropy term representing t increases, thereby offsetting reduction due to the cubic term.

Figure A plots Equation 3 as a function of V_{dd} when σ_n is equal to 300 millivolts and R is 1 Gbit per second. Note that the supply voltage at which dynamic power dissipation is minimized (V_{dd-opt}) is greater than the minimum supply voltage for which a reliable implementation of the AND gate exists (V_{dd-min}). Reliability and energy efficiency are thus elegantly linked together. In Figure 3, V_{dd-opt} is equal to 0.8655 volts and V_{dd-min} is 0.6303 volts; the energy consumed when V_{dd} is equal to V_{dd-opt}—the lower bound on energy dissipation—is E_{b-min}, or 15.5 femtojoules per bit.

As equations 2 and 3 show, noise ε and the information transfer rate R play a key role in determining the lower bounds on power dissipation and the supply voltage. Past attempts at quantifying such lower bounds have ignored this dependence and thus are incomplete.

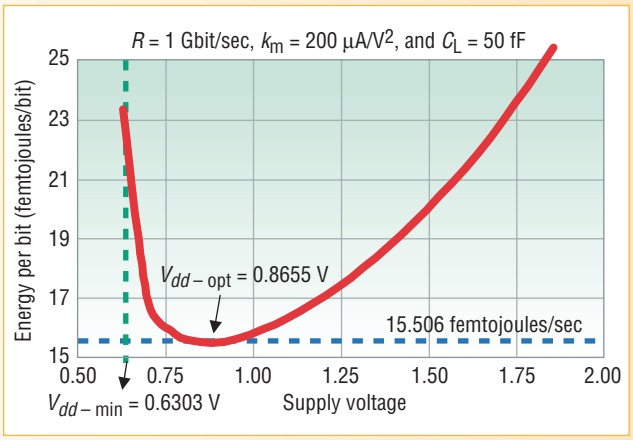

Figure A. Lower bounds on computational energy per bit for single output logic gate. The supply voltage at which dynamic power dissipation is minimized (V_{dd-opt}) is greater than the minimum supply voltage for which a reliable implementation of the AND gate exists (V_{dd-min}).

A noisy circuit with *information transfer capacity* C bits per second can reliably process a data source with an *information transfer rate* of R bits per second provided C is greater than R. The lower bound on energy efficiency is obtained when the capacity C is quantifiably close to R. The "Energy-Efficiency Bounds for DSM Circuits" sidebar describes this relationship in more detail. In the specific context of bus transmission, achievable energy efficiencies are a factor of 24 times below current-day systems.

Research has also shown that it is possible to compute reliably even when the signal and noise powers are comparable.[1-4,9] This implies that if computation is to occur near the limits of energy efficiency, noise tolerance is the correct design philosophy to optimize energy consumption while maintaining reliability.

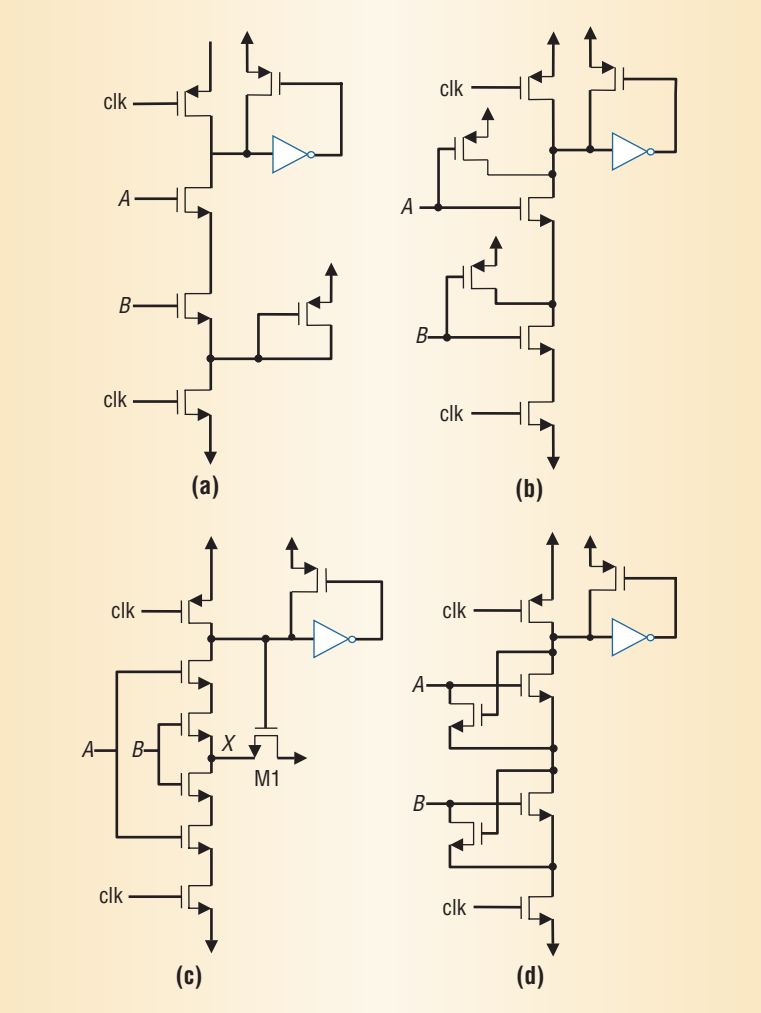

Figure 3. Noise-tolerant circuit techniques: (a) PMOS pull-up, (b) CMOS inverter, (c) mirror, and (d) twin-transistor.

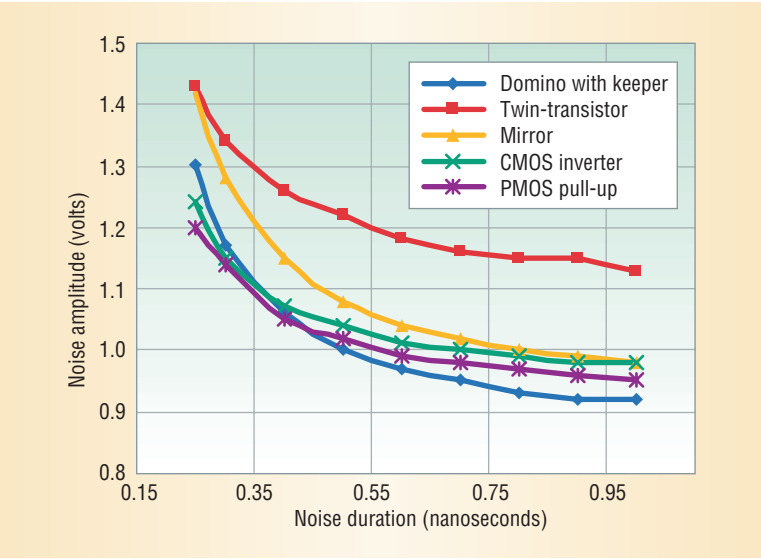

Figure 4. Noise-immunity curves for a two-input AND gate implemented using five dynamic circuit styles in a 0.18-μm, 1.8-V CMOS process. All of the circuit styles are designed to operate at a delay of 280 picoseconds when driving a load of 20 femtofarads.

Noise-tolerant circuit design

Dynamic circuits provide a convenient platform for studying the tradeoffs between energy efficiency and noise tolerance. Enhancing noise tolerance in a circuit carries an energy and power penalty or tax; effective techniques minimize this tax when they achieve a specified level of noise tolerance.

Figure 3 illustrates four dynamic circuit styles designed to provide noise tolerance: positive-channel MOS (PMOS) pull-up, complementary MOS (CMOS) inverter, mirror, and twin-transistor.

The *PMOS pull-up* technique[10] utilizes a pull-up device to increase the source potential of the negative-channel metal oxide semiconductor (NMOS) device, thereby increasing the transistor threshold voltage V_{tn} and hence the switching threshold voltage V_{sw} of the gate during the evaluation phase. This technique suffers from large static power dissipation.

The *CMOS inverter* technique[11] utilizes a PMOS transistor for each input, thereby adjusting V_{sw} to equal that of a static circuit. This technique cannot be used for NOR-type circuits as certain input combinations can generate a direct path from supply to ground.

The *mirror* technique utilizes two identical NMOS evaluation networks and one additional NMOS transistor M1 to pull up the source node of the upper NMOS network to $V_{dd} - V_{tn}$ during the precharge phase, thereby increasing V_{sw}. This technique guarantees zero DC power dissipation, but a speed penalty is incurred if the transistors are not resized.

The *twin-transistor* technique[5] represents the state of the art in noise-tolerant dynamic circuit design. It employs an extra transistor for every transistor in the pull-down network to pull up the source potential in a data-dependent manner. The twin-transistor technique consumes no DC power and has a limited impact on speed and power.

Noise immunity curves

Conservative static noise margin metrics do not account for the inherently low-pass nature of logic circuits, which can filter out noise pulses that either have a small amplitude or short duration. Thus, a comparison of circuit styles requires using metrics such as the noise immunity curves (NICs) shown in Figure 4.

A digital gate's NIC is a locus of tuples (V_n, T_n), where V_n is the noise amplitude in volts and T_n is the noise duration in seconds, representing noise pulses that generate a logic error for that gate. Noise pulses above the curve are guaranteed to gen-

erate logic errors, while those below are guaranteed not to do so.

A logic error is said to occur when the output crosses a predefined voltage threshold, usually halfway between the supply rails, or when the output glitch amplitude equals the input noise-pulse amplitude. The *average noise threshold energy*[5] is a convenient metric that can be derived from the NIC by averaging the energy of the noise pulses that cause an error. Normalizing the ANTE with the energy consumption provides the NANTE[5] metric, a measure of the noise-tolerance circuit technique's effectiveness.

Table 1 quantifies the ANTE, energy, and NANTE metrics for each of the four noise-tolerant techniques as well as a conventional domino design when implemented in a 0.18-µm, 1.8-volt CMOS process. The twin-transistor technique has the best ANTE and NANTE metrics, indicating that it provides the highest noise immunity per unit of energy consumption. Both the mirror and twin-transistor techniques have been proved experimentally in the past via the design and test of prototype chips in 0.35-µm CMOS technology.

More research is needed to develop dynamic circuit styles that are not only tolerant to various noise sources but also have low-noise-generation features. Circuit-level techniques are not sufficient, especially when energy efficiency also is a concern. Noise tolerance techniques are required at the architectural, algorithmic, and system levels of the design hierarchy.

Table 1. Noise tolerance versus energy efficiency.

Dynamic circuit technique	Average noise threshold energy (picojoules)	Energy (picojoules)	Normalized ANTE
Conventional domino	586.6	0.2688	2,182
Twin-transistor	859.1	0.2862	3,002
Mirror	633.4	0.3158	2,006
CMOS inverter	622.3	0.2752	2,261
PMOS pull-up	606.2	0.4826	1,256

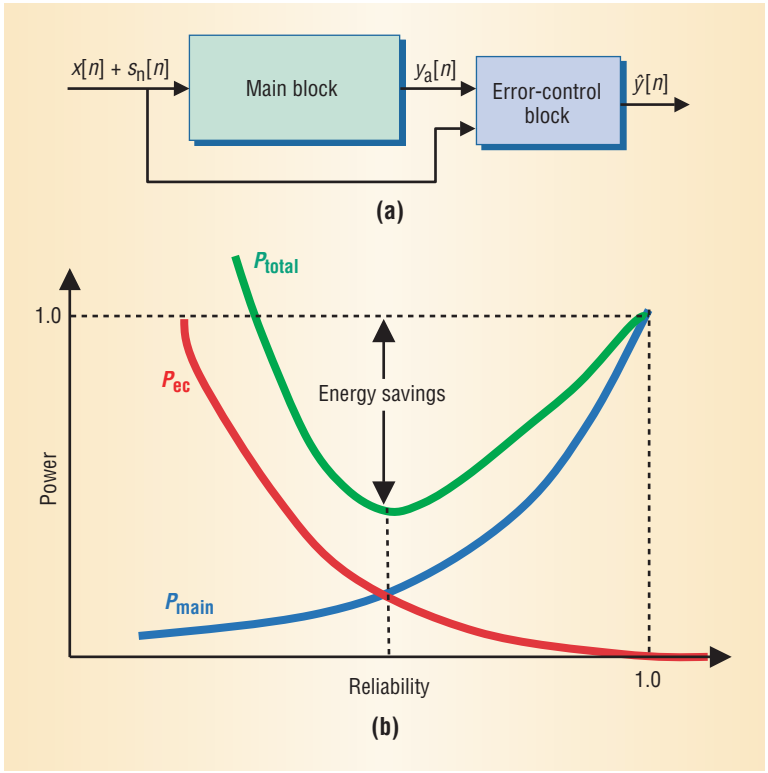

Figure 5. Algorithmic noise tolerance. (a) An ultra-energy-efficient main block executes most required computations; a reliable and thus energy-inefficient error-control block detects and corrects intermittent errors by the main block. (b) As the main block's energy efficiency increases, its reliability decreases, necessitating the use of increasingly complex error control.

ALGORITHMIC NOISE TOLERANCE

The key idea behind *algorithmic noise tolerance* (ANT),[6] illustrated in Figure 5a, is that an ultra-energy-efficient main block executes most of the required computations. The main block can make intermittent errors as long as they occur infrequently.

An *error-control block*, which is reliable and thus energy inefficient, detects and corrects these errors. As Figure 5b shows, as the main block's energy efficiency increases, its reliability decreases, necessitating the use of increasingly complex error control. Thus, the error-control block's power dissipation increases.

The total power dissipation of the main block and the error-control block reaches its minimum when the main block achieves a specific level of reliability, which is determined by the error frequency and the error control's effectiveness.

ANT techniques

Effective error-control techniques provide a high level of error detection and correction with low hardware complexity. Thus, error-control blocks can have a relaxed delay and power constraint.

Significantly relaxing delay and power constraints permits enhancement of noise immunity in the error-control block. For example, researchers can use the noise-tolerant twin-transistor technique[5] along with noise-analysis and noise-mitigation techniques to design a robust error-control block that has minimal impact on speed and power. Instead of focusing on the entire SoC, a designer need only ensure that the few error-control blocks in a complex design are robust at the circuit level.

Developers of error-control techniques for signal-processing kernels can use statistical performance metrics, such as the signal-to-noise ratio, to exploit the signals' statistical structure. In such

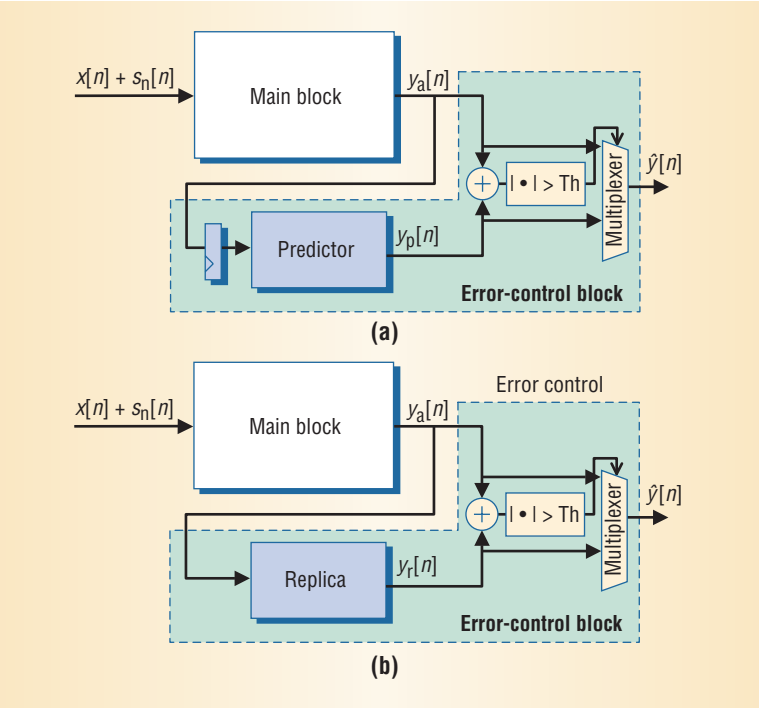

Figure 6. Algorithmic noise-tolerance techniques. (a) A predictor uses the past outputs to generate a statistical estimate of the main block. (b) Reduced-precision redundancy employs a replica of the main filter as an estimator.

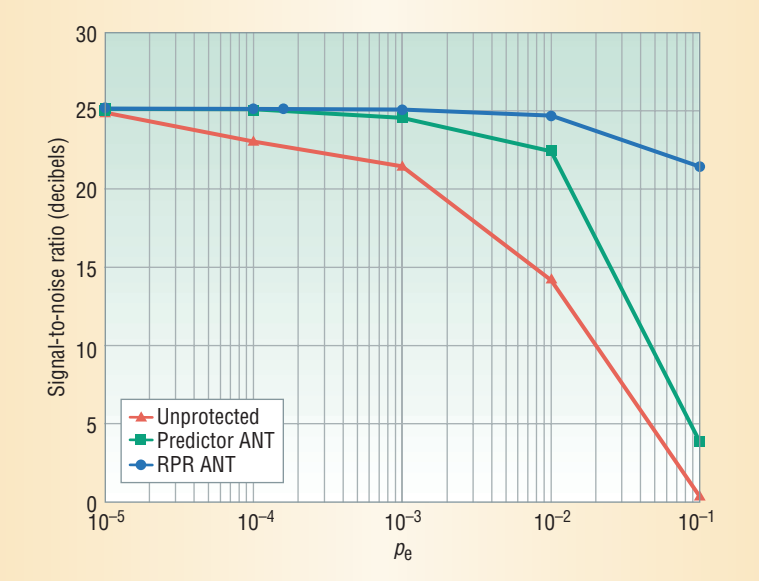

Figure 7. Predictor and reduced-precision redundancy ANT technique performance in the presence of random noise. The signal-to-noise ratio improves by 10 decibels even when each output bit of the filter is independently flipped at an average rate of once every 1,000 samples.

cases, designers can use signal-estimation techniques for error control.

The following examples all assume that the main block is a digital filter that makes intermittent errors.

In the *predictor* technique, shown in Figure 6a, another filter (predictor) uses the past outputs $y_a[n-1], y_a[n-2], \ldots, y_a[n-N_p]$ to generate a statistical estimate $y_p[n]$ of the main block. Designers can use standard statistical signal-processing techniques to determine optimal predictor coefficients that minimize the mean-squared error between the predictor and main filter outputs. The error-control block detects errors by comparing the predictor and main filter outputs. When the difference between the two outputs exceeds a prespecified threshold, the error-control block declares an error. In the event of an error, the error-control block selects the predictor output $y_p[n]$. Both error detection and correction are approximate, which is reasonable given the interest in maintaining the signal-to-noise ratio.

Figure 6b illustrates *reduced-precision redundancy*, another simple, yet effective, error-control technique. RPR employs a replica of the main filter as an estimator. The error-detection and -correction steps are the same as those in the predictor technique except that the prespecified threshold must be greater than the quantization noise floor.

Figure 7 shows predictor and RPR performance when the main filter output bits are randomly flipped with frequency p_e. The signal-to-noise ratio improves by 10 decibels even when each output bit of the filter is independently flipped at an average rate of once every 1,000 samples.

Voltage overscaling

One way to study the tradeoff between energy efficiency and reliability is to reduce the supply voltage below the minimum required for correct operation. This *voltage overscaling*[6] results in delay violations, causing output errors whenever the user applies an appropriate input sequence. VOS, which has recently been used in the design of reliable low-power microprocessors,[12] improves energy efficiency beyond what present-day supply scaling can achieve.

Although VOS errors are systematic, they are modeled as being random when using ANT techniques, which are known to be effective for random errors. In addition, error-control techniques that exploit the systematic nature of VOS errors are too complex to be of any practical use.

A voltage-overscaled digital filter chip incorporating the predictor technique and implemented in a 0.35-µm CMOS process has demonstrated up to 70 percent savings in energy over a filter operating at critical supply voltage.

ANT-based techniques can be used instead of triple-modular redundancy (TMR) to provide robustness against soft errors due to particle hits. Much greater energy efficiencies are achievable by employing one main filter block and two estimators. In contrast, TMR would require three main

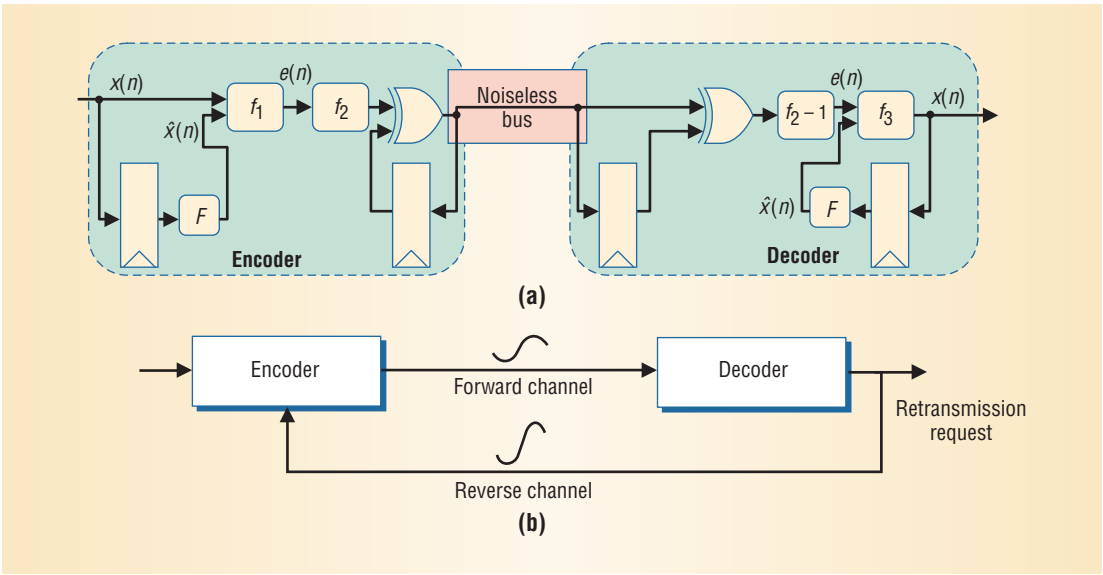

Figure 8. Bus coding. (a) A bus-coding framework based on source coding consists of a predictor F, differentiator f_1, and mapper f_2. (b). Use of Hamming and Reed-Muller codes along with a reverse retransmission request channel reduce signaling levels in the frequently used forward channel.

filter blocks, resulting in very high overhead. In such a case, the assumption of an error-free error-control block can be relaxed.

NOISE-TOLERANT BUS TRANSMISSION

Buses are key SoC components. Coupling between adjacent wires, supply bounce at the receiver, and other sources can cause noise in buses. Energy consumption in buses occurs mainly due to transitions on bus lines, including charging and discharging the self and coupling capacitances.

Early bus power-reduction techniques ignored coupling, focusing instead on reducing transition activity in individual bit lines. As Figure 8a shows, a bus-coding framework[13] based on source coding—for example, video compression employed in multimedia communication networks—has three key elements. The predictor F can be an identity or an increment function; the differentiator f_1 can be an XOR or a subtractor; and the mapper f_2 can employ a probability-based mapping, value-based mapping, inversion, or identity function.

Assigning different functionalities to F, f_1, and f_2 results in a family of coding schemes demonstrating this framework's power. For example, the following assignment can derive the well-known *bus-invert* scheme from this framework: F = identity, f_1 = XOR, and f_2 = inversion. Another useful coding scheme for address buses is *INC-XOR*, obtained by the following assignment: F = increment, f_1 = XOR, and f_2 = identity.

These and other similar techniques that focus on reducing transition activity in individual bus lines ignore the problem of coupling found in DSM processes. Recent work uses coding to minimize transitions on adjacent bus lines, thereby reducing delay. However, none of these techniques address the noise problem.

As Figure 8b shows, the first work addressing noise and energy efficiency in high-speed SoC bus transmission[2] used Hamming and Reed-Muller codes along with a reverse retransmission request channel to reduce signaling levels in the frequently used forward channel. Simulations demonstrated a three- to fourfold power savings, but again this work does not address coupling. Use of error-detection and retransmission was recently proposed for reliable communications in networks- on-a-chip.[14]

A remaining challenge is to develop noise-tolerant bus transmission codes that jointly address coupling, self-capacitance, and noise.

The semiconductor industry faces numerous challenges in developing reliable, energy-efficient SoC designs that are on a par with modern communications systems. Researchers must explore the tradeoffs between reliability and energy efficiency at the device, circuit, architectural, algorithmic, and system levels to develop a reliability-energy "knob" that can be synergistically tuned to meet these requirements at each level of the design hierarchy. Elegant and practical solutions will require the application of coding and communication-theoretic techniques to the design of SoC components. In addition, researchers must develop statistical approaches to design and verification as well as statistical performance metrics. ■

Acknowledgments

This work was supported by the MARCO-sponsored Gigascale Silicon Research Center and National Science Foundation grants CCR-000987 and CCR-0095929. I also thank Krisnamurthy Soumyanath, Ram Krishnamurthy, Tanay Karnik, and Shekhar Borkar at Intel as well as current and former graduate students Sumant Ramprasad, Raj Hegde, Lei Wang, Byonghyo Shim, and Ganesh Balamurugan for their contributions.

References

1. N.R. Shanbhag, "A Mathematical Basis for Power-Reduction in Digital VLSI Systems," *IEEE Trans. Circuits and Systems,* Part II, vol. 44, no. 11, 1997, pp. 935-951.
2. R. Hegde and N.R. Shanbhag, "Towards Achieving Energy Efficiency in Presence of Deep Submicron Noise," *IEEE Trans. VLSI Systems,* vol. 8, no. 4, 2000, pp. 379-391.
3. P.P. Sotiriadis, V. Tarokh, and A. Chandrakasan, "Energy Reduction in VLSI Computation Modules: An Information-Theoretic Approach," *IEEE Trans. Information Theory,* vol. 49, no. 4, 2003, pp. 790-808.
4. P. Abshire and A.G. Andreou, "Capacity and Energy Cost of Information in Biological and Silicon Photoreceptors," *Proc. IEEE,* vol. 89, no. 7, 2001, pp. 1052-1064.
5. G. Balamurugan and N.R. Shanbhag, "The Twin-Transistor Noise-Tolerant Dynamic Circuit Technique," *IEEE J. Solid-State Circuits,* vol. 36, no. 2, 2001, pp. 273-280.
6. R. Hegde and N.R. Shanbhag, "Soft Digital Signal Processing," *IEEE Trans. VLSI Systems,* vol. 9, no. 6, 2001, pp. 813-823.
7. K.L. Shepard and V. Narayanan, "Conquering Noise in Deep-Submicron Digital ICs," *IEEE Design and Test of Computers,* vol. 15, no. 1, 1998, pp. 51-62.
8. C.E. Shannon, "A Mathematical Theory of Communications," *Bell System Technical J.,* vol. 27, nos. 3 and 4, 1948, pp. 379-423, 623-656.
9. S. Winograd and J.D. Cowan, *Reliable Computation in the Presence of Noise,* MIT Press, 1963.
10. G.P. D'Souza, *Dynamic Logic Circuit with Reduced Charge Leakage*, US patent 5483181, Patent and Trademark Office, 1996.
11. J.J. Covino, *Dynamic CMOS Circuits with Noise-Immunity*, US patent 5650733, Patent and Trademark Office, 1997.
12. T.M. Austin, "DIVA: A Reliable Substrate for Deep Submicron Microarchitecture Design," *Proc. 32nd Ann. Int'l Symp. Microarchitecture,* IEEE CS Press, 1999, pp. 196-207.
13. S. Ramprasad, N.R. Shanbhag, and I.N. Hajj, "A Coding Framework for Low-Power Address and Data Busses," *IEEE Trans. VLSI Systems,* vol. 7, no. 2, 1999, pp. 212-221.
14. D. Bertozzi, L. Benini, and G. De Micheli, "Low Power Error Resilient Encoding for On-Chip Data Busses," *Proc. 2002 Design, Automation and Test in Europe Conf. and Exposition,* IEEE CS Press, pp. 102-109.

Naresh R. Shanbhag *is an associate professor in the Department of Electrical and Computer Engineering and the Coordinated Science Laboratory at the University of Illinois at Urbana-Champaign. His research focuses on the design and implementation of digital integrated circuits and systems for signal processing and communications. Shanbhag received a PhD in electrical engineering from the University of Minnesota. He is a member of the IEEE Solid-State Circuits Society, IEEE Circuits and Systems Society, and IEEE Signal Processing Society. Contact him at shanbhag@uiuc.edu.*