

Exploring Very Low-Energy Logic: A Case Study

L. P. Alarcón, T.-T. Liu, M. D. Pierson, and J. M. Rabaey*

Berkeley Wireless Research Center, University of California, Berkeley, CA 94704, USA

(Received: 2 October 2007; Accepted: xx Xxxx xxxx)

This paper shows leakage as a limit to the effectiveness of voltage scaling as a means of reducing the energy per operation in a digital circuit. Methods of decreasing operational or dynamic leakage are then discussed. The design and simulation results of a sense amplifier-based pass transistor logic (SAPTL) circuit topology as a low leakage and low energy alternative is presented and then compared to standard static 90-nm CMOS implementations.

Keywords: 90-nm CMOS, Pass Transistors, Stack Effect, Leakage Current, Self-Timed Circuits, Body Bias.

1. INTRODUCTION

The continued scaling of transistor feature sizes leads to an increase in integration density, which brings about a corresponding increase in compute density.¹ This scaling also results in an increase in overall circuit power consumption and this increase in power that accompanies this scaling trend is preventing us from truly harnessing the benefits of decreasing transistor feature sizes.² For applications that are severely energy limited, such as those using implantable electronics, the energy per operation must continue to decrease, allowing for years of battery life at relatively low operating frequencies and power levels.^{3,4} The best way to reduce the energy per operation is to reduce the supply voltage, V_{dd} .

However, as V_{dd} is scaled down further, leakage energy increases due to increased delay, Δt_{op} , and reduced activity factors, α , resulting in higher overall energy per operation as seen in Figure 1 and can be expressed as:

$$E_{total} = V_{dd}(\alpha C_{sw} V_{dd} + I_{leak} \Delta t_{op}) \quad (1)$$

Therefore, for a certain logic operation and circuit topology, a minimum energy point exists at some optimal V_{dd} . For most static CMOS circuits, this optimum V_{dd} is less than the transistor threshold voltage, V_{th} , placing this minimum energy point in the subthreshold region of operation. In this region of operation, which comes with a high performance penalty, the impact of variability and process variations are more pronounced due to the exponential behavior of the transistors. Thus, in this light, it is more desirable to operate transistors in the superthreshold regime.

* Author to whom correspondence should be addressed.
 Email: jan@eecs.berkeley.edu

The introduction of high- K gate dielectrics has led to significantly reduced gate leakage current, I_g , thus, only subthreshold source-to-drain leakage current, I_{off} is considered in this paper.⁵ In order to reduce I_{off} , threshold voltages are not scaling as fast as V_{dd} .² This asymmetric scaling of V_{th} with respect to V_{dd} however decreases the performance by reducing the available drive current of the transistor.⁶ In transistors with channel lengths below 50 nm, band-to-band tunneling (BTBT) can dominate I_{off} , negating the effect of using relatively higher threshold voltages.⁷

A review of several low leakage techniques and how they affect the total energy of a CMOS gate is conducted in Section 2, then Section 3 presents pass transistor logic (PTL) as a good alternative to low leakage logic design. Section 4 introduces the sense amplifier-based pass transistor logic (SAPTL) topology as a low leakage circuit alternative that allows continued energy reduction through voltage scaling even in the presence of leakage and discusses the organization and synchronous timing operation of the SAPTL. Section 5 shows the simulated leakage, energy and delay characteristics of the SAPTL. An asynchronous timing scheme for the SAPTL is discussed in Section 6. Future directions are presented and conclusions are then drawn at the end of the paper.

2. LOW LEAKAGE TECHNIQUES

The MOS transistor subthreshold leakage current, including the effect of drain-induced barrier lowering (DIBL), can be expressed as:⁸

$$I_{D, off} = \frac{\mu_{eff} C_{ox}}{L_g} (m_2 - 1) \left(\frac{kT}{q} \right)^2 e^{q(V_g - V_{th, sat} + \lambda_D V_{DS}) / m_2 kT} \times (1 - e^{-qV_{DS} / kT}) \quad (2)$$

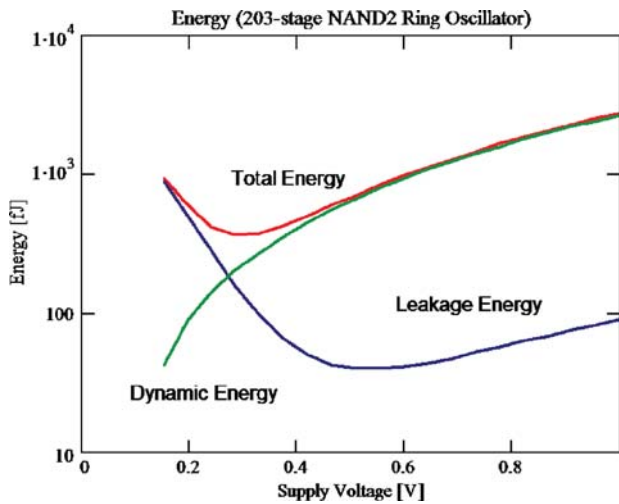


Fig. 1. The total energy of a 203-stage NAND2 ring oscillator as a function of supply voltage. Also shown are its switching and leakage components.

where m_2 is given by:

$$m_2 = 1 + \frac{1}{C_{ox}} \sqrt{\frac{\epsilon_{Si} q N_a}{4\phi_B}} \quad (3)$$

Subthreshold leakage current can thus be reduced (1) by increasing the channel length, L_g , (2) by reducing the supply voltage and consequently the drain-to-source voltage, V_{ds} or (3) by increasing the threshold voltage, $V_{th,sat}$.

A sampling of currently used low leakage circuit techniques aside from voltage scaling include the use of (1) non-minimum channel lengths; (2) stacked transistors and

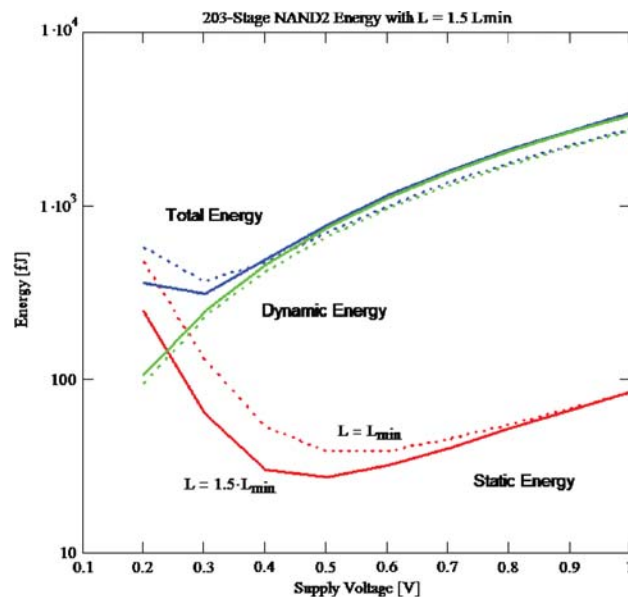


Fig. 2. The effect of increasing the channel length on leakage energy and its corresponding impact on the total energy per cycle at very low supply voltages. Here, $L_g = 1.5L_{min}$.

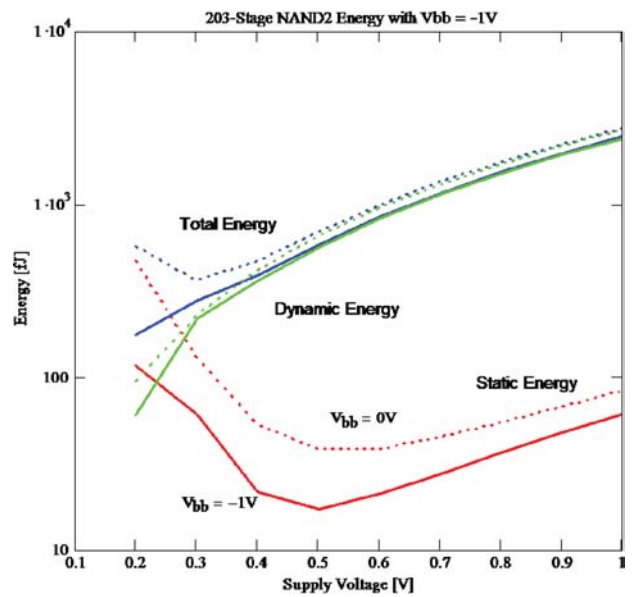


Fig. 3. The effect of reverse body bias on leakage energy and its corresponding impact on the total energy per cycle at very low supply voltages. Here, $V_{bb} = -1$ V.

(3) various header and footer switch topologies and are summarized in Refs. [9] and [2].

Since the focus of this paper is on dynamic or operational leakage reduction, standby leakage reduction techniques such as header and footer switches are not considered. In order to illustrate the effect of leakage on the energy per cycle or per operation, these two low leakage techniques are applied to a 203-stage NAND2 ring oscillator in 90 nm CMOS. Figure 2 shows the switching and leakage energy components as a function of the supply

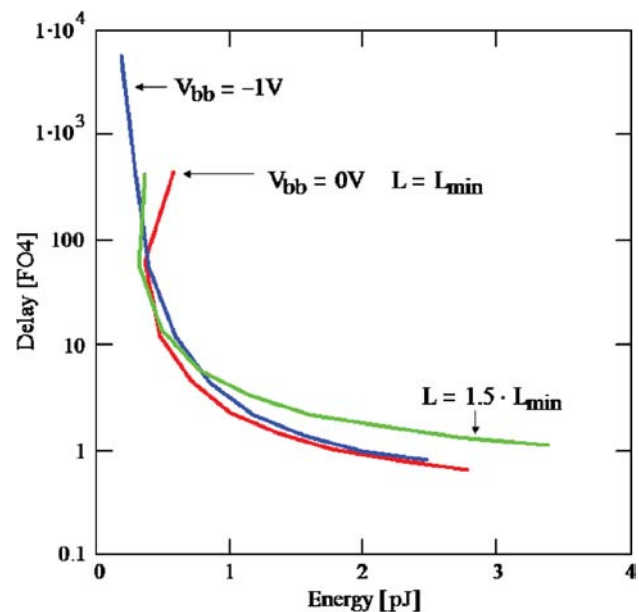


Fig. 4. The energy-delay tradeoffs for a 203-stage NAND ring oscillator for the reference circuit ($L = L_{min}$ and $V_{bb} = 0$ V) and the effect of increasing the channel length and applying reverse body bias.




Fig. 5. Using more complex gates reduces the number of leakage paths at the cost of increasing the effective resistance of the path from the supply rails.

voltage when the channel length is increased to 1.5 times the minimum channel length. It can be seen that the total energy at low voltages is reduced as a result of reduced leakage energy. This is again due to the fact that leakage energy dominates at very low voltages.

The same effect is observed when reverse body bias is used. Figure 3 shows the energy as a function of supply voltage for V_{bb} of -1 V. Again, the total energy at very low voltages is reduced.

The impact of these low leakage circuit techniques on circuit delay can be best seen on an energy-delay diagram, shown in Figure 4. This figure shows how much delay is incurred to achieve a certain reduction energy. It can be seen that by reducing the operating leakage energy, lowering the supply voltage can continue to reduce the total circuit energy at the expense of increased circuit delay.

Increasing the transistor channel length reduces leakage by increasing the effective resistance of the leakage path between the supply rails, resulting in delay penalties. Using reverse body bias to increase V_{th} also leads to the same results, however due to increased channel doping,¹⁰ BTBT limits the effectiveness of body bias to reduce leakage, and hence the effective power reduction is reduced by approximately $4\times$ per technology generation.¹¹

Another way to reduce leakage current is to use more complex gates as seen in Figure 5. By combining more functionality in a single gate, less gates are used resulting in less leakage paths from V_{dd} to ground. Again, this would increase the effective resistance of the leakage path from the supply rails.

3. PASS TRANSISTOR LOGIC

One circuit alternative is to use pass transistor networks to reduce leakage current. Pass transistor logic is a simple and compact circuit topology and in some cases, outperforms static CMOS circuits.¹²⁻¹⁴ The pass transistor network itself does not have V_{dd} and ground connections, thus drastically reducing the number of leakage paths as shown in Figure 6. In pass transistor logic (PTL), leakage is confined to the driving and level restoring circuitry associated with the pass transistor network. These circuits are used to recover the voltage swing and delay degradation inherent in PTL circuits. Figure 7 shows a conventional pass transistor network that implements logic functions

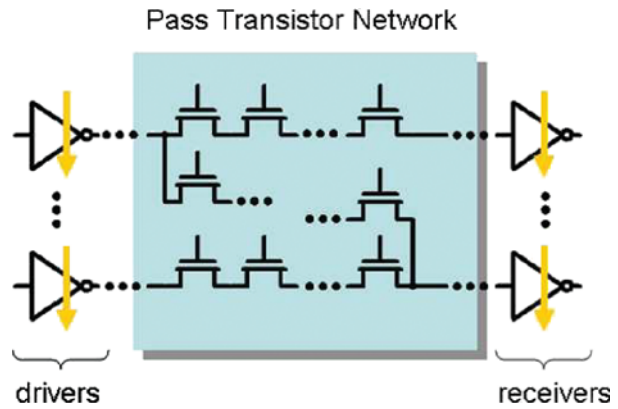


Fig. 6. A generic pass transistor network showing no subthreshold leakage paths in the logic paths. To first order, all leakage is confined to the peripheral circuitry used to recover signal swing and delay.

based on multiplexer or binary decision diagram (BDD) tree structures. The main drawback of these type of tree structures is that sneak paths exist allowing leakage current to flow. Pass transistor networks can be made more complex, thus reducing the total number of drivers and level restorers in order to reduce the number of leakage paths, but unfortunately, the number of sneak paths in the pass transistor tree increases exponentially with the number of logic inputs, i.e., 2^N sneak paths for N levels. Note that the delay is also dependent on the number of levels and is proportional to N^2 .

Pass transistors also increase the effective channel length (and thus resistance of the leakage path) between the supply rails. However, PTL has the potential to offer more computational density for a given leakage path resistance than simply increasing the transistor channel length.

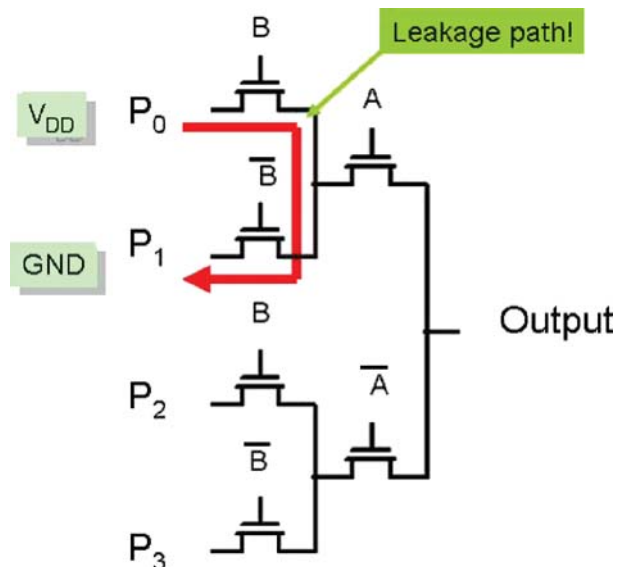


Fig. 7. A conventional pass transistor tree showing sneak leakage paths. Note that there are 2^N sneak leakage paths for a fully decoded tree of depth N .

In order to overcome the limitations of conventional pass transistor logic, a sense amplifier-based pass transistor logic (SAPTL) family is offered as an alternative low leakage and low energy circuit topology.

4. THE SAPTL

The basic organization of the SAPTL circuit is shown in Figure 8. It consists of

- (1) the pass transistor tree, called the *stack*, which computes the required logic function;
- (2) a root node driver that injects signals into the *stack* and
- (3) a sense amplifier that is used to recover both voltage swing and performance.

4.1. The Stack

To mitigate the limitations of conventional multiplexer-based pass transistor trees due to sneak paths, and recognizing that pass transistors are inherently bidirectional circuits, an inverted pass transistor tree, which is referred to as the *stack* is utilized, and shown in Figure 9. The *stack* still has no supply rail connections and has predictable delay paths, and in addition, has pseudo-differential outputs, where a signal or current is present in either S or \bar{S} , but not both at the same time.

The input capacitances of the *stack* can be made equal by making the transistors closer to the root input larger. This also has the effect of decreasing the delay of the *stack*, by reducing the resistance of the signal path near the root of the tree.¹⁵

Since the input can only propagate from the root of the *stack* to the output, there are no sneak paths that exist, and thus to first order, reducing V_{th} to near zero is possible. Also, this reduction in threshold voltage also reduces the resistance, and thus, the propagation delay from the root of the *stack* to the outputs S and \bar{S} without any corresponding increase in leakage current drawn from the supply rails. The absence of sneak paths also allows the construction of deeper and more complex stacks, again without an increase in supply rail leakage. This V_{th} reduction and complexity increase, however, imposes stricter




Fig. 8. The SAPTL organization showing (1) the pass transistor tree network or the *stack*, (2) the root node driver and (3) the sense amplifier.

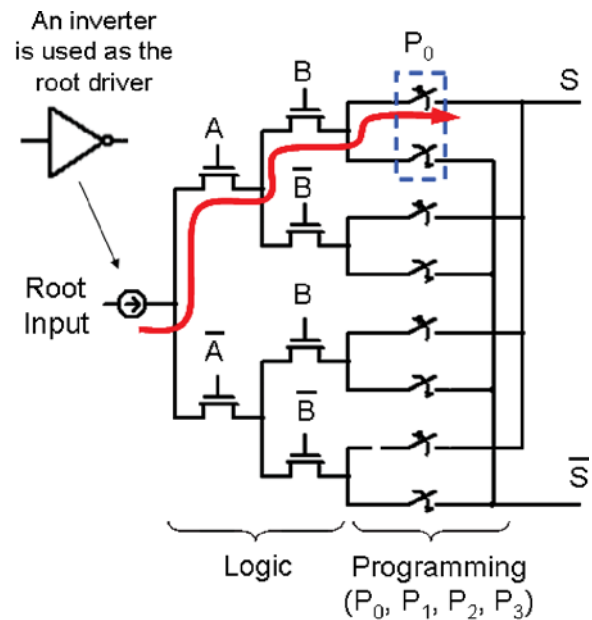


Fig. 9. The *stack* - an inverted pass transistor tree wherein a drive signal is injected into the root of the tree and observed pseudo-differentially at the outputs S and \bar{S} . Note that sneak leakage paths do not exist since the signal can only flow from the root to the outputs.

input resolution requirements on the sense amplifier, due to the lower I_{on}/I_{off} ratio at its inputs.

Each path from the root node to the output of the *stack* represents a minterm of a logic function, thus to program the *stack*, each branch representing the minterms contained in the desired logic function to be implemented is connected to the output S and each maxterm is connected to \bar{S} . Figure 10 shows how a 2-input *stack* can be configured to generate a boolean function of two variables. In this paper, the depth of the *stack*, N_{stack} , is defined as the number of transistors in series from the root node to the output, and due to the nature of the *stack*, it is the same for every path. Note that the input capacitance of the

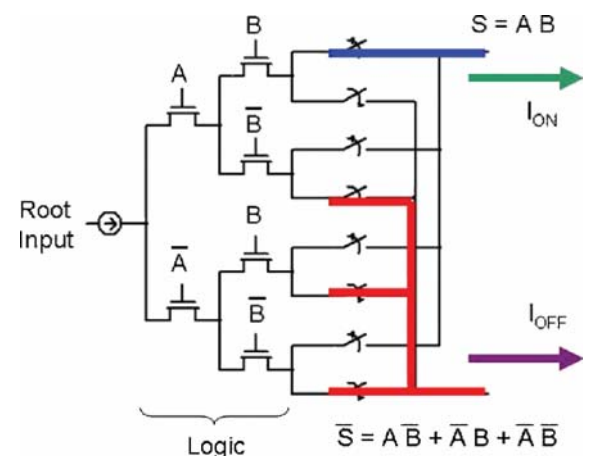


Fig. 10. The *stack* showing the programming for the boolean function $S = AB$. When the function is true, current is observed at the S output or at \bar{S} when the function is false.




Fig. 11. The sense amplifier used to recover voltage swing and delay lost due to the pass transistor tree or the stack.

stack, C_{in} is proportional to $2^{N_{stack}}$. In this paper, an inverter is used as the root driver.

4.2. The Sense Amplifier

The maximum voltage that can appear at either the S output or the \bar{S} output of the stack is $V_{dd} - V_{th}$, due to the voltage drop needed by the first transistor in the chain to maintain drive current flow. The time it takes for this degraded signal to reach the output is determined by the stack, and can be modeled using Elmore's delay equation¹⁵ and is strongly dependent on N_{stack} .

In order to (1) recover this voltage degradation; (2) improve the performance of the SAPTL and (3) provide sufficient buffering in order to drive a reasonable load capacitance, a sense amplifier (SA), shown in Figure 11, is added at the outputs of the stack. The SA consists of an input stage that acts as a preamplifier and a cross-coupled latch.

Other sense amplifier topologies, such as current-sense amplifiers, were evaluated. However, due to the low I_{on}/I_{off} ratios at the output of the stack, a two-stage topology is

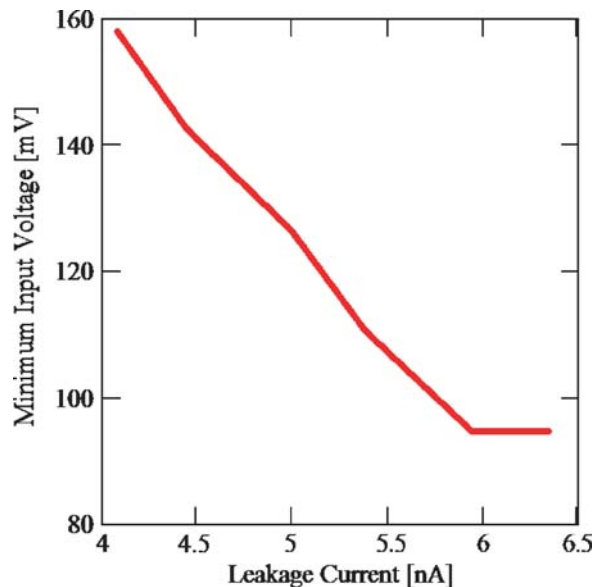


Fig. 12. The sense amplifier leakage versus minimum input voltage tradeoff in the presence of width, length and threshold voltage mismatch.

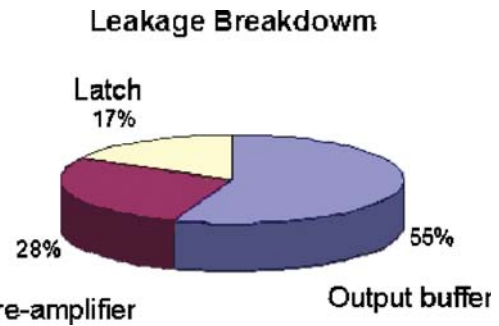


Fig. 13. The leakage breakdown by component of the sense amplifier, showing the significant leakage share of the output buffers.

used. Using this two-stage topology provides a sufficient transconductance gain in order to reduce the effect of the offset voltage of the latch due to transistor mismatch and process variations, allowing the detection of smaller voltage differences across S and \bar{S} without using very large transistors. Thus, this increased input sensitivity enables the sense amplifier to trigger earlier, decreasing the overall propagation delay. The sense amplifier is kept in a precharged state before triggering and its timing is determined by the CK input. Separating the SA timing from its functionality and using a precharge cycle reduces the required gain and supply voltage constraints, allowing the use supply voltages as low as 300 mV. Another advantage of this sense-amplifier topology is that it allows the retention of the latched data even when both of its inputs are set to zero. This is an important component of the SAPTL timing as discussed in the next subsection.

The pseudo-differential outputs of the stack matches well with the differential SA, making it more robust in the presence of noise and other interfering common-mode signals for better detection of small differential voltages. Since the sense amplifier is the SAPTL's main source of leakage current, SA sensitivity and speed is traded off to balance the amount of area, power and leakage current that it consumes.

Figure 12 shows how the sense amplifier leakage current is traded off for input sensitivity allowing for width, length and threshold voltage mismatch. In this design, the SA transistors are sized such that an input differential

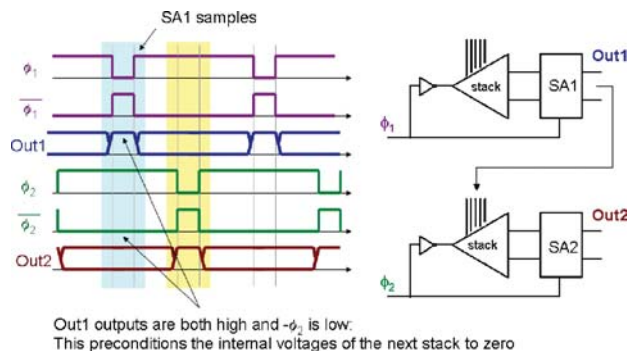


Fig. 14. The timing diagram for the SAPTL using a two-phase non-overlapping clocking scheme.




Fig. 15. The layout of the SAPTL in a 90 nm CMOS process with $N_{stack} = 5$. Note the regularity of the *stack* structure.

voltage of $V_{dd}/3$ can be resolved reliably. The output buffers account for a significant amount of the overall sense amplifier leakage as seen in Figure 13. Thus, leakage can also be reduced by reducing the output drive strength of the sense amplifier at the expense of increased delay.

4.3. Synchronous Timing

One approach to providing timing information to the SAPTL is by using global two-phase non-overlapping clock signals. The timing diagram for this fully synchronous operation is shown in Figure 14.

Due to the possibility of charge build-up within the non-energized *stack* paths, two-phase clocking is used in order to precondition all the internal nodes of the *stack* to ground prior to applying the root node drive signal. A *stack* can be preconditioned by setting all the inputs to the *stack* to V_{dd} and the root node to ground, thus forcing all nodes to be discharged. This ensures that there are no unwanted charge sharing events that occur inside the *stack* during the evaluation phase that could possibly cause the sense amplifier to make an incorrect decision.

Taking advantage of the fact that the outputs of the sense amplifier are both at V_{dd} during their precharge cycle, the succeeding SAPTLs being driven by this SA will be preconditioned if they are on a opposite clock domains.

4.4. Implementation

The layout of the SAPTL in a 90 nm CMOS process with $N_{stack} = 5$ is shown in Figure 15. The regularity of the pass transistor array results in increased robustness in the presence of variability.¹⁶ The area of the SAPTL is approximately 38% less than the area occupied by the equivalent CMOS LUT implemented using standard cells. This is due to the fact that the *stack* does not use PMOS transistors, and the fact that automated place and route tools consume 20%–40% more area than handcrafted designs.

5. SIMULATION RESULTS

Figure 16 shows the simulated transient behavior of a 5-input SAPTL using a commercial 90 nm CMOS process.

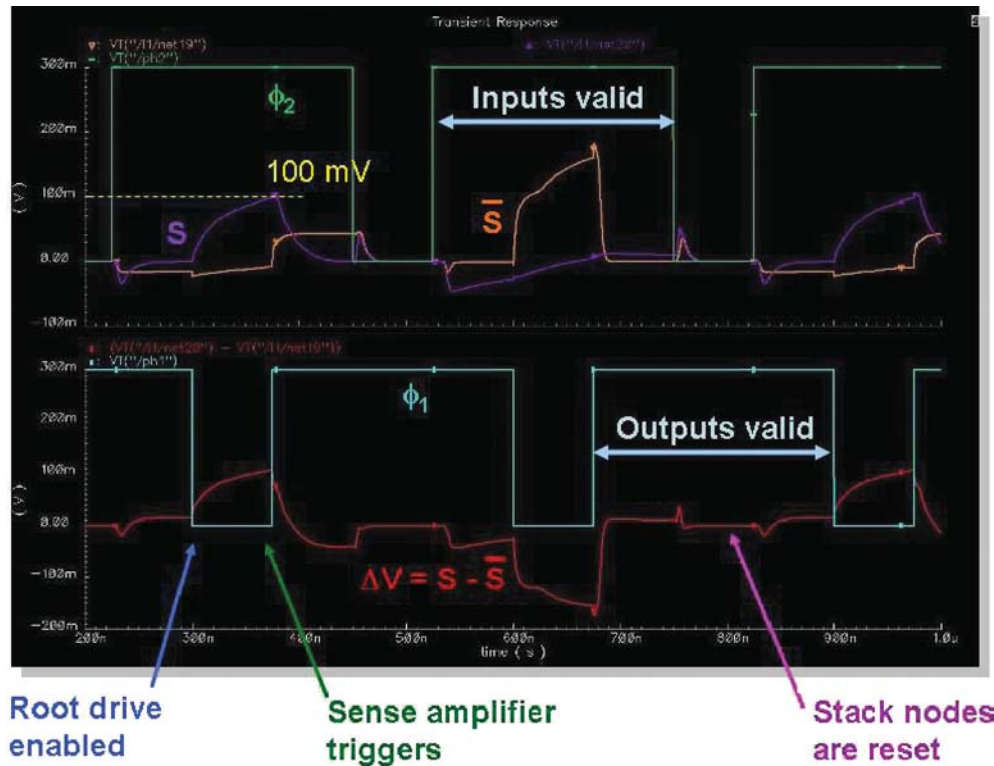


Fig. 16. The transient simulation of a 5-input SAPTL using two-phase non-overlapping clocks for a supply voltage of 300 mV.

The signals are referenced to the edges of the 2-phase non-overlapping clock, showing the precondition or precharge phase and the output of the *stack* for a supply voltage of 300 mV. The sense amplifier is triggered when the voltage at either *S* or \bar{S} reaches a third of the supply voltage, or in this case, 100 mV.

5.1. Leakage Simulations

The synchronous SAPTL is compared to a CMOS LUT and a transmission-gate (TG) LUT with the same number of inputs. The circuits are designed using a commercial 90-nm CMOS process and simulated using the Spectre circuit simulator. The TG LUT is designed via schematic entry while the CMOS LUT is described using VHDL then synthesized using the Synopsys Design Compiler. The Astro place and route tool is then used for placement and routing, utilizing a commercial 90-nm standard cell library.

Figure 17 shows the leakage behavior of the SAPTL5 (SAPTL with $N_{stack} = 5$), the CMOS LUT5 and the TG LUT5 as a function of supply voltage. At $V_{dd} = 300$ mV, the CMOS LUT and the TG LUT leakage current is 30× and 4.7× greater than the synchronous SAPTL respectively, while at $V_{dd} = 1$ V, it is 42× and 11× larger respectively. This relatively small SAPTL leakage is due to its significantly reduced supply rail connections, and thus its leakage current is dominated mainly by the sense amplifier. Note that this does not take into account the leakage from the implied clock distribution network associated with the SAPTL.

5.2. Energy-Delay Simulations

From Figure 18 it can be seen that at lower supply voltages, the SAPTL performs better than static CMOS, as expected, due to the reduced leakage penalty. At a supply voltage of 1 V, the minimum delay of the SAPTL5 is approximately 7× slower than the minimum CMOS LUT5

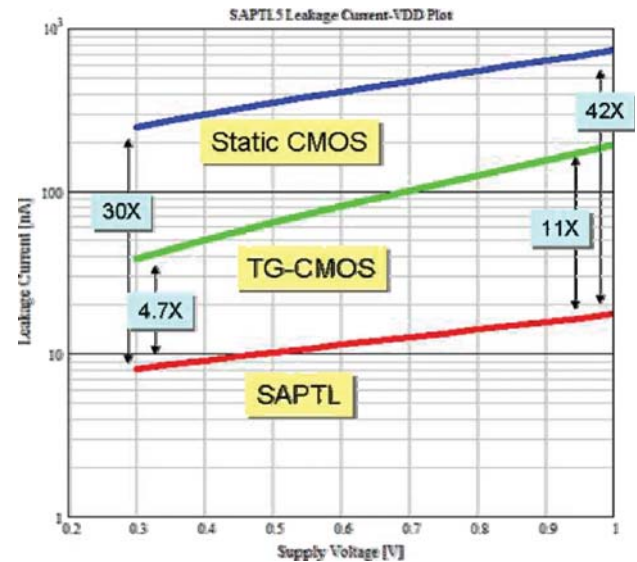


Fig. 17. The leakage characteristics of the SAPTL with $N_{stack} = 5$ and the equivalent CMOS LUT5 and TG LUT5 as the supply voltage is varied from 300 mV to 1 V.

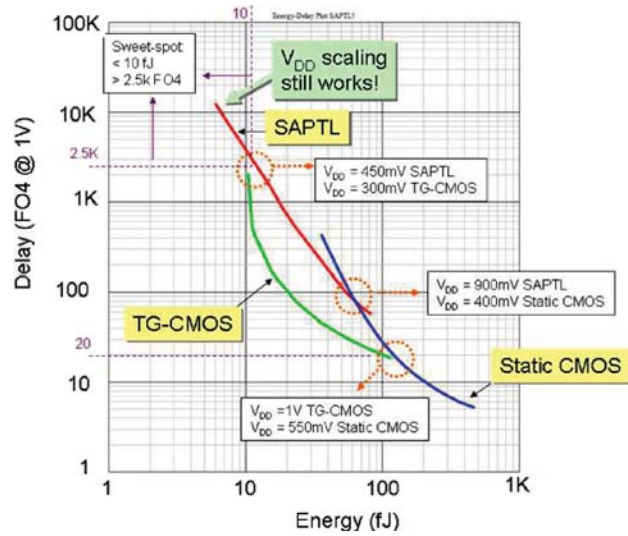


Fig. 18. The energy-delay characteristics of the SAPTL with $N_{stack} = 5$ and the equivalent CMOS LUT5 and TG LUT5 as the supply voltage is varied from 300 mV to 1 V.

delay, while at $V_{dd} = 300$ mV, the energy consumed by the CMOS LUT5 is $6\times$ larger than the SAPTL5.

As expected, since the SAPTL exhibits lower leakage currents, the reduction in supply voltage continues to result in reduced energy per operation. Thus, the preferred region of operation of the SAPTL is at energy levels of below 10 fJ and relatively low speeds on the order of 2500 fanout-of-four (FO4) delays. Again note that this does not take into account the leakage from the implied clock distribution network associated with the SAPTL.

Note that the SAPTL5 with a supply voltage of 450 mV has roughly the same energy-delay point as the TG LUT5 with a supply voltage of 300 mV, and the SAPTL5 with $V_{dd} = 900$ mV is at the same energy-delay point as the CMOS LUT5 with $V_{dd} = 400$ mV. This highlights the fact that as leakage energy starts to degrade the performance of the CMOS LUT and TG LUT at low supply voltages, the SAPTL becomes an attractive alternative, while still operating in superthreshold.

Since the SAPTL favors superthreshold operation, the exponential dependence of delay on physical transistor

parameters, such as threshold voltage, can be avoided, resulting in increased robustness to process variability.

5.3. Threshold Voltage Reduction

Figure 19 shows the minimum energy-delay-product (EDP) of the SAPTL favoring higher supply voltages and lower threshold voltages, thus operating the transistors in strong inversion. This supports the premise that the device threshold voltage can be reduced to improve performance while reducing the impact of the overall SAPTL leakage current. Standard static CMOS circuits on the other hand favors low supply voltages near the transistor threshold voltage.¹⁷

6. ASYNCHRONOUS OPERATION

The ability of any clock distribution network to guarantee timing across the whole integrated circuit is dependent on the amount of variability present, and as technology continues to scale, the delay and power overhead associated with fully synchronous designs may be too prohibitive.¹⁸ By using local timing references rather than global signals (1) a significant reduction in power can be achieved by completely removing the clock distribution network and (2) a possibility of performance improvement if the worst-case delay paths are not used all the time. These advantages are accompanied by (1) an increase in block complexity due to the added handshake logic and (2) an increase wiring density due to the increased number of signals that need to be routed locally.

The sense amplifier topology used in the SAPTL allows a relatively easy way to generate completion detection signals, thus making the asynchronous operation of the SAPTL a straightforward next step in reducing power consumption and increasing performance. One possible implementation of the SAPTL is the self-timed SAPTL.

The Self-Timed SAPTL (ST-SAPTL)¹⁹ shown in Figure 20 replaces the global clock input to the sense amplifier and root node drive input with (1) a root drive enable circuit; (2) a delay line and (3) a Muller C-element.

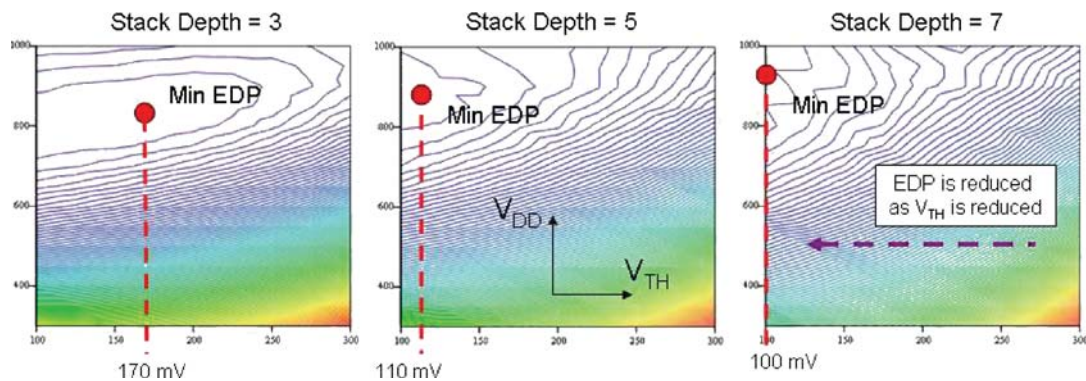


Fig. 19. The energy-delay product (EDP) of the synchronous SAPTL as a function of supply and threshold voltage for $N_{stack} = 3, 5$ and 7 .

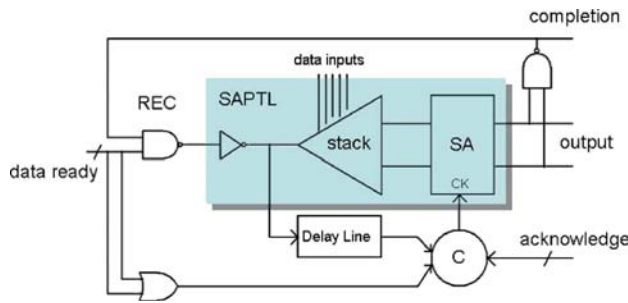


Fig. 20. The self-timed SAPTL (ST-SAPTL).

A completion detection circuit is also added at the output of the sense amplifier to indicate that the sense amplifier has completed its decision process and the latched decision value is currently available at its outputs.

The root enable circuit (REC) drives the root node of the *stack* as well as the input of the delay line. Its output is asserted when all the completion signals, associated with each input to the *stack*, are asserted, and is deasserted when the SA asserts its own completion signal. This ensures that all the inputs are valid before *stack* evaluation can begin, and that the root node is raised to V_{dd} only for the appropriate amount of time. The REC also guarantees that succeeding data does not corrupt the current computation.

The delay line (DL) produces a signal matched to the worst-case delay of the *stack*. This delayed signal is used to trigger the sense amplifier, guaranteeing that the differential voltage at the output of the *stack* is at least $V_{dd}/3$.

The C-element determines whether the sense amplifier is in precharge or in hold mode, and controls the transition between these two states. When the input completion signals are all asserted, the REC launches a signal down the *stack* and the DL. The output of the DL signals the C-element that the output of the *stack* has reached the correct voltage values, and that the SA can now trigger. The C-element then asserts the CK input of the SA, taking the SA from the precharge state into the hold state. The SA performs the estimation operation during this transition and latches its resulting decision. This decision is now made available at the SA outputs and the completion detection circuit asserts the *COMPLETE* signal. The C-element ensures that the SA outputs are valid until all the completion signals from all the fan-outs are asserted. The SA is returned to the precharge state when all the fan-out blocks signify completion and all the fan-in blocks have deasserted their completions, preventing new data from interfering with the present computation.

The ST-SAPTL reduces the overall power consumption by eliminating the need for global clock signals and by turning off the root node drive signal as soon as the sense amplifier latches the outputs. Turning off the root drive signal prevents the internal nodes of the *stack* from rising further, speeding up the preconditioning process and wasting less energy. Self-timed operation, however, results in

- (1) more overhead circuitry, resulting in additional leakage power;
- (2) increased wire count due to the addition of the completion signals;
- (3) added design complexity in matching the delay line to the *stack*;
- (4) the fan-out dependency of the C-element, leading to increased fan-out dependent delay since the completion signals drive loads from both the fan-out and fan-in blocks.

7. FUTURE WORK

To be a feasible alternative to static CMOS logic, a SAPTL design flow that spans design entry, synthesis, optimization and physical layout must exist. Currently, a partial initial design flow has been created and will continue to be an integral part of the development of the SAPTL family.

The effect of various clocking schemes and memory elements cannot be truly evaluated at the building block level since these are in essence, functional and system level concerns. Thus, comparisons between the SAPTL and other logic families must be in terms of larger systems and functional units. In light of this, SAPTL models are being created to facilitate system-level design space exploration to quickly evaluate the applicability of the SAPTL to a particular application. Two 90 nm CMOS test chips have been designed and are being tested in order to validate these ideas and help create these models and tools.

It is interesting to note that the leakage reduction techniques presented in Section 2 are device-level techniques and that they are also applicable to the SAPTL's sense amplifier and drive inverter and will be used to optimize the SAPTL at the system and functional level.

The *stack* presented in this paper is a fully decoded pass transistor tree that allows the implementation of any logic function of N_{stack} inputs. System- and functional-level simplification of these pass transistor trees can also be done and will also be used for system-level optimization.

This work presents a first attempt at evaluating the properties of the SAPTL as an alternative low energy circuit topology. The advantages of the SAPTL over static CMOS logic are expected to be more significant for more advanced technologies such as 65 nm CMOS and beyond.

8. CONCLUSION

The SAPTL is an alternative circuit topology that allows the reduction of energy per operation via voltage scaling even in the presence of leakage. It allows aggressive threshold voltage scaling since the V_{th} of the *stack* transistors can now be decoupled with its subthreshold leakage, allowing for very low threshold voltages and thus allowing the *stack* transistors to remain in the superthreshold region. In addition, the differential signaling used by the SAPTL lends itself to synchronous and asynchronous operation

and the inherent layout regularity points to the SAPTL as a very good candidate for robust ultra low energy operation.

Acknowledgment: The authors wish to acknowledge the contributions of the students, faculty and sponsors of the Berkeley Wireless Research Center, the National Science Foundation Infrastructure Grant No. 0403427, the wafer fabrication donation of STMicroelectronics, and the Gigascale Research Center (GSRC) for their support of this research.

References

1. The International Technology Roadmap for Semiconductors (ITRS) 2005 Edition, Technical Report, <http://www.itrs.net/Links/2005ITRS/Home2005.htm>.
2. T. Sakurai, Perspectives on power aware electronics. *ISSCC Digest of Technical Papers* 1, 26 (2003).
3. J. Rodrigues, T. Olsson, L. Sörnmo, and V. Öwall, Digital implementation of a wavelet-based event detector for cardiac pacemakers. *IEEE Trans. on Circuits and Systems I: Regular Papers* 52 (2005).
4. M. Ghovanloo and K. Najafi, A modular 32-site wireless neural stimulation microsystem. *IEEE Journal of Solid-State Circuits* 39 (2004).
5. N. Kim, T. Austin, D. Blaauw, T. Mudge, K. Flautner, J. Hu, M. Irwin, M. Kandemir, and V. Narayanan, Leakage current: Moore's law meets static power. *IEEE Computer* 36, 68 (2003).
6. J. W. McPherson, Scaling-induced reductions in CMOS reliability margins and the escalating need for increased design-in reliability efforts. *International Symposium on Quality Electronic Design* (2001), pp 123–130.
7. A. Agarwal, K. Kang, S. Bhunia, J. Gallagher, and K. Roy, Effectiveness of low power dual- V_t designs in nano-scale technologies under process parameter variations. *International Symposium on Low Power Electronics and Design*, August (2005), pp. 14–19.
8. Y. Taur and T. Ning, *Fundamentals of Modern VLSI Devices*, Cambridge University Press, Cambridge, UK (1998).
9. B. Chatterjee, M. Sachdev, S. Hsu, R. Krishnamurthy, and S. Borkar, Effectiveness and scaling trends of leakage control techniques for sub-130 nm CMOS technologies. *International Symposium on Low Power Electronics and Design*, August (2003), pp. 122–127.
10. S. Yu, Extremely scaled planar bulk CMOS: Challenges and options. *7th International Conference on Solid-State and Integrated Circuits Technology*, October (2004), pp. 41–46.
11. A. Keshavarzi, S. Narendra, S. Borkar, C. Hawkins, K. Royi, and V. De, Technology scaling behavior of optimum reverse body bias for standby leakage power reduction in CMOS IC's. *International Symposium on Low Power Electronics and Design* (1999), pp. 252–254.
12. K. Yano, Y. Sasaki, K. Rikino, and K. Seki, Top-down pass-transistor logic design. *IEEE Journal of Solid-State Circuits* 31, 792 (1996).
13. R. Shelar and S. Sapatnekar, BDD decomposition for delay oriented pass transistor logic synthesis. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 13, 957 (2005).
14. P. Buch, A. Narayan, A. Newton, and A. Sangiovanni-Vincentelli, Logic synthesis for large pass transistor circuits. *IEEE/ACM International Conference on Computer-Aided Design*, November (1997), pp. 663–670.
15. W. C. Elmore, The transient analysis of damped linear networks with particular regard to wideband amplifiers. *J. Appl. Phys.* 19, 55 (1948).
16. V. Kheterpal, V. Rovner, T. Hersan, D. Motiani, Y. Takegawa, A. Strojwas, and L. Pileggi, Design methodology for IC manufacturability based on regular logic-bricks. *Proceedings of the 42nd Design Automation Conference*, June (2005), pp. 353–358.
17. A. Wang and A. Chandrakasan, A 180-mV subthreshold FFT processor using a minimum energy design methodology. *IEEE Journal of Solid State Circuits* 40, 310 (2005).
18. A. J. Martin and M. Nystrom, Asynchronous techniques for system-on-chip design. *Proceedings of the IEEE*, June (2006), Vol. 94, pp. 1089–1120.
19. T.-T. Liu, L. Alarcón, M. Pierson, and J. Rabaey, Asynchronous computing in sense amplifier-based pass transistor logic, Berkeley Wireless Research Center (BWRC) Internal Technical Report (2007), unpublished.

Louis P. Alarcón

Louis P. Alarcón was born in Quezon City, the Philippines. He finished received his B.S. and M.S. degrees in electrical engineering from the University of the Philippines, Diliman, in 1995 and 2002 respectively. He has been a faculty member of the University of the Philippines since 1995, and is on an extended study leave. Starting the fall of 2004, he has been with the Department of Electrical Engineering and Computer Science at the University of California, Berkeley, where he is currently working towards his Ph.D. degree. His research is on circuit design techniques for very low energy applications and is a member of the Berkeley Wireless Research Center.

Tsung-Te Liu

Tsung-Te Liu received the B.S. and M.S. degrees in electrical engineering from the National Taiwan University, Taipei, Taiwan, R.O.C., in 2002 and 2004, respectively. He is currently working toward the Ph.D. degree at the University of California, Berkeley, where he is a member of the Berkeley Wireless Research Center. From 2004 to 2005, he was with MediaTek Inc., HsinChu, Taiwan, R.O.C., where he was involved in the system-level and analog circuit design for wireless communications. His current research interests include low-energy circuit and system design for wireless communications.

Matthew D. Pierson

Matthew D. Pierson received the B.S. degree in electrical engineering from The University of Texas at Austin in 2004, and the M.S. degree in electrical engineering from The University of California at Berkeley in 2007. His main research interests include low-energy and asynchronous circuit design. He currently works at Texas Instruments in Dallas, Texas designing high-performance digital signal processors.

Jan M. Rabaey

Jan M. Rabaey received the EE and Ph.D. degrees in applied sciences from the Katholieke Universiteit Leuven, Belgium. From 1983–1985, he was connected to the UC Berkeley as a Visiting Research Engineer. From 1985–1987, he was a research manager

at IMEC, Belgium, and in 1987, he joined the faculty of the Electrical Engineering and Computer Science department of the University of California, Berkeley, where he is now holds the Donald O. Pederson Distinguished Professorship. He has been a visiting professor at the University of Pavia (Italy), Waseda University (Japan), Technical University Delft (Netherlands), Victoria Technical University and the University of New South Wales (Australia). From 1999 until 2002, he was the Associate Chair of the EECS Department at Berkeley. He is currently the scientific co-director of the Berkeley Wireless Research Center (BWRC), as well as the director of the MARCO GigaScale Systems Research Center (GSRC). Professor Rabaey received numerous scientific awards, including the 1985 IEEE Transactions on Computer Aided Design Best Paper Award (CAS), the 1989 Presidential Young Investigator award, the 1994 Signal Processing Society Senior Award, and the 2002 ISSCC Jack Raper Award. He is an IEEE Fellow. He has served as associate editor for the IEEE Journal of Solid State Circuits and the TODAES ACM Journal. He was one of the founders and first technical program chair of the International Symposium on Low Power Electronic Design (ISLPED). From 1994 until 2002, he served on the Executive Committee of the Design Automation Conference, of which he was both technical program chair and general chair. Prof. Rabaey serves on the technical advisory board of a range of companies and research institutes focused in the areas of design automation, semiconductor intellectual property and wireless systems. His current research interests include the conception and implementation of next-generation integrated wireless systems.