Circuits & Systems

Test Circuits for Characterization of Process, Device, and Interconnect Variation

K. Balakrishnan, A. H. Chang, N. Drego, D. S. Boning
Sponsorship: SRC/FCRP IFC, SRC/FCRP C2S2, TSMC, Samsung Electronics

Due to the continuous and aggressive scaling of CMOS technology, the variability in parameters that are most critical in determining the quality and robustness of a device must be accurately characterized. For process- and device-level characterization, a test chip that characterizes contact plug resistance variability has been designed, fabricated, and measured. In addition, a test circuit has been designed to characterize layout-induced systematic variability in transistor saturation current.

One trend affecting device interconnect is the impact of contact resistance, which is becoming an increasingly larger component of the total resistance in a MOSFET. The characterization of contact resistance variability will allow for the generation of a compact model that incorporates sensitivities to various parameters into the determination of individual contact resistances. A test chip that can characterize the variability of resistance in contact plugs is implemented in a 90nm CMOS process; the chip contains over 40,000 devices under test. Figure 1 shows the distribution of normalized contact plug resistance across one die. These results are unique in measuring components of plug resistance within a transistor environment, as opposed to measurements using contact/via chains. Results show that the within-die resistance has a $\sigma/\mu \approx 4.6\%$ and the die-to-die resistance has a $\sigma/\mu \approx 4.4\%$. In addition, there is a clear systematic offset in the contact resistance between two regions of the chip, possibly caused by variability in the step-and-scan lithography process. The contact plug resistance is also a function of the distance from the contact to both the polysilicon gate and the edge of the diffusion region. Device simulations to further investigate this effect are ongoing.

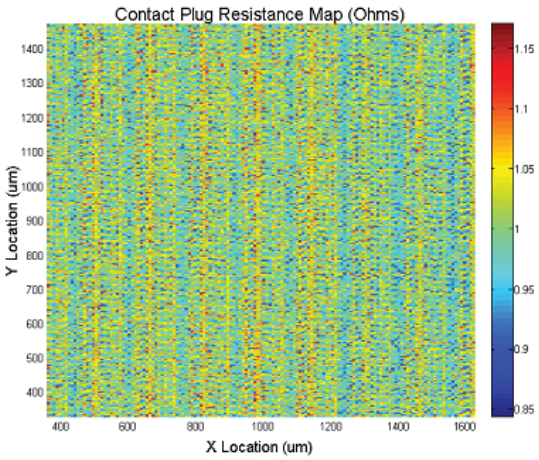

FIGURE 1: Normalized contact plug resistance map for a representative test chip.

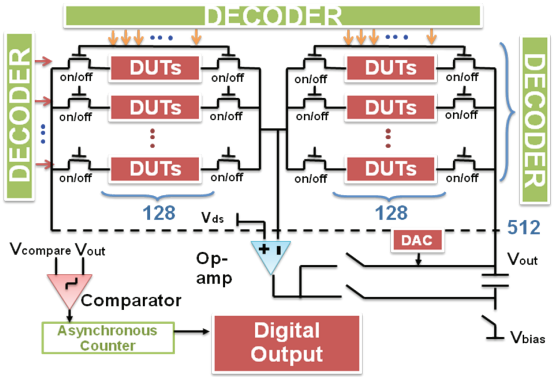

FIGURE 2: Hierarchical design of on-chip current measurement.

A test chip has also been designed to study layout-induced systematic variation in transistor behavior. Two pattern densities are chosen in our design: polysilicon density and shallow-trench isolation (STI) density. A test structure is designed to study these pattern density effects, as well as systematic spatial dependency between transistors. One goal is to decouple the impact of different variation sources on transistor characteristics and determine the farthest layout distance that should be taken into account when determining local transistor characteristics. More accurate transistor models based on surrounding layout details can then be built using our results.

The test structure is divided into six blocks, each with a different polysilicon density or STI density. A rapid change of pattern density between blocks is designed to emulate a step response for future modeling. The two pattern densities are chosen to reflect the introduction of a new process technology, such as strain engineering or rapid thermal annealing. The test structure is designed to have more than 260K devices under test (DUT). The test circuit is designed to enable on-chip current measurement with a high dynamic range analog-to-digital converter (ADC). The ADC, shown in Figure 2, has a dynamic range of over four orders of magnitude, to measure current from 50nA to 1mA. The test chip also implements a hierarchical design with a minimum amount of periphery circuitry around the DUTs, so most of the chip area is dedicated to the transistors under test.

Variation-Induced Energy Overhead Reduction in Multicore Processors

N. Drego, A. P. Chandrakasan, D. S. Boning
Sponsorship: SRC/FCRP C2S2

In modern process technologies it is clear that performance, power/energy, cost and variation are intricately linked to each other. As system architects value core homogeneity in multicore processors, the burden of finding energy-efficient variation mitigation solutions increases. Balancing other constraints, such as area overhead and design complexity, must also play into any mitigation technique. Evaluation of some of the more common mitigation schemes, in the context of these constraints, reveals a need for an energy-efficient technique capable of guaranteeing both performance and yield constraints while introducing minimal overhead.

Introducing additional system voltages (fewer than one per core) provides a compromise between excess design/area overhead and energy efficiency while maximizing performance for multicore systems with N cores. An analytic framework capable of optimal voltage selection, to minimize energy and reduce the energy overhead required to mitigate variation, is developed and analyzed in-depth. Specifically, a simple, efficient Minimum Energy Voltage Selection (MEVS) algorithm to select optimal voltages forms the basis of this framework, allowing for multiple types of analysis. Though the optimization problem is not convex in general and simplifying approximations are used, the algorithm is nevertheless able to find optimal (for N = 2) or near optimal solutions. Furthermore, the behavior of the algorithm is mathematically bounded and shown to perform according to the bounds.

Using a custom simulation methodology and the MEVS algorithm, a core designed specifically for multicore contexts is simulated to observe the magnitude of performance variation and the impact of introducing additional voltages to a massively parallel 1K-core processor. Analysis shows that a single additional power-supply voltage provides the greatest incremental impact, with 59-75% reduction in the variation-induced energy overhead and 6-16% total energy reduction, as seen in Figure 1. The desired yield constraint also has a significant impact on energy reduction: though counterintuitive, turning off a small fraction of the thousand cores can provide a positive trade-off between performance and energy, as depicted in Figure 2, with multiple system voltages further improving this trade-off. Lastly, when voltage regulator efficiencies are properly accounted for, using only a few system voltages ($2 \leq N \leq 10$) provides greater energy reduction than supplying each core with its own voltage.

REFERENCES

[1] N. Drego, "Characterization and Mitigation of Process Variation in Digital Circuits and Systems," PhD Thesis, Massachusetts Institute of Technology, Cambridge, 2009.

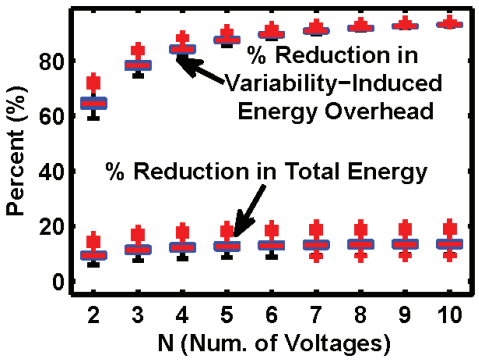

FIGURE 1: Energy reduction by adding voltages to a multicore system. A single additional power-supply voltage provides the most incremental benefit in such systems, with asymptotic gains in energy reduction for additional power-supply voltages.

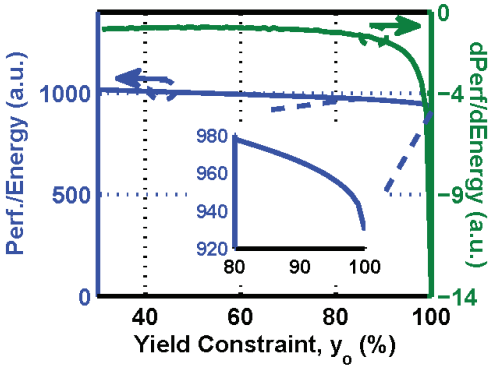

FIGURE 2: Joint performance/energy metric as a function of yield. Due to the exponential tails of the performance distributions, turning off the last 5-15% of poor-performing cores leads to a positive trade-off. The incremental change in performance for a given increment of energy severely degrades at the yield constraint is increased beyond 85% (green plot, right axis).

45nm Direct-battery DC-DC Converter for Mobile Applications

S. Bandyopadhyay, Y. K. Ramadass, A. P. Chandrakasan
Sponsorship: MTL, MIT

Mobile applications use lithium-ion batteries as the power supply. With the aggressive downscaling of transistors, it is becoming increasingly difficult to interface the low-voltage digital core with the battery, which may be at voltages as high as 5.5V. The interfacing can be done by a DC-DC converter, which may be a separate IC designed on an older generation process capable of handling high voltages. In this work, a buck converter has been designed in a 45-nm process so that it can be integrated with the 45-nm digital core on the same die. This design gives us cost and performance advantages for a single chip solution. This work presents a wide load range DC-DC converter that gives more than 80% efficiency from 10μA to 100mA of load currents. This efficiency is accomplished by using both the PFM and PWM control schemes [1], [2]. Further, the converter requires Switched Capacitor (SC) converters [3] to generate the voltages for stacking and for the control circuitry. Figure 1 shows the block diagram of the cellular system with the battery and the direct-battery DC-DC converter.

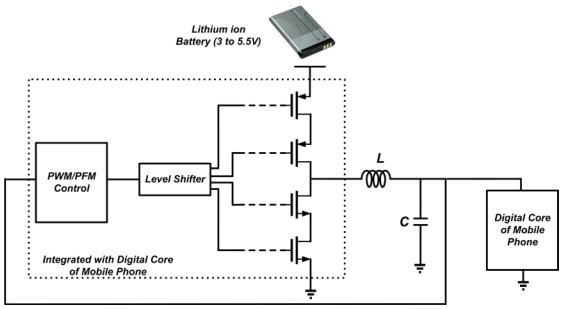

FIGURE 1: Block Diagram showing the Lithium ion battery with DC-DC converter and the digital core of mobile phone at 45nm.

REFERENCES

[1] J. Xiao et.al "A 4-µA Quiescent-Current Dual-Mode Digitally Controlled Buck Converter IC for Cellular Phone Applications," *IEEE Journal of Solid State Circuits*, Vol. 39, No. 12, December 2004

[2] Y. K. Ramadass and A.P. Chandrakasan, "Minimum Energy Tracking Loop with Embedded DC-DC Converter Delivering Voltages down to 250mV in 65nm CMOS," *International Solid State Circuits Conference*, February 2007, San Francisco.

[3] J. Kwong, Y. K. Ramadass, N. Verma, and A.P. Chandrakasan, "A 65nm Sub-Vt Microcontroller With Integrated SRAM and Switched Capacitor DC-DC Converter," *IEEE Journal of Solid State Circuits*, Vol.44 No.1, Jan 2009.

A Low-power AES Engine with Resistance to Differential Power Analysis Side-channel Attacks

H. W. Chung, A. P. Chandrakasan

Security concerns for transmission or storage of data by battery-operated wireless systems require the development of an energy-efficient encryption coprocessor. However, even with the security ICs, core information can be discovered by attackers since the ICs are vulnerable to side-channel attacks. Among all the side-channel attacks, differential power analysis (DPA) attack is effective in finding a secret key. Measuring the current from power supply and then performing statistical analysis of the measured power traces can lead to discovery of the secret key. Therefore, development of an energy-efficient encryption processor that is immune to differential power analysis attack is required for the secure transmission and storage of the data in battery-operated security ICs.

The Advanced Encryption Standard algorithm [1] is a block cipher that converts 128-bit plaintext to ciphertext with selectable key lengths (128, 192, or 256 bits). The algorithm is organized as a repeated "round transformation" that includes four types of sub-operations, i.e., "S-Box," "ShiftRows," "MixColumns," and "AddRoundKey"(Figure 1). The DPA attack occurs at the transition of the RB registers (Figure 1). Correlation between the modeled power trace based on a guessed secret key and the measured transition power of the RB registers can give enough information to find out a secret key.

A novel architecture is suggested to guarantee DPA-immunity. First, S-Box is designed with decoder-encoder block [2], and then the RB register is moved between the decoder and encoder of S-Box so that the number of 0 to 1 transitions is the same for all the possible encryption data. With this design, attackers cannot get any information about the secret key since there is no data-dependency in the measured current trace of the RB registers. Area overhead exists area for this architecture, but the power overhead is not significant with proper clock-gating of the RB registers. Moreover, maximizing parallel operation of data can lower the supply voltage into the sub-threshold region with satisfying the performance specification. Sub-threshold operation is beneficial not only for energy-efficient operation, but also for improvement of DPA-immunity.

REFERENCES

[1] NIST, "Federal Information Processing Standards PUB 197 Advanced Encryption Standard" (2001, Nov.). [Online]. Available: http://www.nist.go/aes/

[2] G. Bertoni, M. Macchetti, and L. Negri, "Power-Efficient ASIC Synthesis of Cryptographic Sboxes," *ACM Great Lake Symposium on VLSI*, pp. 277-281, April. 2004

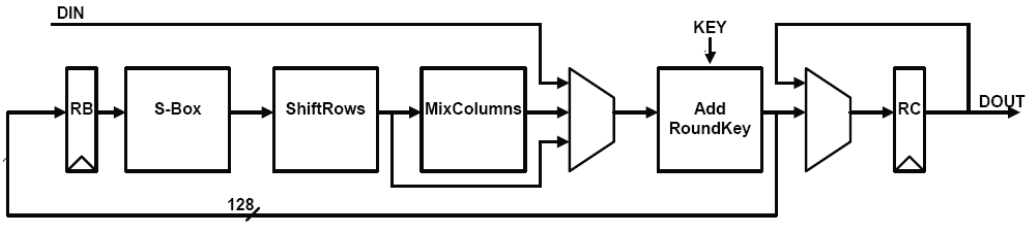

FIGURE 1: Architecture of AES core: DPA attack mainly occurs at the transition time of the RB registers.

A Pulsed UWB Receiver SoC for Insect Motion Control

D. C. Daly, P. P. Mercier, M. Bhardwaj, H. Liang, A. P. Chandrakasan
Sponsorship: DARPA, NSERC

For decades, scientists and engineers have been fascinated by cybernetic organisms, or cyborgs, that fuse artificial and natural systems. Cyborgs enable the harnessing of biological systems that have been honed by evolutionary forces over millennia to achieve astounding feats. Male moths can detect a single pheromone molecule, a sensitivity of roughly 10^{-21} grams. Thus, cyborgs can perform tasks at scales and efficiencies that would ordinarily seem incomprehensible. Semiconductor technology is central to realizing this vision because it offers powerful processing and communication capabilities as well as low weight, small size, and deterministic control. An emerging cyborg application is moth flight control, where electronics and MEMS devices are placed on and within a moth to control flight direction. To receive commands on the moth, a lightweight, low-power and low-volume receiver is required. Figure 1 presents an overview of the moth flight control system being developed in collaboration with other scientists and researchers at MIT, the University of Washington, and the University of Arizona.

A critical component of the hybrid-insect system is the wireless communication link, which provides flight control commands to the moth. Pulsed ultra-wideband (UWB) wireless signaling is employed as UWB radios can achieve highly integrated, energy-efficient operation in nanometer CMOS processes [1]-[3]. Power, weight and volume are all highly constrained, necessitating a highly integrated solution with minimal off-chip components. Data is transmitted by PPM modulation in one of three 500-MHz channels in the 3-to-5-GHz band. Figure 2 presents a block-diagram of the wireless receiver system-on-chip. The non-coherent receiver amplifies, squares, and integrates received pulses to measure the amount of energy received in a given time period. A differential, inverter based front end is employed to reduce current consumption while allowing for a single 1V core power supply. The highly duty cycled RX requires 0.5-to-1.4nJ/bit and achieves a sensitivity of -76dBm at a data rate of 16 Mb/s (10^{-3} BER). The moth stimulator generates a multi-channel, digital, pulse-width modulated signal to control flight direction. The radio has been integrated on a miniature, 1g PCB; and preliminary flight control in a wind tunnel has been demonstrated.

REFERENCES

[1] P. P. Mercier, D. C. Daly, and A. P. Chandrakasan, "A 19pJ/pulse UWB Transmitter with Dual Capacitively-Coupled Digital Power Amplifiers," *IEEE Radio Frequency Integrated Circuits Symposium*, pp. 47-50, June 2008.

[1] P. P. Mercier, M. Bhardwaj, D. C. Daly, A. P. Chandrakasan, "A 0.55V 16Mb/s 1.6mW Non-Coherent IR-UWB Digital Baseband with ±1ns Synchronization Accuracy," *IEEE International Solid-State Circuits Conference (ISSCC)*, pp. 252-253, Feb. 2009.

[2] D. C. Daly, P. P. Mercier, M. Bhardwaj, A. L. Stone, J. Voldman, R. B. Levine, J. G. Hildebrand, A. P. Chandrakasan, "A Pulsed UWB Receiver SoC for Insect Motion Control," *IEEE International Solid-State Circuits Conference (ISSCC)*, pp. 200-201, Feb. 2009.

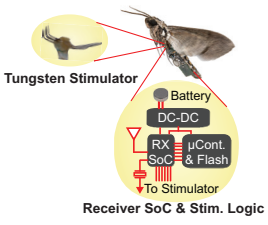

FIGURE 1: Overview of hybrid-insect flight control system.

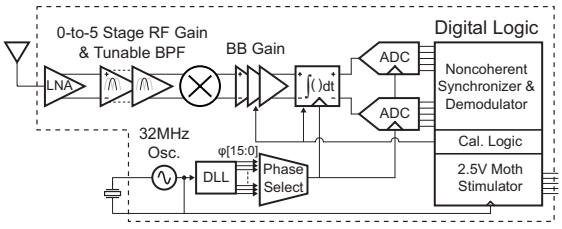

FIGURE 2: Block diagram of the wireless receiver SoC.

A Low-voltage, Fault-tolerant Microprocessor

N. Ickes, Y. Koken, F. Pappalardo (STMicroelectronics), A. P. Chandrakasan
Sponsorship: STMicroelectronics

Digital logic circuits are most energy-efficient when operated at very low (near subthreshold) voltages. However, many severely energy-constrained applications (e.g., implanted medical devices) also require high reliability, and the rate of radiation-induced soft-errors increases significantly at low voltages [1]. Existing techniques for improving soft-error resilience come with significant power overhead. The purpose of this project is to investigate error detection and correction mechanisms, for both memory and logic, which are specifically optimized for micropower, low-voltage systems.

Soft error events affect both combinational and sequential logic gates, as shown in Figure 1. Due to the necessarily tight power supply voltage margins at low voltage, power supply droop can also generate errors by causing signals to arrive late. Flip-flop and latch designs capable of detecting these errors [2], [3] have been previously demonstrated by others. However, their work has focused on high-performance processors with significant speculative state, so that errors can be recovered from simply by flushing speculative instructions from the pipeline. Micro-power processors have little or no speculative state, so we are exploring alternative error recovery mechanisms.

SRAMs make up the majority of the area of most microprocessor chips and must be continuously powered for data retention. Designing SRAMs for low voltage operation is therefore particularly important. However, scaling down power supply voltage not only increases susceptibility to radiation-induced soft errors but also degrades bit-cell stability due to device variation effects. Simple SECDED Hamming codes are quite effective at recovering from radiation-induced errors. We are exploring the use of higher-order BCH codes capable of correcting multiple bits per word, in order to address both radiation and bitcell variation induced errors (Figure 2). In particular, we are looking at how error-correcting capability might be efficiently scaled with the operating voltage.

As a test vehicle for these techniques, we are modifying an existing 32b microcontroller for operation at 0.5V in 90nm CMOS. In addition to adding error detection and correction logic, we are exploring the design of miniature instruction and data caches to both improve performance at very low voltages, and to reduce memory access energy.

REFERENCES

[1] T. Heijmen, P. Roche, G. Gasiot, K. R. Forbes and D. Giot, "A Comprehensive Study on the Soft-Error Rate of Flip-Flops From 90-nm Production Libraries," *Device and Materials Reliability, IEEE Transactions on,* vol. 7, pp. 84-96, 2007.

[2] S. Das, C. Tokunaga, S. Pant, W. Ma, S. Kalaiselvan, K. Lai, D. M. Bull and D. T. Blaauw, "RazorII: In Situ Error Detection and Correction for PVT and SER Tolerance," *Solid-State Circuits, IEEE Journal of,* vol. 44, pp. 32-48, 2009.

[3] K. A. Bowman, J. W. Tschanz, Nam Sung Kim, J. C. Lee, C. B. Wilkerson, S. L. Lu, T. Karnik and V. K. De, "Energy-Efficient and Metastability-Immune Resilient Circuits for Dynamic Variation Tolerance," *Solid-State Circuits, IEEE Journal of,* vol. 44, pp. 49-63, 2009.

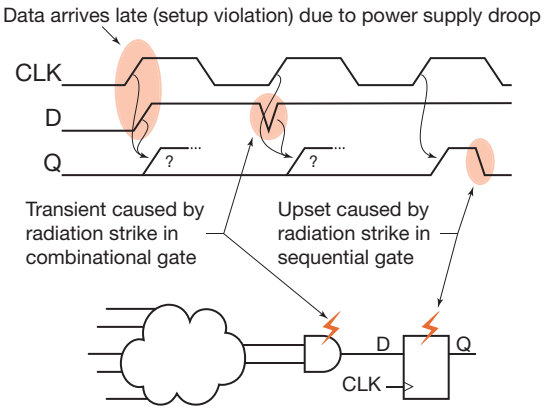

FIGURE 1: Soft error mechanisms in logic: delay violations, transients in combinational logic, and upsets in sequential gates.

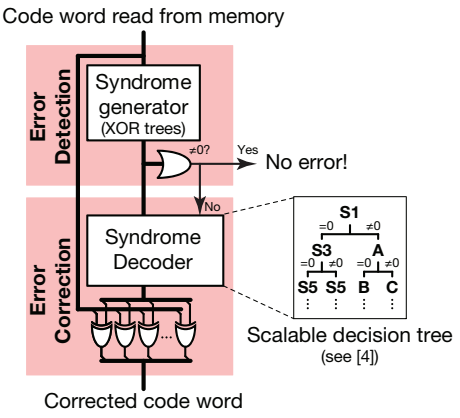

FIGURE 2: BCH decoding, showing the separate phases of error detection and correction.

Wearable Medical Monitoring Platform

J. Kwong, P. Mercier, M. Yip, A. P. Chandrakasan

Sponsorship: Texas Instruments, DARPA, SRC/FCRP C2S2

Advances in mobile electronics are fueling new possibilities in medical monitoring in which sophisticated, wearable devices can monitor a subject's vital signs. As illustrated in Figure 1, these signals can be securely transmitted over the internet via a local relay (often a cell phone or PDA) for preventative medicine, diagnostics, or emergency monitoring purposes. In these applications, comfort and convenience are important considerations, motivating a high level of integration to achieve small form factors and long operating lifetimes from a small battery or scavenged energy. Fortunately, the low rates of biological signals, which are typically on the order of tens to hundreds of Hz, make basic monitoring applications amenable to low-power processing [1], [2]. To support the wide range of signals, sensors, and algorithms, we propose a reconfigurable and energy-efficient platform for medical monitoring.

A block diagram of the platform is shown in Figure 2. The system requires a flexible sensor front-end and ADC that can interface with different types of sensors. The sensor front-end must have adjustable gain, bandwidth, and noise settings, as signals from different sensors can vary by several orders of magnitude in amplitude and frequency. The ADC should have configurable resolution from 8 to 12 bits to enable a variety of applications. A processor retrieves the digitized data from the ADC and performs local processing tasks specified in software. To save energy, the processor can operate at a low voltage and frequency when executing simple algorithms, but it maintains flexibility by elevating its voltage when high performance is needed. The architecture includes hardware accelerators to speed up computations. When not in use, each accelerator can be powered off to reduce idle leakage. To further save energy, local sensor nodes transmit processed data only to the local relay, instead of directly to the cellular network. The short transmission distance (1-2m) and inherently low data rates (1-10kbps) promote the design of highly-digital, energy-efficient communication circuits involving wired clothing networks [3] and/or wireless ultra-wideband [4]. A high-powered radio, as typically found in a PDA, can then consolidate sensor node data and transmit them over the cellular network to the internet for monitoring purposes.

REFERENCES

[1] A. P. Chandrakasan, N. Verma, D.C. Daly, "Ultralow-Power Electronics for Biomedical Applications," *Annual Review of Biomedical Engineering*, vol. 10, pp. 247-274, August 2008.

[2] A.C.-W. Wong, D. McDonagh, G. Kathiresan, et al., "A 1V, Micropower System-On-Chip for Vital-Sign Monitoring in Wireless Body Sensor Networks," *IEEE International Solid-State Circuits Conference*, February 2008, pp. 138-139.

[3] S. Lee, J. Yoo, and H-J. Yoo, "A Healthcare Monitoring System with Wireless Woven Inductor Channels for Body Sensor Network," *IEEE International Workshop on Wearable and Implantable Body Sensor Networks*, June 2008.

[4] P.P. Mercier, D.C. Daly, and A.P. Chandrakasan, "A 19pJ/pulse UWB Transmitter with Dual Capacitively-Coupled Digital Power Amplifiers," *IEEE Radio Frequency Integrated Circuits (RFIC) Symposium*, June 2008, pp. 47-50.

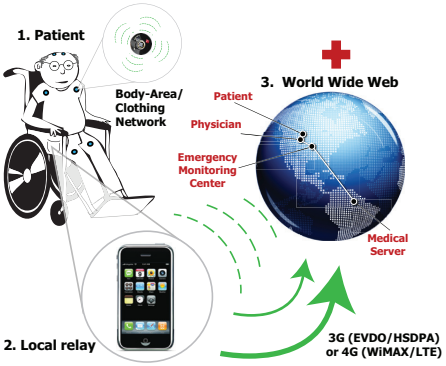

FIGURE 1: The vision of connected health care enabled by the wearable medical monitoring platform, a local relay and the World Wide Web.

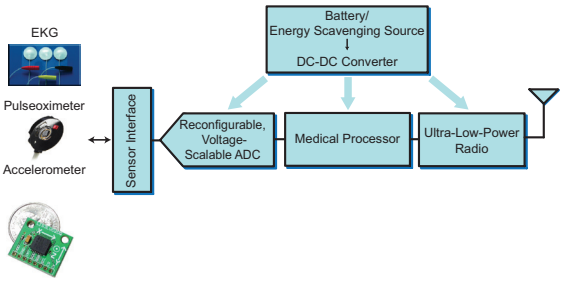

FIGURE 2: Block diagram of the wearable medical monitoring platform.

A Non-coherent Pulsed-UWB Digital Baseband Employing Quadratic Correlation

P. P. Mercier, M. Bhardwaj, D. C. Daly, A. P. Chandrakasan
Sponsorship: DARPA

Pulsed-ultra-wide-band (UWB) radios are finding increasing use in low-data-rate sensing applications, in part because they can be easily duty-cycled to achieve extreme energy efficiency. Within pulsed radios, non-coherent RF front-ends that use simple square-and-integrate samplers offer significant energy-per-bit savings over their coherent counterparts [1], [2]. However, such samplers lose phase information and accumulate squared noise over the integration period. While this loss increases the SNR required to relay a bit reliably, the greater challenge is achieving signal synchronization. Telemetry applications often have small payloads (10-100 bits) where synchronization time dominates. Furthermore, synchronization performance is being continually pushed to enable positioning capability. Hence, the ultimate advantage of non-coherent receivers relies on their ability to synchronize efficiently.

The proposed UWB digital baseband, whose top-level block diagram is shown in Figure 1, reduces synchronization overhead without any increase in RF front-end power or complexity using modified synchronization codes and quadratic correlators (QCORRs) in place of matched filters [3]. These features result in synchronization times that are 11x shorter than repetition codes, or, equivalently, require 10dB lower SNR than repetition codes, as shown in Figure 2. To further reduce energy and decrease synchronization latency, the baseband is highly parallelized, and operates down to a supply voltage as low as 0.55V. Implemented in 90nm CMOS, the baseband occupies an area of 2.55mm^2 and consumes an average power of 1.6mW during a typical preamble.

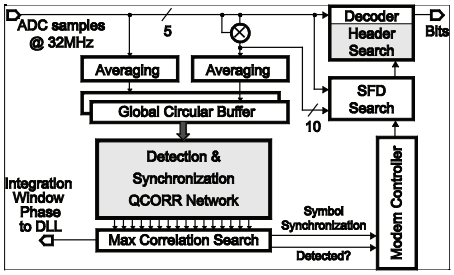

FIGURE 1: A top-level block diagram of the fabricated UWB modem.

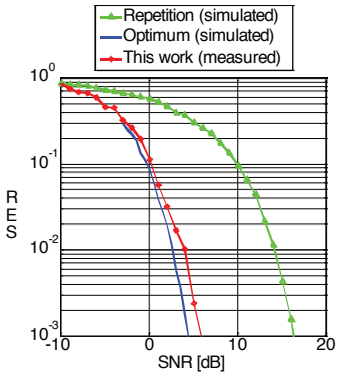

FIGURE 2: Synchronization error rate (SER) curves for the proposed baseband plotted with a simulated repetition-code receiver and an ideal maximum-likelihood receiver.

REFERENCES

[1] F.S. Lee, and A.P. Chandrakasan, "A 2.5nJ/b 0.65V 3-to-5GHz Subbanded UWB Receiver in 90nm CMOS," *IEEE ISSCC Dig. Tech. Papers*, pp. 116-117, Feb. 2007.

[2] D.C. Daly, P. P. Mercier, M. Bhardwaj, A. L. Stone, J. Voldman, R. B. Levine, J. G. Hildebrand, A. P. Chandrakasan, "A Pulsed UWB Receiver SoC for Insect Motion Control," *IEEE ISSCC Dig. Tech. Papers*, pp. 200-201, February 2009.

[3] P. P. Mercier, M. Bhardwaj, D. C. Daly, and A. P. Chandrakasan, " A 0.55V 16Mb/s 1.6mW Non-Coherent IR-UWB Digital Baseband with ±1ns Synchronization Accuracy," *IEEE ISSCC Dig. Tech. Papers*, pp. 252-253, Feb. 2009.

Statistical SRAM Optimization

M. Qazi, A. P. Chandrakasan
Sponsorship: SRC/FCRP C2S2

Process variation in deep-submicron technologies is especially pronounced for embedded SRAM, which must meet demands for higher density and higher performance at increased levels of integration. The very low failure probabilities required of the highly repeated memory cells must be predicted—both in terms of functionality and performance. To address these issues, we have developed a novel statistical simulation methodology for bitcell stability. Next, we extend this algorithm with new circuit insight into evaluating the worst-case performance of the SRAM critical path. The overall goal is to provide a comprehensive and efficient statistical design methodology of embedded memory from memory cell to block-level architecture.

We have developed and reported a novel importance sampling via norm minimization algorithm that achieves up to 10,000x speedup over conventional Monte Carlo simulation for very low failure probabilities [1]. We are also working on a statistical simulation tool flow that evaluates worst-case performance of SRAM critical paths.

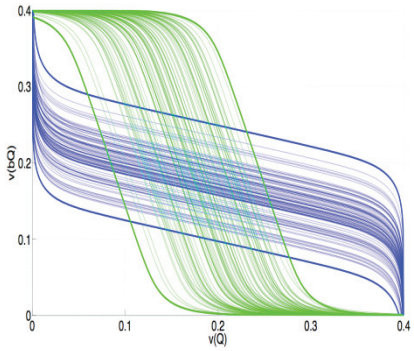

FIGURE 1: The SRAM memory cell exhibits a significant amount of variation as the supply voltage is reduced. Shown here is the variation under the retention condition of 32nm SRAM at 400mV supply for predictive technology models [2].

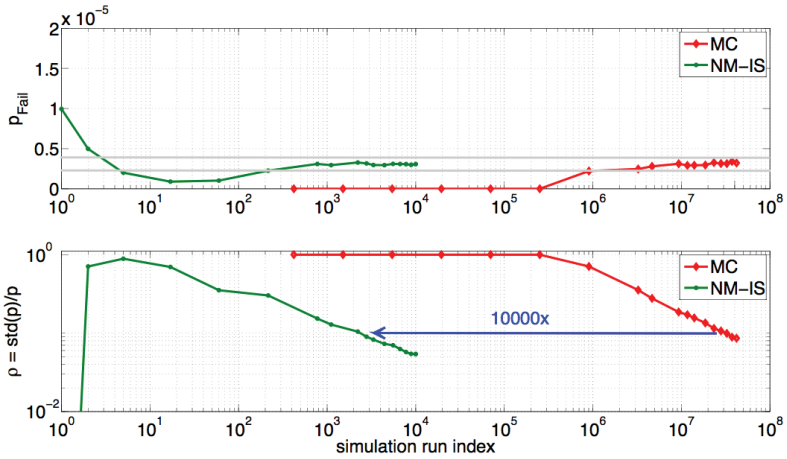

FIGURE 2: The top plot shows the comparison of the retention failure as evaluated by Monte Carlo simulation and the Norm Minimization Importance Sampling algorithm. 10,000X fewer runs are required with the sped-up algorithm for the same level of confidence as depicted in the lower plot.

REFERENCES

[1] L. Dolecek, M. Qazi, D. Shah and A. Chandrakasan, "Breaking the simulation barrier: SRAM evaluation through norm minimization," in *IEEE/ACM International Conference on Computer-Aided Design,* 2008, pp. 322-329.

[2] W. Zhao and Y. Cao, "New generation of predictive technology model for sub-45nm design exploration," *IEEE Transactions on Electron Devices,* vol. 53, pp. 2816-2823, 2006.

An Efficient Piezoelectric Energy-harvesting Interface Circuit

Y. K. Ramadass, A. P. Chandrakasan
Sponsorship: DARPA

Energy harvesting is an emerging technology with applications to handheld, portable, and implantable electronics. Harvesting ambient vibration energy through piezoelectric (PE) means is a popular energy-harvesting technique that can potentially supply 10's-100's of μW of available power [1]. One of the limitations of existing PE-harvesters is in their interface circuitry. Commonly used full-bridge rectifiers and voltage doublers [2] severely limit the electrical power extractable from a PE-harvesting element. Further, the power consumed in the control circuits of these harvesters reduces the amount of usable electrical power. This work presents a bias-flip rectifier that can improve upon the power extraction capability of existing full-bridge rectifiers by greater than 4X. An efficient control circuit with embedded DC-DC converters that can share their filter inductor with the bias-flip rectifier, thereby reducing the volume and component count of the overall solution, is demonstrated.

Figure 1 shows a conventional full-bridge rectifier circuit together with the implemented bias-flip rectifier circuit. The main limitation of the full-bridge rectifier is that, even when ideal diodes are considered, most of the current available from the harvester does not go into charging the output capacitor CRECT at high values of VRECT. The shaded portion of the current waveform in Figure 1 shows the time spent in charging or discharging CP every half-cycle. This loss in charge limits the amount of electrical power that can be extracted using the full-bridge rectifier. The bias-flip rectifier consists of an inductor LSHARE that is connected in parallel with the PE-harvester. When switches M1 and M2 of the bias-flip rectifier are turned ON, the inductor helps in flipping the voltage VBF across CP. After the switches close, the PE current IP needs to supply a smaller amount of charge to CP to bring it up to ±VRECT. This reduction in charge lost significantly improves the amount of power extractable from the harvester. The inductor used in the rectifier is shared efficiently with other DC-DC converters in the system. The entire chip was fabricated in a 0.35-μm CMOS process [3]. Figure 2 shows the measured power obtained at the output of the rectifier for the different rectifier scenarios with off-chip diodes. The effectiveness of the bias-flip rectifier improves as the inductance is increased. An 820μH inductor provides a 4.2X improvement in power extracted compared to the full-bridge rectifier. This power improvement increases to above 7X when on-chip diodes are used. The DC-DC converters employed in the system achieve greater than 85% efficiency with shared inductors, in the micro-watt power levels output by the piezoelectric energy- harvester.

REFERENCES

[1] S. Roundy, P.K. Wright and J. Rabaey, Energy Scavenging for Wireless Sensor Networks with Special Focus on Vibrations, Kluwer Academic Press, 2003.

[2] T. T. Le et al., "Piezoelectric Micro-Power Generation Interface Circuits," *IEEE J. Solid-State Circuits*, pp. 1411-1420, June 2006.

[3] Y.K. Ramadass and A. Chandrakasan, "An Efficient Piezoelectric Energy Harvesting Interface Circuit Using a Bias-Flip Rectifier and Shared Inductor," *IEEE Int. Solid-State Circuits Conference*, pp. 296-297, Feb. 2009.

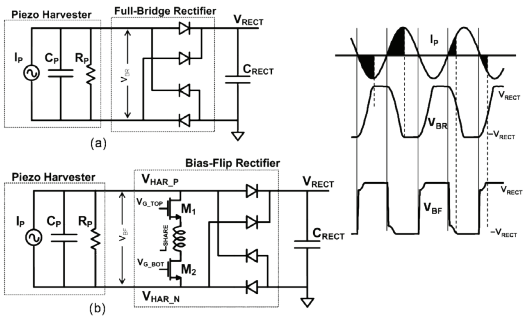

FIGURE 1: (a) Conventional full-bridge rectifier, (b) bias-flip rectifier. The right-hand side shows simulated input current and voltage waveforms for the different scenarios. The shaded portions of the current waveforms depict the amount of charge not delivered to the output capacitor C_{RECT}.

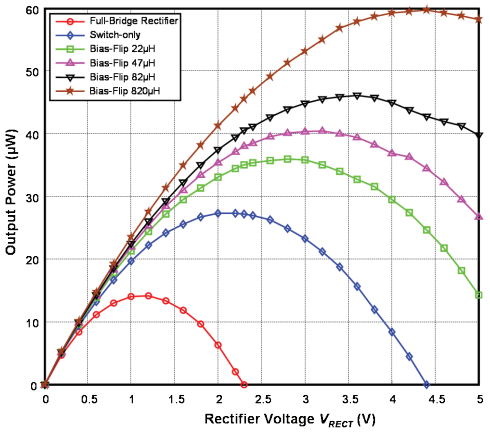

FIGURE 2: The measured electrical power output by the piezoelectric-energy-harvester as a function of the rectifier output voltage (V_{RECT}). The optimal value of V_{RECT} for maximal power transfer increases with the switch-only and bias-flip rectifier schemes.

Statistical Static Timing Analysis for Sub-0.5V Operation

R. Rithe, S. Chou, D. Buss, A. P. Chandrakasan

Statistical process variations have long been an important design issue. Until recently, process variations have been assumed to be global [1], [2]. With transistor geometries shrinking below 65nm, it is no longer valid to assume that transistor parameters are constant across a die because there are "local" or intra-die variations [1], [2]. The traditional corner-based analysis is not enough to correctly determine the performance of the integrated circuits. Despite the importance of SSTA for accurate performance analysis, the biggest challenge is to develop the statistical models and develop a computationally efficient algorithm for performing SSTA. At low voltage (V_{DD} ~0.5V), circuit delay is a non-linear function of the transistor random variables. This trait complicates the statistical analysis because the circuit delay is no longer Gaussian [3]. This work is aimed at developing such an algorithm that can perform accurate path-based SSTA in the regime where delay is a highly non-linear function of the random variables.

In this work, the dimensionality of the problem is significantly reduced by mapping the non-Gaussian delay PDF on to a Gaussian parameter through a non-linear function called Cell/Arc Delay Function. The novel concept of Operating Point (the point where joint PDF of the cell delays in a timing path attains the maxima) is introduced which allows us to determine the most probable combination of cell delays that would result into the *3-Sigma* delay (or in general *f-Sigma* delay) of the Timing Path (TP) without the need to compute the entire delay PDF.

Cell characterization is computationally efficient because we do not calculate entire delay PDF. We calculate the minimum information required to combine cell delay statistics into Timing Path delay. No expensive Monte Carlo (MC) simulation is required during characterization or timing closure. Our SSTA analysis runs in linear time with respect to number of stages, whereas MC run time increases exponentially. The concept of Operating Point that has been introduced here greatly simplifies computations despite non-linearities without sacrificing accuracy. The approach is tested on a library implemented using commercial scaled 32nm technology. We have been able to achieve accuracy of within 5% compared to MC simulation with more than 120 times improvement in run-time. The approach can be extended to perform hold-time analysis and can also be used to determine the design strategy to minimize stochastic delays.

REFERENCES

[1] A. Agarwal, D. Blaauw and V. Zolotov, "Statistical timing analysis for intra-die process variations with spatial correlations," *Computer Aided Design, 2003. ICCAD-2003. International Conference on*, pp. 900-907, 2003.

[2] B. Liu, "Gate Level Statistical Simulation Based on Parameterized Models for Process and Signal Variations," *Quality Electronic Design, 2007. ISQED '07. 8th International Symposium on*, pp. 257-262, 2007.

[3] L. Cheng, J. Xiong and L. He, "Non-Linear Statistical Static Timing Analysis for Non-Gaussian Variation Sources," *Design Automation Conference, 2007. DAC '07. 44th ACM/IEEE*, pp. 250-255, 2007.

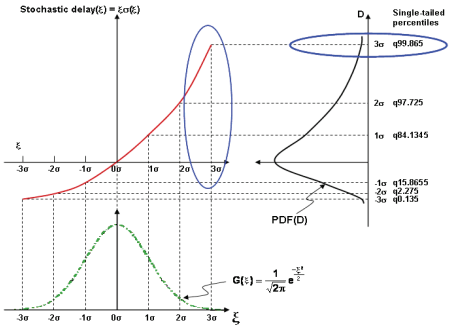

FIGURE 1: Mapping of non-Gaussian delay PDF on to a Gaussian parameter ξ through the non-linear Cell/Arc Delay Function. The CADF is stored as the output of cell characterization.

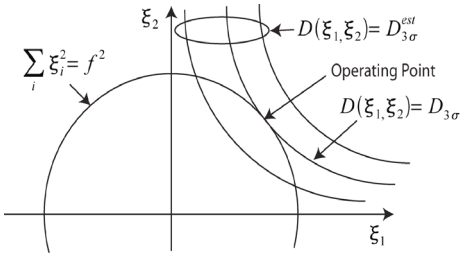

FIGURE 2: The operating point for a 2-stage TP that gives f-Sigma delay at the output. The operating point allows us to easily determine the stochastic delay imparted by each cell by referring to the stored CADFs without having to look into the transistor RVs.

An 8T Reconfigurable SRAM in 65-nm CMOS Achieving 0.25-1.2V Operating Voltage Range

M. E. Sinangil, A. P. Chandrakasan
Sponsorship: DARPA

In modern integrated circuits (ICs), the trend of integrating more on-chip memories on a die has led SRAMs to account for a large fraction of total chip area as well as the total chip energy. Hence, applying energy-saving schemes such as dynamic voltage scalability (DVS) to SRAMs is an important research area. However, optimizing circuit operation over a large voltage range is not trivial due to conflicting trade-offs of low-voltage (moderate and weak inversion) and high-voltage (strong inversion) transistor characteristics. Specifically, low-voltage operation requires various assist circuits for functionality, which might severely impact high-voltage performance. Previous work in literature focused on either sub-threshold operation [1] or DVS in only the above-threshold regime [2]. In this work [3], SRAM design for a very large voltage range including both sub-threshold and above-threshold regions is investigated. Reconfigurable circuit assists are proposed as a solution to the problem of optimizing circuits over a large voltage range with minimum performance penalty and power overhead.

A test chip is implemented in 65-nm low-power CMOS process to demonstrate an ultra-dynamic voltage scalable (U-DVS) SRAM. The 64-kbit memory array is structured in 8 blocks, each with 64 rows and 128 columns of memory cells. Figure 1 shows the die photo of the test chip. Read and write functionality is achieved from 1.2V down to 0.25V at 200MHz and 20kHz, respectively, as shown in Figure 2. Over this voltage range, leakage power scales down by more then 50x. The minimum energy point of this work lies close to 0.4V at less than 0.1pJ/bit/access.

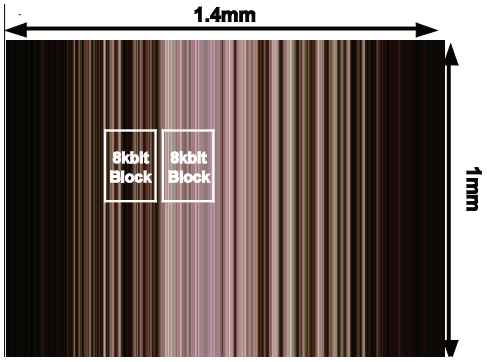

FIGURE 1: Die photo of the 64kbit SRAM test chip fabricated in 65-nm CMOS process. Die area is 1mm x 1.4mm.

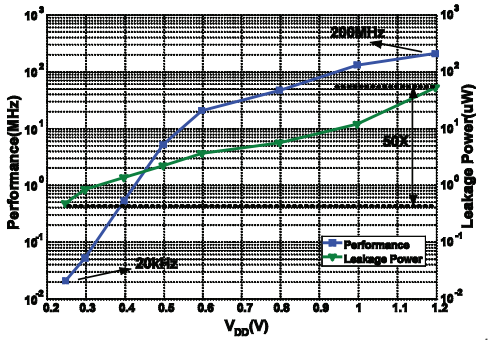

FIGURE 2: Measured performance and leakage power plots for the 64-kbit SRAM array. Over the 0.25-1.2V operating voltage range, leakage power scales by more than 50X.

REFERENCES

[1] N. Verma and A. P. Chandrakasan, "A 65nm 8T sub-vt SRAM employing sense-amplifier redundancy," in *IEEE International Solid-State Circuits Conference, Digest of Technical Papers,* 2007, pp. 328-606.

[2] L. Chang, Y. Nakamura, R. K. Montoye, J. Sawada, A. K. Martin, K. Kinoshita, F. H. Gebara, K. B. Agarwal, D. J. Acharyya, W. Haensch, K. Hosokawa and D. Jamsek, "A 5.3GHz 8T-SRAM with operation down to 0.41V in 65nm CMOS," in *IEEE Symposium on VLSI Circuits,* 2007, pp. 252-253.

[3] M. E. Sinangil, N. Verma and A. P. Chandrakasan, "A reconfigurable 65nm SRAM achieving voltage scalability from 0.25–1.2V and performance scalability from 20kHz–200MHz," in *European Solid-State Circuits Conference,* 2008, pp. 282-285.

A Low-Power 0.7-V H.264/AVC 720p Video Decoder

V. Sze, D. Finchelstein, A. P. Chandrakasan
Sponsorship: Nokia, Texas Instruments

The H.264/AVC video coding standard can deliver high compression efficiency at a cost of large complexity and power [1]. The increasing popularity of video capture and playback on portable devices requires that the power of the video codec be kept to a minimum.

This work proposes several architecture optimizations such as increased parallelism, multiple voltage/frequency domains, and custom voltage-scalable SRAMs that enable low voltage operation to reduce the power of a high-definition decoder. Additionally, this work highlights the importance of on-chip caching in reducing the power and performance penalty of accessing a large off-chip memory. Dynamic voltage and frequency scaling can efficiently adapt to the varying workloads by leveraging the low voltage capabilities and domain partitioning of the decoder. The decoder architecture, shown in Figure 1, is organized as a pipeline that processes 4x4 luma and chroma blocks in parallel [2]. The pipeline units are decoupled by variable-length FIFOs, in order to average out the variable latency of the units, and minimize pipeline stall cycles.

An H.264/AVC Baseline Level 3.2 decoder ASIC was fabricated in 65-nm CMOS and verified. For high definition 720p video decoding at 30 frames per second, it operates down to 0.7 V and has a measured power of 1.8 mW, which is over an order of magnitude lower than previously published results, as shown in Figure 2. The decoder highly scalable and is capable of operating down to 0.5 V for decoding QCIF at 15 frames per second with a measured power of 30 μW. During P-frames (temporally predicted), the decoder power is dominated by the motion compensation (MC) and deblocking filter (DB). The on-chip SRAM caches take up 75% of the active die area, while the parallelism proposed less than 3% area overhead.

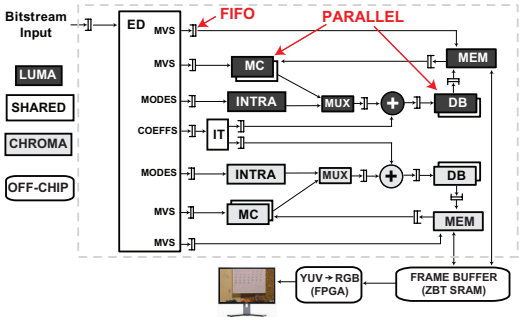

FIGURE 1: H.264 decoder architecture. Pixels are processed using a 4x4 pipeline, separated by FIFOs to average out workload. Luma and chroma pipelines run concurrently, with the exception of ED and IT.

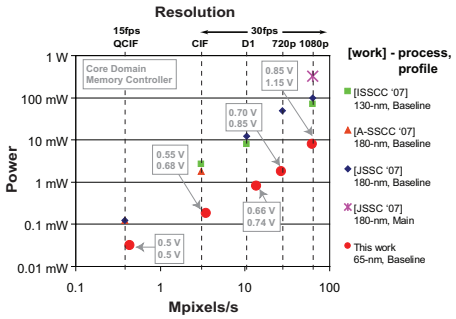

FIGURE 2: Comparison of this chip with other publications. At 720p, our decoder uses 13x lower power than the previous low, though in a more advanced technology (65nm vs. 130nm). Our decoder can scale from QCIF@15 to 1080p@30, with the core voltage scaling from 0.5V to 0.85V.

REFERENCES

[1] J. Ostermann, Bormans, P. List, D. Marpe, M. Narroschke, F. Pereira, T. Stockhammer, and T. Wedi, "Video coding with H.264/AVC: Tools, Performance, and Complexity," *IEEE Circuits and Systems Magazine*, vol. 4, pp. 7– 28, 2004.

[2] D. Finchelstein, V. Sze, M. Sinangil, Y. Koken, A. Chandrakasan, "A Low-power 0.7-V H.264 720p Video Decoder," in *IEEE Asian Solid State Circuits Conference*, November 2008.

A Micro-Power EEG Acquisition SoC with Integrated Seizure Detection Processor for Continuous Patient Monitoring

N. Verma, A. Shoeb, J. Guttag, A. P. Chandrakasan

Sponsorship: Intel Foundation Ph.D. Fellowship Program, CICS, NSERC, National Semiconductor

Epilepsy, a neurological disorder affecting 50 million people worldwide, causes sudden seizures that result in convulsions, loss of coherence, or even death. Seizure detection, before the onset of these symptoms, can improve the lives of patients tremendously by providing an early warning to them and their caregivers, or by triggering therapy to stop the seizure. Early detection, however, requires sophisticated processing to separate normal and abnormal neural activity, which varies greatly from patient-to-patient.

We present a system-on-chip, integrating an ultra-low-power instrumentation amplifier, ADC, and digital processor [1]. The chip continuously senses a patient's neural firings through non-invasive electrodes on the scalp (i.e. EEG). The neural signals are processed to extract the subtle information necessary to detect seizure onset. Seizures are then detected through further processing using a machine-learning classifier. By compressing the neural information, it can be transmitted wirelessly with 14x lower system power, eliminating hazardous cables from the patient's scalp. The detection algorithm has been tested through 536 hours of patient tests [2], and the chip consumes less than 10 micro-Watts/channel (depending on the patient, up to 18 channels may be used).

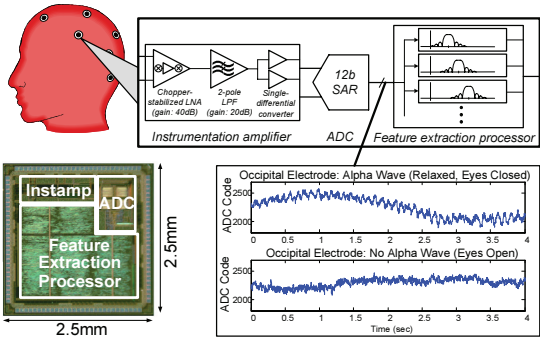

FIGURE 1: Chip block diagram showing instrumentation amplifier, ADC, and seizure feature extraction processor; measured EEG are also shown, sensing neural firings of relaxed eyes closed state.

REFERENCES

[1] N. Verma, A. Shoeb, J. Guttag, and A. Chandrakasan, "A Micro-power EEG Acquisition SoC with Integrated Seizure Detection Processor for Continuous Patient Monitoring," *VLSI Symposium of Circuits*, 2009.

[2] A. Shoeb, B. Bourgeois, S. T. Treves, S. C. Schachter and J. Guttag, "Impact of Patient-Specificity on Seizure Onset Detection Performance," *Engineering in Medicine and Biology Society, 2007. EMBS 2007. 29th Annual International Conference of the IEEE*, pp. 4110-4114, 2007.

Stable Model Reduction: A Semi-definite System-identification Approach for Nonlinear Systems

B. Bond, Z. Mahmood, A. Megretski, L. Daniel
Sponsorship: DARPA, SRC/FCRP IFC

During recent years, a great effort has been made by researchers of the Electronic Design Automation community to develop new techniques for automatically generating accurate compact models of nonlinear system blocks. The majority of existing methods for creating stable reduced models of nonlinear systems, such as [1], require knowledge of the internal structure of the system, as well as access to the exact model formulation for the original system. Unfortunately, this information may not be available if a designer is using a commercial design tool, or may not even exist if the system to be modeled is a physical fabricated device.

As an alternative approach to nonlinear model reduction, we have proposed a system-identification procedure. This procedure requires only data available from simulation or measurement of the original system, such as input-output data pairs. By enforcing incremental passivity, as shown in [2], it is possible to formulate a semi-definite optimization problem whose solution is a stable nonlinear model that optimally matches the given data pairs from the original system. In addition, the proposed optimization formulation allows us to specify completely the complexity of the identified reduced model through the choice of both model order and nonlinear function complexity.

Applications for the proposed modeling technique include analog circuit building blocks such as operational amplifiers and power amplifiers, MEMS devices, and individual circuit elements such as transistors. The resulting compact models may then be used in a higher-level design optimization process of a larger system. One such example of an analog circuit block is the low-noise amplifier shown in Figure 1; it contains both nonlinear and parasitic elements. For this example, input-output training data was generated from a commercial circuit-simulator and used to identify a compact nonlinear model. The output responses of the original system and the identified model are compared in Figure 2.

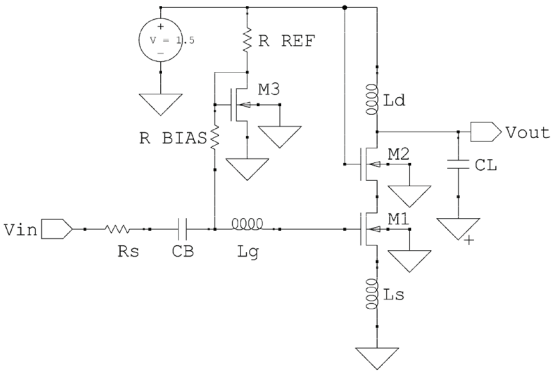

FIGURE 1: Application example: Low-noise amplifier designed in CMOS technology.

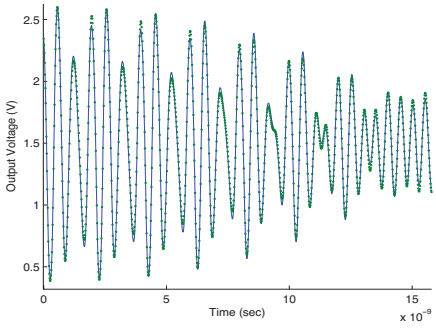

FIGURE 2: Comparison of the output response from a commercial circuit simulator (solid blue line) and the output response from a stable nonlinear reduced model created with the proposed approach (green dots).

REFERENCES

[1] B. Bond and L. Daniel, "Stabilizing schemes for piecewise-linear reduced order models via projection and weighting functions," *Proc. of the IEEE Conference on Computer-Aided Design*, San Jose, CA, Nov. 2007, pp. 860-867.

[2] A. Megretski, "Convex optimization in robust identification of nonlinear feedback," *in Proc. of the IEEE Conference on Decision and Control*, Cancun, Mexico, Dec. 2008, pp. 1370-1374.

Stable Model Reduction: A Projection Approach for Indefinite and Possibly Unstable LINEAR Systems

B. Bond, L. Daniel
Sponsorship: SRC/FCRP IFC

Although stable model reduction for linear systems has been an extremely popular topic for many years and has generated many useful results, there is still a need for more robust and efficient techniques. In [1], a very efficient model-reduction technique preserves stability and passivity only for systems described by system matrices with a particular structure. In [2], a stability-preserving technique is presented for unstructured stable linear systems, but the method is computationally prohibitively expensive. Additionally, existing methods focus only on preserving stability, when in fact it is sometimes necessary to create stability through the model-reduction process. It is not uncommon, for instance, to obtain large unstable models of stable physical systems from field-solvers, as a result of numerical error arising from discretization.

In our work, we have developed a stability-preserving projection framework for model reduction of linear systems that is cheap in terms of computation complexity, puts no requirements on the structure of the original system, and can even generate accurate stable reduced models from originally unstable models of stable physical systems. Specifically, given a right-projection matrix, we derive a set of linear constraints for the left-projection matrix resulting in a projection framework that is guaranteed to generate a stable and passive reduced model. We formulate the problem of computing the stabilizing projection matrix as a semi-definite program that can be solved efficiently using existing techniques, resulting in an optimal stabilizing projection framework. Details of this procedure can be found in [3].

Our algorithms have been tested on a large variety of typical VLSI applications, including field-solver-extracted models of RF inductors for analog applications, power-distribution grids for large VLSI digital integrated circuits, and MEMS devices for sensing and actuation applications. One such application, a spiral RF inductor, is shown in Figure 1. For this example, an electro-magneto-quasi-static (EMQS) field solver was first used to extract a large linear system of order 647 (which turned out to be slightly numerically unstable). The proposed optimization-based method was then used to compute a stabilizing projection framework, finally resulting in a stable 10th order reduced model. Figure 2 plots the quality factor of the large extracted model (solid blue line) and the stable reduced model (red crosses) generated with the proposed approach.

REFERENCES

[1] A. Odabasioglu, M. Celik, and L.T. Pileggi, "PRIMA: passive reduced order interconnect macromodeling algorithm," in *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, vol. 17, no. 8, pp. 645-654, Aug. 1998.

[2] B. Moore, "Principal component analysis in linear systems: controllability, observability, and model reduction," *IEEE Transactions on Automatic Control*, vol. 26, no. 1, pp. 17-32, Feb. 1981.

[3] B. Bond and L. Daniel, "Guaranteed stable projection-based model reduction for indefinite and unstable linear systems," *Proc. of the IEEE Conference on Computer-Aided Design*, San Jose, CA, Nov. 2008, pp. 728-735.

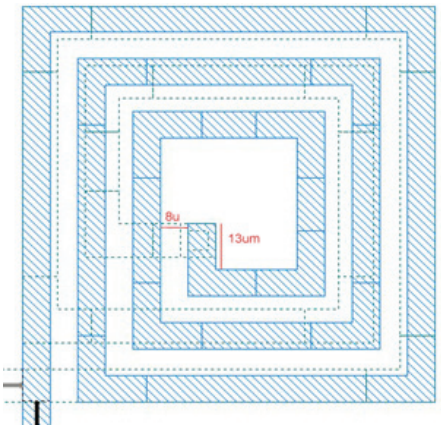

FIGURE 1: Application example: Spiral RF inductor.

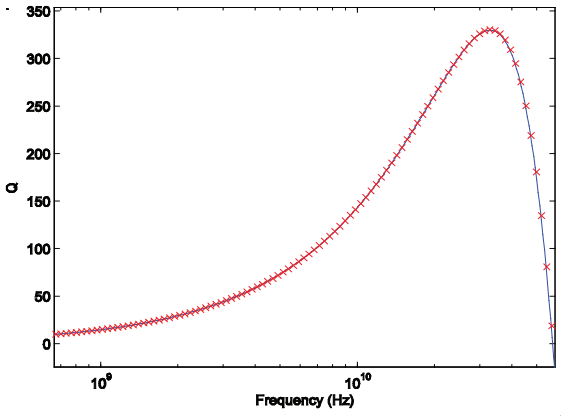

FIGURE 2: Comparison of quality factor for order 647 model extracted from EMQS field solver (solid blue line) and a 10th-order reduced model created with the proposed approach (red crosses).

Stable Model Reduction: A Projection Approach for NONLINEAR Systems in the More General Descriptor Form

B. Bond, L. Daniel
Sponsorship: SRC/FCRP IFC

The ability to generate accurate reduced-order models (ROMs) of nonlinear dynamical systems, such as analog circuits and micro-electro-mechanical systems (MEMS), is a crucial first step in the automatic design and optimization of such systems. One popular approach to model order reduction (MOR) of highly nonlinear systems employs trajectory-based methods, such as the piecewise-linear (PWL) approach. Despite substantial recent interest in such methods [1], [2], trajectory-based models (TBMs) have failed to gain widespread acceptance due to a lack of theoretical statements concerning the accuracy of the resulting ROMs. In this work we address one such theoretical issue: guaranteed stability. Specifically, we present a scheme for preserving stability in PWL models, whose system matrices possess a certain structure. We also propose a projection scheme that allows us to extend some of these stability results to PWL systems composed of arbitrary unstructured matrices.

The stability of nonlinear systems can be certified for instance by the existence of a Lyapunov function. Our stabilizing scheme ensures stability by constructing the projection matrices such that there exists a Lyapunov function for the resulting ROM. In the case where all the Jacobians of the linearized systems possess a certain structure, examples of which are given in [3], we present a projection routine that guarantees the existence of a quadratic Lyapunov function for both the large PWL model and the ROM. In the case where the system's Jacobians have no structure and it is not known whether a Lyapunov function exists for the large PWL model, we utilize a new nonlinear projection to create a collection of stable reduced local models. The resulting nonlinear model is guaranteed to be at least locally stable. One example of a system that produces unstructured Jacobians, and thus potentially unstable TBMs, is a MEMS switch (shown in Figure 1). Figure 2 shows a sample output from the MEMS switch, a stable TBM generated by our approach, and an unstable TBM generated by the traditional approach. For further details on the stabilizing procedure, see [3].

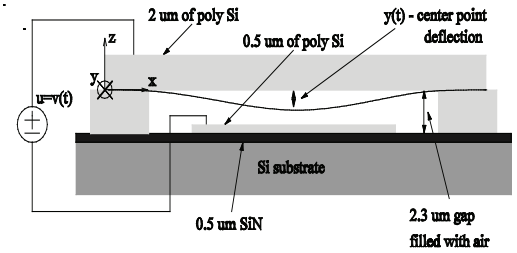

FIGURE 1: Application example: MEM switch realized by a polysilicon beam fixed at both ends and suspended over a semiconducting pad and substrate expansion.

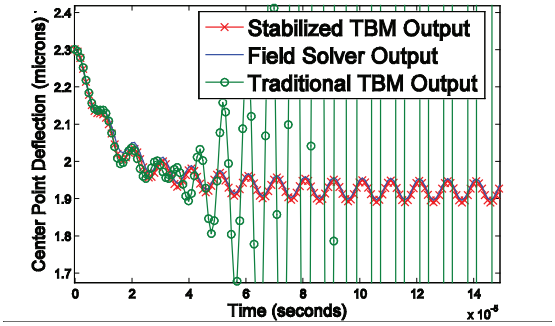

FIGURE 2: Center-point deflection predicted by our stabilized reduced model (red crosses), compared to a finite-difference detailed simulation (solid blue lines) and the traditional TBM approach (green circles).

REFERENCES

[1] M. Rewienski and J.K. White, "A trajectory piecewise-linear approach to model-order reduction and fast simulation of nonlinear circuits and micromachined devices," in *Proc. of IEEE/ACM International Conference on Computer Aided-Design*, San Jose, CA, Nov. 2001, pp 252-257.

[2] B. Bond and L. Daniel, "A piecewise-linear moment-matching approach to parameterized model-order reduction for highly nonlinear systems," *IEEE Trans. on Computer-Aided Design*, vol. 26, no. 12, pp. 2116-2129, Dec. 2007.

[3] B. Bond and L. Daniel, "Stabilizing schemes for piecewise-linear reduced order models via projection and weighting functions," *Proc. of the IEEE Conference on Computer-Aided Design*, San Jose, CA, Nov. 2007, pp. 860-867.

Efficient Capacitance Solver for 3D Interconnect Based on Template-instantiated Basis Functions

Y-C. Hsiao, L. Daniel
Sponsorship: SRC/FCRP IFC, Mentor Graphics, AMD

Integrated circuit performance and signal integrity can be largely affected by interconnect parasitic capacitance, and they require fast and accurate extraction tools. Satisfying simultaneously both constraints is, however, an extremely challenging task. The current state-of-the-art in efficient extraction methods involves 2D cross-section scanning, determining wire adjacency, calculating 2D capacitance in a table lookup approach, and then reconstructing quasi-3D capacitance. Such an approach is indeed fast, yet it is accurate only for 2D structures. Full 3D structures (e.g., crossing wires in adjacent metal layers) need the accuracy of field solvers such as FastCap [1] and Precorrected FFT [2]. Such tools are based on piece-wise constant basis functions and are accelerated by fast matrix-vector products that have a significant computational overhead but scale almost linearly with the number of conductors. Hence they are ideal for very large-scale examples.

This project is targeted instead at efficient small-to-medium scale capacitance extraction. The key idea is to exploit the highly restrictive design rules of the recent sub-micro to nano-scale technologies. In this scenario, a limited number of pre-computed surface charge distributions can be used as a set of fundamental template basis functions. Figure 1 shows an example of charge distribution "stretchability," enabling the instantiation of basis functions for every practical interconnect structure. Using a total of just 72 template-instantiated basis functions, the example in Figure 2 demonstrates a worst-case relative error of 3% with respect to the total capacitance of each conductor, compared to the result extracted by FastCap in a very fine discretization with tens of thousands of unknowns. FastCap requires 732 unknowns to produce the same 3% error in a coarser discretization. Hence, for the same 3% accuracy, our algorithm requires approximately 10x fewer unknowns. In such medium-size examples, the overhead of the FastCap multipole expansion makes the linear acceleration ineffective, while in our approach, analytical formulas and numerical tabulation of the Galerkin coefficients for our template basis functions can effectively limit the setup overhead, producing a *two-orders-of-magnitude* improvement in both simulation time and memory requirement.

REFERENCES

[1] K. Nabors and J. White, "FastCap: A Multipole-Accelerated 3-D Capacitance Extraction Program," *IEEE Transactions on Computer-Aided Design*, vol.10, no. 10, pp. 1447-1459, November 1991.

[2] J.R. Phillips and J. K. White, "A Precorrected-FFT method for Electrostatic Analysis of Complicated 3-D Structures," *IEEE Trans. on Computer-Aided Design*, vol. 16, no.10, p.059-1072, October 1997.

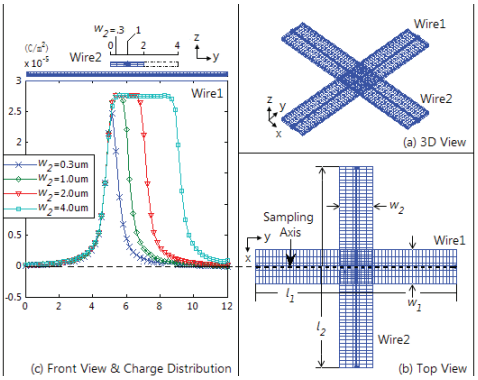

FIGURE 1: A pair of crossing wires: (l_1, w_1, l_2, w_2) = (12, 2, 12, 2) (um). (a) 3D view, (b) Top view (x-y plane), (c) Front view (y-z plane), and charge distribution sampled along the sampling axis of Wire1 shown in (b) as a broken line. Note the charge distributions are "stretchable" with w_2 while preserving decaying shapes off the edge of Wire2.

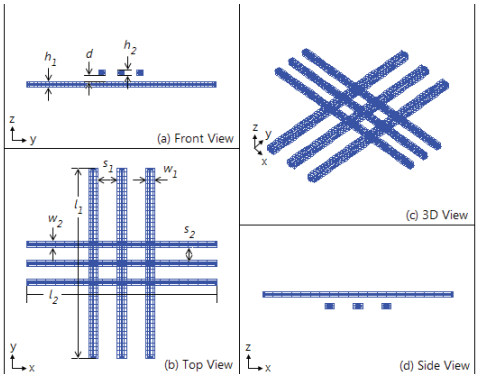

FIGURE 2: Six wires in two metal layers: (l_1, w_1, h_1, s_1 | l_2, w_2, h_2, s_2 | d) = (10, 0.2, 0.3, 0.2 | 10, 0.2, 0.3, 0.2 | 0.3) (um). (a) Front view (y-z plane), (b) Top view (x-y plane), (c) 3D view, and (d) Side view (x-z plane). The width and spacing are exaggerated for clarity.

Variation-aware Parasitic Extraction: A Deterministic Approach

T. A. El-Moselhy, L. Daniel
Sponsorship: SRC/FCRP IFC, IBM

On-chip and off-chip fabrication processes may typically generate interconnect structures of irregular geometries. Such irregularities are not deterministic and are produced by several different manufacturing steps such as etching, chemical mechanical polishing (CMP), electro-deposition, and photolithography. However, as a result of technology scaling, such manufacturing uncertainties significantly affect the electrical characteristics of the interconnect structures. The effect of such variations on the electrical characteristics can be efficiently extracted using what we call "variation-aware" parasitic extraction tools. Such solvers can in general be divided into two categories, namely, those based on deterministic algorithms and those based on stochastic algorithms.

In this project we have developed new ***deterministic*** variation-aware extraction algorithms based on the well-known floating random walk (FRW) algorithm [1]. First, we have presented a new finite-difference-based sensitivity analysis [2] within the improved FRW algorithm to efficiently compute capacitance sensitivities with respect to a large number of small parameter variations. We have demonstrated that the expected complexity of computing the nominal capacitance and all the sensitivities is less than 2 times that of computing only the nominal capacitance regardless of the number of parameters. The complexity of our sensitivity algorithm is therefore independent of the number of varying parameters (unlike standard finite difference sensitivity analysis) and independent of the number of output capacitances (unlike standard adjoint sensitivity analysis [3]). Second, we have presented a new incremental FRW algorithm [2] to efficiently compute the capacitances of similar geometrical configurations resulting from simultaneous large perturbations of the geometrical parameters of a common geometrical topology. The new algorithm satisfies a major objective of variation-aware parasitic extraction, namely, that the average time required to solve a single geometrical configuration within a set of similar configurations is reduced as the cardinality of the set is increased. We have observed that the average simulation time of a single configuration from a set of similar configurations of cardinality 100,000 is reduced by three orders of magnitude. Consequently, we were able to solve more than 130,000 similar configurations in the time required to solve just 50 independent configurations. We believe that the latter result will naturally fit in a litho- and CMP-aware extraction flow.

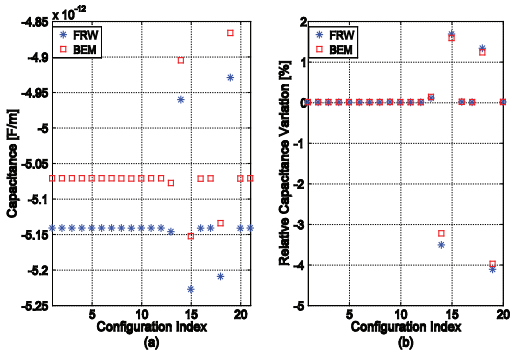

FIGURE 1: (a) Capacitance, (b) Relative capacitance variation, i.e., the difference between the capacitance of any configuration and that of the nominal configuration divided by the capacitance of the nominal, as computed using our FRW sensitivity algorithm and the standard boundary element method (BEM) for different geometrical configurations.

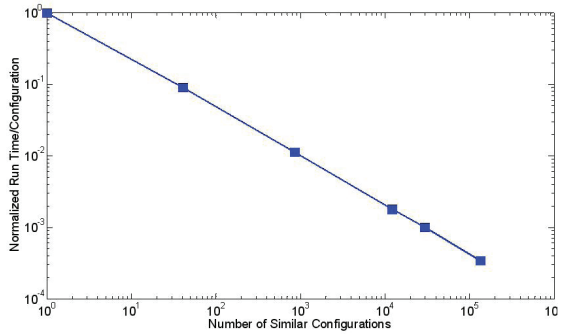

FIGURE 2: Normalized average simulation time for a single geometry (configuration) in a set of similar geometries versus the total number of configurations in the set.

REFERENCES

[1] R. B. Iverson and Y. L. Le Coz, "A Stochastic Algorithm for High Speed Capacitance Extraction in Integrated Circuits" *Solid-State Electronics*, vol. 35, no. 7, pp. 1005-1012, 1992.

[2] T. Moselhy, I. M. Elfadel and L. Daniel, "A Capacitance Solver for Incremental Variation-Aware Extraction," in *Proc. IEEE/ACM International Conference on Computer Aided Design*, 2008.

[3] J. Wang and J. White, "Fast Algorithms for Computing Electrostatic Geometric Sensitivities" in *Proc. International Conference on Simulation of Semiconductor Processes and Devices*, pp. 121-123, 1997.

Variation-aware Parasitic Extraction: A Stochastic Approach

T. A. El-Moselhy, L. Daniel
Sponsorship: SRC/FCRP IFC, IBM

On-chip and off-chip fabrication processes may typically generate interconnect structures of irregular geometries. Such irregularities are not deterministic and are produced by several different manufacturing steps such as etching, chemical mechanical polishing (CMP), electro-deposition, and photolithography. However, as a result of technology scaling, such manufacturing uncertainties significantly affect the electrical characteristics of the interconnect structures. The effect of such variations on the electrical characteristics can be extracted using what we call "variation-aware" extraction tools. Such solvers can in general be divided into two categories, namely, those based on deterministic algorithms and those based on stochastic algorithms.

In this research we have developed a new ***stochastic*** algorithm to solve large stochastic linear systems typically appearing during variation-aware extraction [1]. We have derived a new theorem to compute the coefficients of the multivariate Hermite expansion using only low dimensional integrals, resulting in a time complexity that is independent of the number of variables and dependent only on the order of the expansion. Practically speaking, for a typical large multivariate expansion, the new theorem provides an improvement in the computation time by 86 orders of magnitude as compared to the standard tensor product rule or by 10 orders of magnitude as compared to the state of the art (Monte Carlo integration or sparse grid integration [2]). Such a theorem is not only useful for our methodology but it can also be applied to any algorithm that relies on expanding a random process, such as the stochastic finite element method (SFE) [3]. We have also provided a new stochastic simulation technique by merging both the Neumann expansion and the polynomial chaos expansion. The main advantages of the resulting technique are the compact size of the system (unlike SFE) and the ease of calculating the statistics of the high order terms (unlike Neumann expansion [4]). In addition, the new simulation algorithm is parallelizable and can therefore take advantage of the multicore platforms readily available in recent processor technologies. We have demonstrated the computational efficiency of the new methodology by solving problems that were completely intractable before. We have demonstrated that our algorithm can be used to compute the complete probability density function of the input impedance of very large problems (up to 400 random variables) in less than 8 hours using Matlab on a standard 4-core machine and using only 121MB RAM.

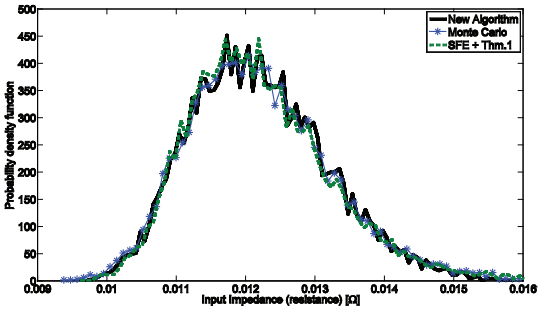

FIGURE 1: Comparison between the probability density function of the microstrip line obtained from our new algorithm and the reference Monte Carlo simulation.

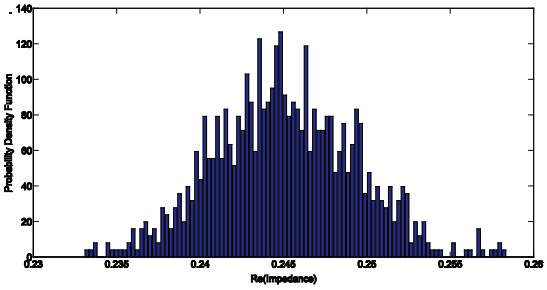

FIGURE 2: Probability density function of the real part of the input impedance at 1GHz for correlation length Lc = 50μm. The resistance of the non-rough surface is 11.3% smaller than the mean of the obtained distribution.

REFERENCES

[1] T. Moselhy and L. Daniel, "Stochastic Integral Equation Solver for Variation-Aware Interconnect Extraction," in *Proc. IEEE/ACM Design Automation Conference,* 2008.

[2] D. Xiu and G. Karniadakis, "The Wiener-Askey Polynomial Chaos Expansion For Stochastic Differential Equations," *SIAM Journal of Scientific Computing*, vol. 24, pp. 619-644, Oct. 2002.

[3] R. Ghanem and P. Spanos, *Stochastic Finite Elements: A Spectral Approach*. Spring-Verlag, 1991.

[4] Z. Zhu and J. White, "FastSies: A Fast Stochastic Integral Equation Solver for Modeling the Rough Surface Effect," in *Proc. IEEE/ACM International Conference on Computer Aided Design,,* pp. 675-682, Nov. 2005.

Digital-to-RF-Phase Converter for AMO Architecture

T. W. Barton, S. Chung, P. A. Godoy, D. J. Perreault, J. L. Dawson
Sponsorship: Desphande Center for Technological Innovation, SiGe Semiconductor

This work presents a digital to RF phase converter (DRFPC) designed specifically for the Asymmetric Multilevel Outphasing (AMO) transmitter architecture [1]. The AMO architecture, shown in Figure 1, has great advantages in efficiency over similar architectures such as Multilevel LINC (linear amplification with non-linear components) without loss of linearity. By allowing the power amplifiers (PA) supplies to vary independently between multiple levels, the DRFPC reduces the burden on outphasing as a way to control output signal amplitude. Instead, outphasing is used primarily for pulse shaping. This use, along with the abrupt phase-changes required when the amplitude path steps between discrete levels, translates to a requirement for a high-speed phase path.

The DRFPC presented in this work is designed to perform the phase modulation required for the AMO phase path. Because of the requirement that this phase path be very wideband, a closed loop approach such as a phase locked loop (PLL) is impractical. Instead, a differential current-steering topology is the preferred approach, in which weighted quadrature signals are added to produce an arbitrary output angle [2], [3]. The design choice made in this work exploits its use in the AMO architecture. In particular, the digital predistortion used in the transmitter allows for reduced linearity requirements in the design and therefore creates a topology that is relatively low-area and low-power. For example, by constraining the possible inputs to the DRFPC so that the amplitudes of the quadrature signals always sum to one, it is possible to use only one set of binary weighted current sources. This constraint results in a nonlinearity in the DRFPC but nearly halves its area and reduces matching requirements as compared to a generalized approach. A simplified DRFPC schematic is shown in Figure 2.

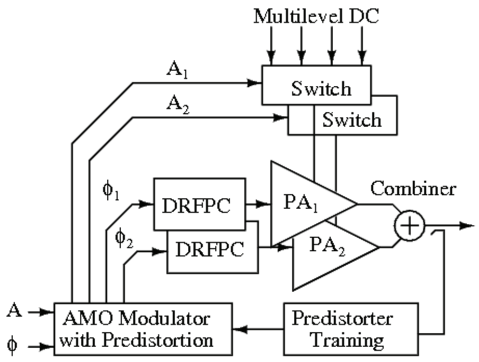

FIGURE 1: The AMO architecture. The ability in the amplitude path to independently switch between discrete voltages moves the outphasing focus from amplitude control to pulse shaping.

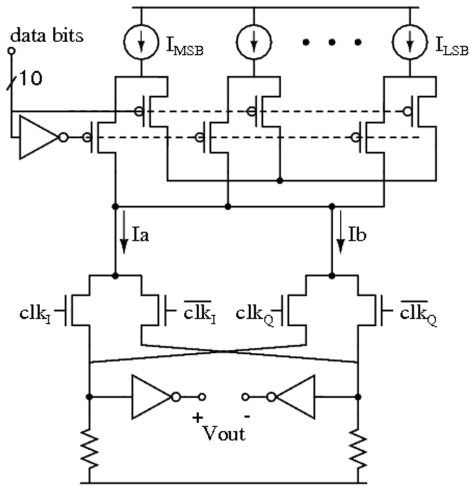

FIGURE 2: Digital to RF phase converter (DRFPC) schematic. The AMO digital predistortion allows for the use of only one set of current sources, nearly halving the required area.

REFERENCES

[1] S. Chung, P. A. Godoy, T.W. Barton, E. W. Huang, D. J. Perreault, and J. L. Dawson, "Asymmetric Multilevel Outphasing Architecture for multi-standard transmitters," *Radio Frequency Integrated Circuits Symposium, 2009. RFIC 2009. IEEE*, 2009.

[2] P. Eloranta, P. Seppinen, S. Kallioinen, T. Saarela and A. Parssinen, "A Multimode Transmitter in 0.13 µm CMOS Using Direct-Digital RF Modulator," *Solid-State Circuits, IEEE Journal of*, vol. 42, pp. 2774-2784, 2007.

[3] M. E. Heidari, M. Lee and A. A. Abidi, "All-Digital Outphasing Modulator for a Software-Defined Transmitter," *Solid-State Circuits, IEEE Journal of*, vol. 44, pp. 1260-1271, 2009.

Digitally-Assisted Analog Front-End for Biomedical Sensors

J. L. Bohorquez, M. Yip, A. P. Chandrakasan, J. L. Dawson

Sponsorship: Texas Instruments, DARPA

Biomedical sensors are used to measure a myriad of biopotential signals including electroencephalogram (EEG), electrocardiogram (EKG), electromyogram (EMG), and neural field potential (NFP) signals [1], [2], [3]. Most of the useful information in these signals resides in the frequency range of 0.5 Hz to 1 kHz, allowing ultra-low power circuits to be used when processing them. This is critical for systems that are implanted, since energy is extremely scarce, and the lifetime of the device must be on the order of 10 years. Unfortunately, these signals are often as small as 10 μVs, and their low frequency location make them vulnerable to aggressors such as DC offset, powerline noise, and flicker noise. DC offset can result from charge accumulation at the interface between the metal electrodes and the skin, and also from amplifier offsets caused by random mismatches. While chopper stabilization has proved effective at mitigating the effects of amplifier DC offset and flicker noise, electrode DC offset cannot be removed through chopping and must be high-pass filtered at the front end of the system to prevent saturation [1], [2], [3]. Powerline noise, typically at 50 or 60 Hz, is mostly a common-mode signal that requires adequate common-mode rejection. However, if there are mismatches or inductive loops in the electrodes, these aggressors can become differential-mode signals, corrupting the desired signal, and potentially saturating the system. In closed-loop deep brain stimulation systems, another aggressor arises from stimulation artifacts [4]. In that case, the NFPs can be much smaller than stimulation artifacts placing stringent requirements on the dynamic range of the system and potentially leading to signal corruption.

We propose a mixed-signal sensor interface that mitigates the effects of all of the aforementioned aggressors in an area efficient manner. Area efficiency is particularly compelling in implantable devices that use tens or hundreds of electrodes, such as neural recording systems [3]. The proposed system, shown in Figure 1, uses a chopper stabilized operational amplifier with capacitive feedback to achieve accurate gain (The system is shown as single-ended for simplicity, but is implemented in a fully differential manner). Figure 2 shows a simplified schematic of the amplifier, including a novel input chopper that creates a switched capacitor resistance between its inputs and a reference voltage. This resistance is shown as Rp in Figure 1 and is used to create a high-pass filter with a corner frequency well below 1 Hz, while setting the common-mode voltage of the input to a desired level. The pole frequency is actually set by the Miller-multiplied feedback capacitor Cf and is inversely proportional to the amplifier's gain A_v, allowing a reduction of many orders of magnitude in component sizes. An additional feedback path is introduced that includes the filter, ADC, DSP, and a feedback DAC. This path can be used to notch out unwanted signals such as powerline noise or stimulation artifacts before they can saturate the system.

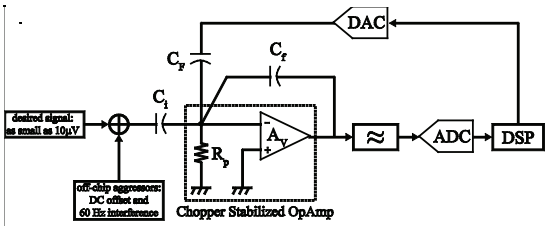

FIGURE 1: Biomedical sensor interface system employing chopper stabilization, and mixed signal feedback to mitigate the effects of DC offset, flicker noise, and other aggressors.

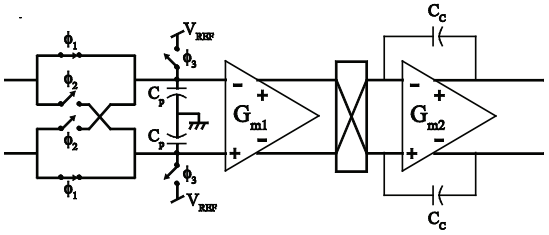

FIGURE 2: Simplified schematic of a chopper stabilized operational amplifier that exploits the parasitic switched-capacitor conductance to implement a high-pass filter.

REFERENCES

[1] T. Denison, K. Consoer, W. Santa, A. Avestruz, J. Cooley, and A. Kelly, "A 2μW 100nV/rtHz chopper-stabilized instrumentation amplifier for chronic measurement of neural field potentials," *IEEE J. Solid-State Circuits*, vol. 42, no. 12, pp. 2934–2945, Dec. 2007.

[2] R. F. Yazicioglu, P. Merken, R. Puers, and C. V. Hoof, "A 200μW eight-channel EEG acquisition ASIC for ambulatory EEG systems," *IEEE J. Solid-State Circuits*, vol. 43, no. 12, pp. 3025–3038, Dec. 2008.

[3] R. R. Harrison *et al.*, "A low power integrated circuit for a wireless 100-electrode neural recording system," *IEEE J. Solid-State Circuits*, vol. 42, no. 1, pp. 123–133, Jan. 2007.

[4] J. Lee, H.–G. Rhew, D. Kipke, and M. Flynn, "A 64 channel programmable closed-loop deep-brain stimulator with 8 channel neural amplifier and logarithmic ADC," in *Symp. VLSI Circuits Dig.*, Honolulu, HI, Jun. 2008, pp. 76–77.

Asymmetric Multilevel Outphasing Architecture for Multi-standard Transmitters

S. Chung, P. A. Godoy, T. W. Barton, Z. Li, T. W. Huang, D. J. Perreault, J. L. Dawson
Sponsorship: Lincoln Laboratory, Deshpande Center for Technological Innovation

In order to increase overall power efficiency of RF power amplifiers (PAs) over a wider output power range and significantly simplify RF/analog front-end for the PAs, we designed a new outphasing transmitter architecture based on asymmetric multilevel outphasing (AMO) modulation [1]. Figure 1 compares the AMO modulation with linear amplification with nonlinear component (LINC) [2] and multilevel (ML) LINC modulation [3]. Independently switching the supply voltage for each PA achieves the smallest outphasing angle with AMO. Discrete supply voltage levels and low oversampling allow wideband transmission such as for WLAN, WiMAX, and 4G LTE systems.

The power-efficiency improvement of AMO transmitter is compared with LINC and ML-LINC in Figure 2. Depending on the probability density distribution of transmission signal amplitude, each of the supply voltage levels can be optimized. For demonstration, an overall transmitter is simulated in a 65-nm CMOS process with HSUPA and WLAN signals. Compared to conventional outphasing modulation, the simulation results show an efficiency improvement from 17.7% to 40.7% for HSUPA at 25.3 dBm output power and from 11.3% to 35.5% for WLAN 802.11g at 22.8 dBm, while still meeting system linearity requirements.

A compact and low-power all-digital modulator, replacing the bulky RF/analog front-end, has been designed to drive a PA in AMO transmitters. For a small silicon footprint and wideband linear operation, digital predistortion technique is applied to compensate for the mismatch existing in an open-loop direct RF phase converter.

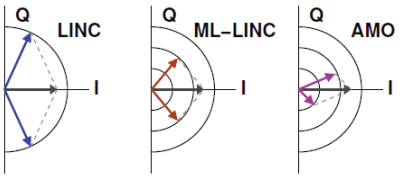

FIGURE 1: Signal component vector diagram for LINC, ML-LINC, and AMO. The smallest outphasing angle is achieved with AMO.

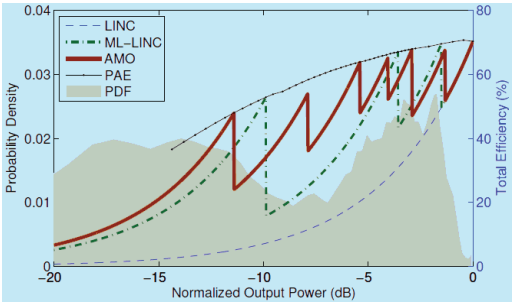

FIGURE 2: Probability density distribution of HSUPA signal amplitude and corresponding efficiency for standard LINC, ML-LINC, and AMO with a 65-nm CMOS switching PA.

REFERENCES

[1] S. Chung, P. A. Godoy, T. W. Barton, Z. Li, T. W. Huang, D. J. Perreault, J. L. Dawson, "Asymmetric multilevel outphasing architecture for multi-standard transmitters," to be presented *at IEEE Symposium on Radio Frequency Integrated Circuits*, Boston, MA, June, 2009.

[2] D. Cox, "Linear amplification with nonlinear components," *IEEE Transactions on Communications*, vol. 12, pp. 1942-1945, Dec. 1974.

[3] Y.-J. Chen, K.-Y. Jheng, A.-Y. Wu, H.-W. Tsao, and B. Tzeng, "Multilevel LINC system design for wireless transmitters," in *Proc. International Symposium on VLSI Design, Automation, and Test*, 2007, pp. 25-27.

[4] M. E. Heidari, M. Lee, and A. A. Abidi, "All digital outphasing modulator for a software-defined transmitter," *IEEE Journal of Solid State Circuits*, vol. 44, pp. 1260-1271, Apr. 2009.

Digitally Assisted Subsampler for RF Power-amplifier Linearization Systems

S. Chung, J. L. Dawson
Sponsorship: SRC/FCRP C2S2

Subsampling is recognized as an energy-efficient signal processing technique for highly digital transceivers [1]-[3]. However, subsamplers are notorious for low SNR performance due to noise folding and for stringent requirements for anti-aliasing prefilters. This combination of faults has largely undermined their use in high-performance receivers. In transmitters, however, the situation is *fundamentally* different. The signal environment has fewer extreme aggressors, such as blockers, and the transmitted data is often known in advance of actual transmission. This last fact enables the use of averaging and other signal processing techniques to overcome the noise-folding problem.

Figure 1 shows a digitally assisted subsampler, which is designed to serve as a downconversion path in adaptive predistortion transmitters with 800MHz-5.8GHz RF power amplifiers [4]. We use digital averaging to overcome the noise-folding problems of subsampling, obtaining a final SNDR of 73.1dB for signals centered around a 2.4GHz carrier. Using quadrature subsampling, we obtain both I and Q samples from the same physical path and thereby eliminate the IQ gain mismatch. When used as part of an adaptive predistortion system, the subsampler enables an EVM improvement of 3.2% and distortion products suppression of up to 7.6dB for 802.11g signals. The subsampler IC, designed in a 90-nm CMOS process, consumes 6.0mW from a 1.2V supply.

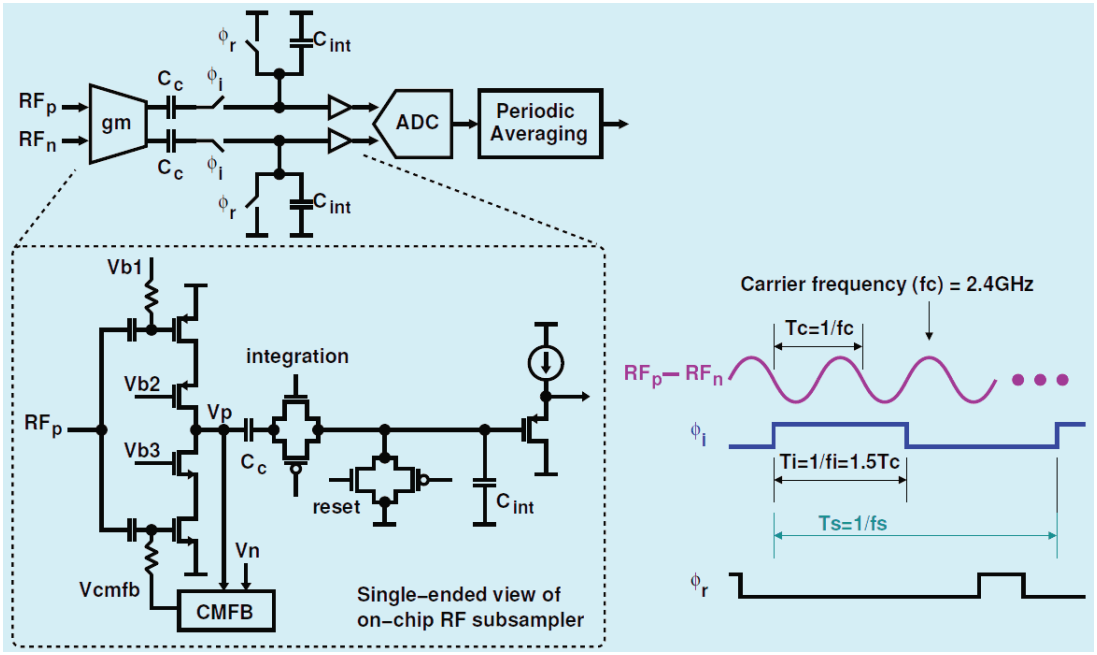

FIGURE 1: Digitally assisted subsampler architecture. Quadrature sampling is employed (i.e., Q samples are offset from I samples by Tc/4), allowing both the I and Q samples to be obtained from the same physical signal path.

REFERENCES

[1] H. Pekau *et al.,* " 2.4GHz CMOS sub-sampling mixer with integrated filtering," *IEEE Journal of Solid-State Circuits*, pp. 2159-2166, Nov. 2005.

[2] S. Karvonen *et al.,* "A quadrature charge-domain sampler with embedded FIR and IIR filtering functions," *IEEE Journal of Solid-State Circuits*, pp. 507-515, Feb. 2006.

[3] R Bagheri *et al.,* "An 800-MHz.6-GHz software-defined wireless receiver in 90-nm CMOS," *IEEE Journal of Solid-State Circuits*, pp. 2860-2876, Dec. 2006.

[4] S. Chung and J. L. Dawson, "73.1dB SNDR digitally assisted subsampler for RF power amplifier linearization systems," to be presented at *IEEE Symposium on VLSI Circuits*, Kyoto, Japan, June, 2009.

Outphasing Energy-recovery Amplifier with Resistance Compression for Improved Efficiency

P. A. Godoy, D. J. Perreault, J. L. Dawson

Sponsorship: Deshpande Center for Technological Innovation, CICS

The outphasing power amplifier dates back to the early 1930's as an approach for the simultaneous realization of high-efficiency and high-linearity amplification [1]. The principle of outphasing, also known as linear amplification of nonlinear components (LINC) [2], is shown in Figure 1(a). It is based on the idea that an arbitrary input signal can be divided into two constant-amplitude, phase-modulated signals that can each be non-linearly amplified and then recombined as a vector sum to produce an output signal that is a linearly amplified version of the input. The key advantage of this approach is that each amplifier operates in an efficient albeit nonlinear mode, and yet the final output can be highly linear, breaking the usual tradeoff between efficiency and linearity. The disadvantage lies in the output-combining network: when the two amplifiers are outphased to vary the amplitude, power is wasted as heat in the isolation resistor [3].

We describe a new outphasing energy recovery amplifier (OPERA) shown in Figure 1(b) that replaces the isolation resistor in the conventional matched combiner with a resistance-compressed rectifier for improved efficiency. The rectifier recovers the power normally wasted in the isolation resistor back to the power supply, while a resistance compression network (RCN) [4] reduces the impedance variation of the rectifier as the output power varies. Because the combiner requires a fixed resistance at the isolation port to ensure matching and isolation between the two outphased power amplifiers (PAs), the RCN serves to maintain high linearity as well as high efficiency in the switching-mode PAs. For demonstration, a prototype OPERA system is designed and implemented with discrete components at an operating frequency of 48MHz, delivering 20.8W peak power with 82.9% PAE. The measurement results show an efficiency improvement from 17.9% to 42.0% for a 50kHz 16-QAM signal with a peak-to-average power ratio of 6.5dB.

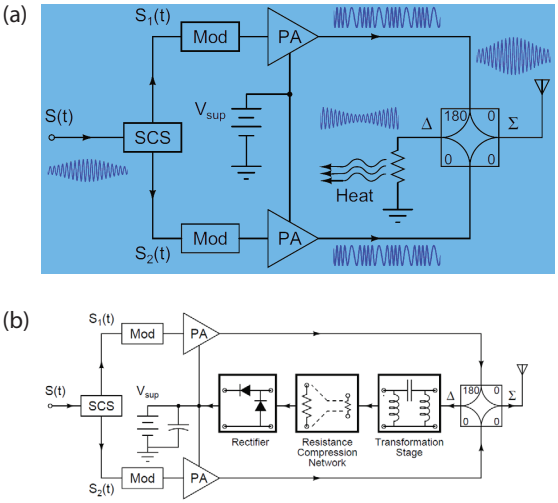

FIGURE 1: (a) Outphasing power amplifier illustrating power wasted in the combining network. (b) Proposed outphasing energy recovery amplifier, utilizing a resistance-compression network in the power-recycling network.

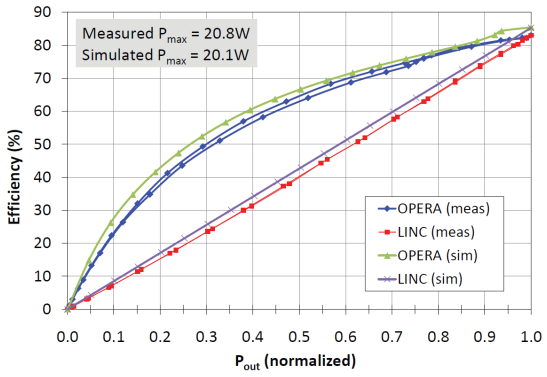

FIGURE 2: Simulated and measured system efficiency vs. output power for the OPERA prototype, with and without energy recovery.

REFERENCES

[1] H. Chireix, "High Power Outphasing Modulation," *Proceedings of the IRE,* vol. 23, pp. 1370-1392, 1935.

[2] D. Cox, "Linear Amplification with Nonlinear Components," *Communications, IEEE Transactions on,* vol. 22, pp. 1942-1945, 1974.

[3] I. Hakala, D. K. Choi, L. Gharavi, N. Kajakine, J. Koskela and R. Kaunisto, "A 2.14-GHz Chireix outphasing transmitter," *Microwave Theory and Techniques, IEEE Transactions on,* vol. 53, pp. 2129-2138, 2005.

[4] Y. Han, O. Leitermann, D. A. Jackson, J. M. Rivas and D. J. Perreault, "Resistance Compression Networks for Radio-Frequency Power Conversion," *Power Electronics, IEEE Transactions on,* vol. 22, pp. 41-53, 2007.

SAR ADC with Local Supply Capacitors and Adiabatic Charging for Use in Medical Implants

T. Khanna, J. L. Dawson

The proposed research program has two primary goals. The first goal is to improve the evaluation and treatment of patients with diabetes and a variety of movement disorders, including Parkinson's disease, restless legs syndrome, and essential tremor, by allowing doctors to continuously monitor relevant biomarkers over much longer time scales and with better precision than currently possible. The second goal is that the proposed implant be a *platform* for electronic sensory monitoring that is inexpensive and flexible and that can be used with a wide variety of sensors and for a wide variety of purposes, such as chemical sensors for monitoring blood chemistry. In this work, we develop an energy-efficient analog-to-digital converter designed to operate with a power management scheme using ultracapacitors as opposed to a battery.

Two techniques are employed to save on energy for the entire system. The first is the use of an integrated capacitor that acts as a local supply for the data conversion circuit. This technique allows for us to duty-cycle the bandgap reference circuit used for power management and can be seen in Figure 1. The second technique is to use adiabatic charging [1], [2] of the capacitors contained in the SAR ADC. This application is ideal for adiabatic techniques because of the low frequency of operation and the ease with which we can reclaim energy from discharging the capacitors. Building on the application in [3], our integrated capacitor acting as a local supply allows us to reclaim energy without having to design any energy-recovery circuitry.

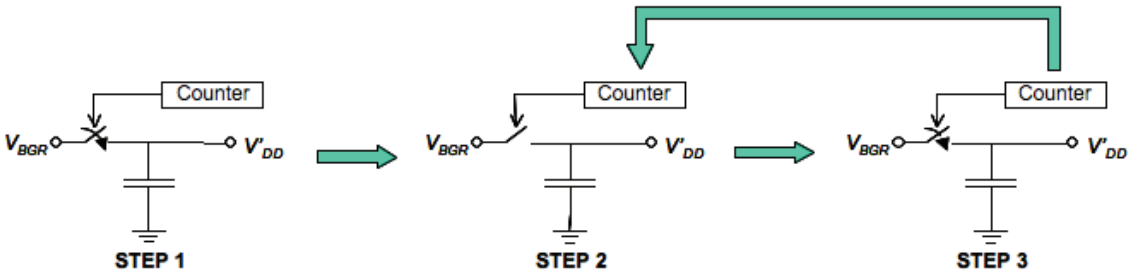

FIGURE 1: Charging algorithm for local supply capacitor.
Step 1: Charge integrated capacitor to bandgap reference (BGR) voltage, V_{BGR}.
Step 2: Run circuits for N clock cycles until temporary supply voltage, V'_{DD}, is reduced by m%.
Step 3: On-chip counter triggers switch after N clock cycles and charges integrated capacitor up to BGR voltage. Return to Step 2.

REFERENCES

[1] L. J. Svensson and J. G. Koller, "Driving a capacitive load without dissipating fCV^2," *Low Power Electronics, 1994. Digest of Technical Papers. , IEEE Symposium,* pp. 100-101, 1994.

[2] J. G. Koller and L. J. Svensson, "Adiabatic charging without inductors," *USC/ISI technical report ACMOSTR-3a*, Feb 8, 1994.

[3] M. van Elzakker, E. van Tuijl, P. Geraedts, D. Schinkel, E. Klumperink and B. Nauta, "A 1.9µW 4.4fJ/Conversion-step 10b 1MS/s Charge-Redistribution ADC," *Solid-State Circuits Conference, 2008. ISSCC 2008. Digest of Technical Papers. IEEE International,* pp. 244-610, 2008.

[4] T. Denison, W. Santa, G. Molnar and K. Miesel, "Micropower Sensors for Neuroprosthetics," *Sensors, 2007 IEEE,* pp. 1105-1108, 2007.

Memory Architecture for µImplant

H. S. Khurana, J. L. Dawson

Sponsorship: Irwin Mark Jacobs and Joan Klein Jacobs Presidential Fellowship

The evaluation and effective treatment of patients suffering from movement disorders such as Parkinson's disease, restless legs syndrome, and others require continuous monitoring and reliable data collection. This monitoring is a challenge due to routine movements of patients and shortcoming of the methods that rely on patients measuring themselves. To overcome these challenges the group uImplant focuses on developing an inconspicuous and minimally invasive IC system encapsulated in a bio-compatible packaging. The complete system will consist of three main areas: power management, signal collection, and data storage and communication. In this direction we have designed a customized low-power sub-threshold SRAM with on-chip features that reduce energy consumption.

The SRAM (Figure 1) is designed with the low-power application in focus. It is 32-kb, 6T bit-cell sequential read/write memory with on-chip power-saving features. Since the data gets written sequentially, the rows are powered just before they get written with valid data. After each row-write, the row written stays on for data retention. The unused rows at any given time stay unpowered, thereby saving precious energy in micro-implant applications. Further savings come by lowering the leakage currents associated with standby cells.

In low-frequency sub-threshold operations, the leakage currents dominate the power consumption and are the main energy sinks. Since leakage currents are a strong function of the supply voltage, it is important to reduce the voltage to the lowest possible levels without losing the bit cell data. In this memory, the system supply voltage is stepped down using a capacitor stack (Figure 2) that switches between a series and a parallel configuration to divide the voltage to desired voltage levels. Besides low standby leakage currents, low dynamic currents give additional savings during the write operation.

The architecture enables independent floating of the supply voltage to each word in the memory. This floating trait reduces power consumption during the write operation. It also reduces the minimum size requirement on the pass transistors connecting the data input bit-lines to the bit cell during the write operation.

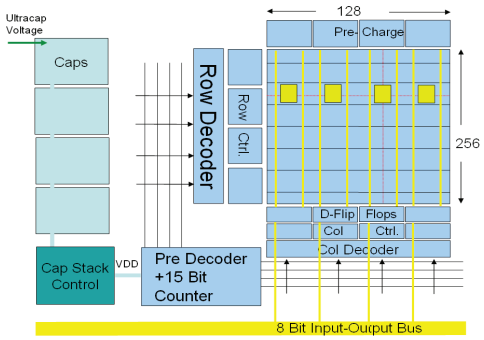

FIGURE 1: Architecture of the 32-kb (128x256) SRAM with energy-saving features.

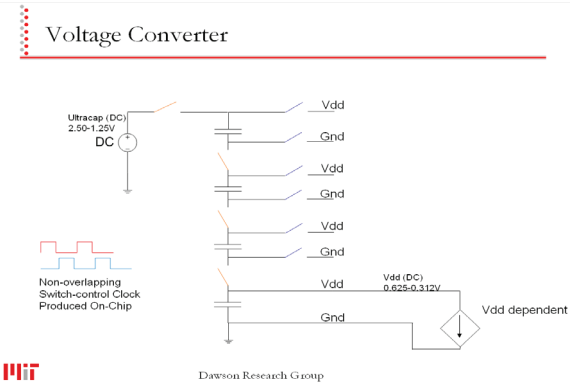

FIGURE 2: Voltage divider using capacitor stack with switches.

REFERENCES

[1] B. H. Calhoun, and A. P. Chandrakasan, "A 256-kb 65-nm Sub-threshold SRAM Design for Ultra-Low-Voltage Operation," *IEEE Journal of Solid-State Circuits*, vol. 42, no. 3, pp. 680-688, Mar. 2007.

[2] Jan M. Rabaey, Anantha Chadrakasan, and Borivoje Nikolic, *Digital Integrated Circuits*. New Jersey: Prentice Hall, 2003.

[3] Jiann S. Yuan, and Li Yang, "Teaching Digital Noise and Noise Margin Issues in Engineering Education," *IEEE Transactions on Education*, vol. 48, pp. 162-167, Feb 2005.

[4] R.J. Baker, *CMOS – Circuit Design, Layout, and Simulation*. John Wiley and Sons: IEEE Press, 2005.

Transmitters for High Efficiency, 10 Gb/s Wireless Communications in the 60 GHz Band

O. T. Ogunnika, J. L. Dawson
Sponsorship: SiGe Semiconductor, Lincoln Laboratory

The purpose of this project is to design an RF transmitter architecture that achieves 10 Gb/s data transfer over a 60 GHz wireless link with high power efficiency. With the availability of 7 GHz of unlicensed bandwidth centered at 60 GHz, this space has emerged as an active area of research. A number of challenges will be faced in the process of bringing this project to completion. Strong atmospheric absorption at 60 GHz lowers the signal-to-noise ratio (SNR) available at the receiver [2]. The low SNR limits the complexity of the constellations that can be used and thus reduces the number of bits per symbol that can be encoded with the modulation strategy. Extremely fast baseband modulators will therefore be required for high data rate transmission because more symbols per second will have to be transmitted. The fundamental challenge of simultaneously obtaining good linearity and high efficiency in power amplifiers is further exacerbated at this carrier frequency, complicating transceiver design. Delivering significant power at 60 GHz requires very fast devices with high f_{max} and f_T. This technological hurdle has been lowered with recent advances in SiGe, III-V semiconductor technology and deeply scaled CMOS.

The technical approach of this project is to exploit complete co-design of the modulation strategy with a new power amplifier concept: Asymmetric Multilevel Outphasing (AMO) [1]. This architecture combines the best properties of polar transmitters and outphasing (LINC) transmitters [5]. The power amplifier's efficiency is improved without significantly degrading its linearity by using the combination of drain voltage modulation and rapid outphasing. A key aspect of this project will be the investigation of energy recovery as a means of further improving the transmitter's efficiency. The use of resistance compression networks [3] as a means of recovering the energy normally lost during outphasing will be critical. To achieve these goals, the most significant research challenges are:

1. achieving baseband modulation commensurate with 10 Gb/s transmission with the new AMO architecture and
2. designing a symbol constellation and modulation strategy that maximally exploits the architecture.

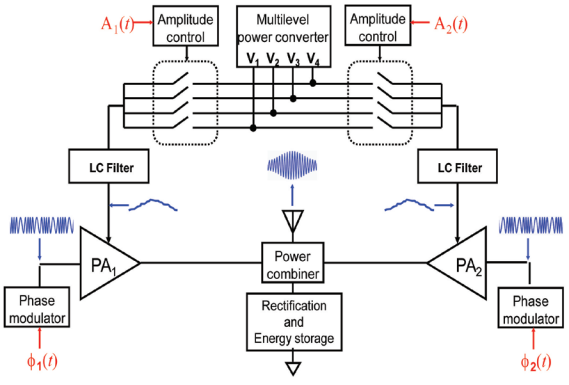

FIGURE 1: System diagram of the Asymmetric Multilevel Outphasing PA. The energy recovery network and amplitude path will be designed to function at a 60 GHz carrier frequency and 10 Gb/s data rate.

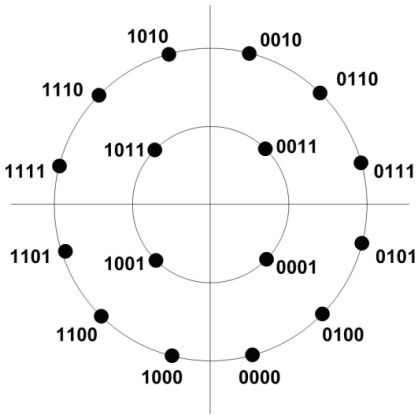

FIGURE 2: Constellation diagram for 16-APSK [4]. This constellation is a potentially suitable modulation scheme because it requires just two amplitude levels.

REFERENCES

[1] S-W Chung, P. A. Godoy, T. W. Barton, E.W. Huang, D. J. Perreault and J. L. Dawson, "Asymmetric Multilevel Outphasing Architecture for Multi-standard Transmitters," IEEE Radio-Frequency Integrated Circuits (RFIC) Symposium, pp. 237-240, June 2009.

[2] C. H. Doan, S. Emami, D. Sobel, A. M. Niknejad, and R.W. Brodersen, "60 GHz CMOS radio for Gb/s wireless LAN," IEEE Radio Frequency Integrated Circuits (RFIC) Symposium, pp. 225-228, June 2004.

[3] Y. Han, O. Leitermann, D. A. Jackson, J. M. Rivas and D. J. Perreault, "Resistance Compression Networks for Radio-Frequency Power Conversion," IEEE Transactions on Power Electronics, Vol. 22, No. 1, pp. 41-53, Jan. 2007.

[4] R. De Gaudenzi, A. G. Fabregas and A. Martinez, "Turbo-coded APSK modulations design for satellite broadband communications," International Journal of Satellite Communications and Networking, Vol. 24, pp. 261-281, 2006.

[5] D. C. Cox, "Linear Amplification with Nonlinear Components," IEEE Transactions on Communications, Vol. 22, No. 12, pp. 1942-1945, Dec. 2004.

A Distributed Power-management Integrated Circuit Based on Ultracapacitors

W. Sanchez, J. L. Dawson
Sponsorship: Gates Millennium Scholars Graduate Fellowship

Presently, continuous monitoring of patients is very difficult and typically intrudes greatly on their lifestyles and daily routines. Methods that rely on patients measuring themselves are unreliable. An unobtrusive, minimally invasive monitoring platform would give medical scientists unprecedented access to continuous long-term data on patients. In this work, we develop a power management integrated circuit (IC) to power minimally invasive monitors using ultracapacitors. The focus of the project is to improve the evaluation and treatment of patients with a variety of movement disorders including Parkinson's disease, restless legs syndrome, and essential tremor by allowing doctors to continuously monitor relevant biomarkers over much longer time scales and with better precision than currently possible.

Because of the strict size requirements of a minimally invasive IC, a simple battery cannot be used due to the area overhead and replacement requirements. As a solution, a wirelessly rechargeable ultracapacitor network can be used as the power supply for sensing and data conversion, storage, and transmission circuitry.

Ultracapacitors can carry 5-10% the energy densities of conventional battery chemistries of comparable weight [3]. With over 10^5 recharge cycles possible, they are attractive candidates for implantable applications employing efficient wireless recharging schemes. Figure 1 is a block diagram for the power management IC. Discrete ultracapacitors from Maxwell (5F, 14mmx24mmx0.5mm) are used in the initial prototype. Figure 2 demonstrates the principle of operation that achieves use of 96.8% of the initial charge. Currently, a prototype whereby the ultracapacitors are recharged using RF rectification is under development.

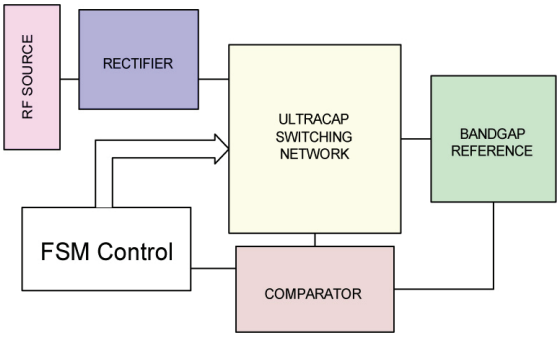

FIGURE 1: Block diagram of power management integrated circuit.

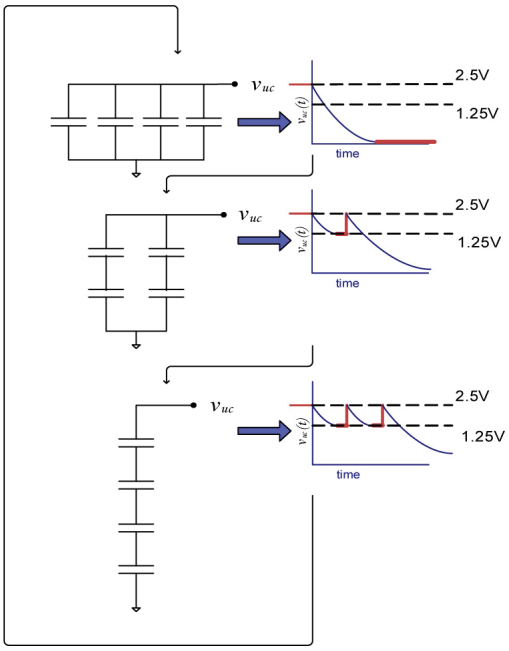

FIGURE 2: Operating principle of ultracapacitor stacking network. Three-step stacking maintains a voltage adequate for circuit operation, utilizing 96.8% of the initial charge stored on the ultracapacitors.

REFERENCES

[1] Nesscap Ultracapacitor - Technical Guide, Nesscap Co., Ltd., 2008

An Ultra-high-speed Zero-crossing-based ADC

A. Chow, H.-S. Lee
Sponsorship: SRC/FCRP C2S2

With an increasing need for higher data rates, both wireless applications and data links are demanding higher speed analog-to-digital converters (ADC) with medium resolution. This work will investigate ADCs with a sampling rate up to 2 Gs/s with 6-8 bits of resolution. Time-interleaved converters achieve their high sampling rate by placing several converters in parallel. Each individual converter, or channel, has a delayed sampling clock and operates at a reduced sampling rate. Therefore, each channel is responsible for digitizing a different time slice. This method requires that the individual converters, which make up the parallel combination, be matched. Mismatches and non-idealities such as gain error, timing error, and voltage offset degrade the performance. Hence channel matching is an important design consideration for time-interleaved ADCs.

Although digital calibration can mitigate many of these non-idealities, a timing mismatch is a non-linear error that is more difficult to remove. At sampling rates up to 2Gs/s, digital calibration would consume a large amount of power. An alternative solution uses a global switch running at the full speed of the converter. Although this technique works reasonably well for medium-high speed ADC's [1], [2], its effectiveness is limited by parasitic capacitance. We have developed a double-global sampling technique to remove the effect of parasitic capacitance on the timing skew. At higher speeds the ability to turn the switch on and off at the full sampling rate becomes a major challenge. The use of scaled CMOS technology and gate-boosted switches still enables multi-GHz input bandwidth.

In this work, we employ a fast, single-slope architecture (Figure 1). Since the single-slope architecture is more sensitive to non-idealities such as ramp nonlinearity, the design uses several innovative ideas to improve the linearity. Offset cancellation is also incorporated in the design. The first silicon fabricated in a 90-nm CMOS technology has been received and is being characterized.

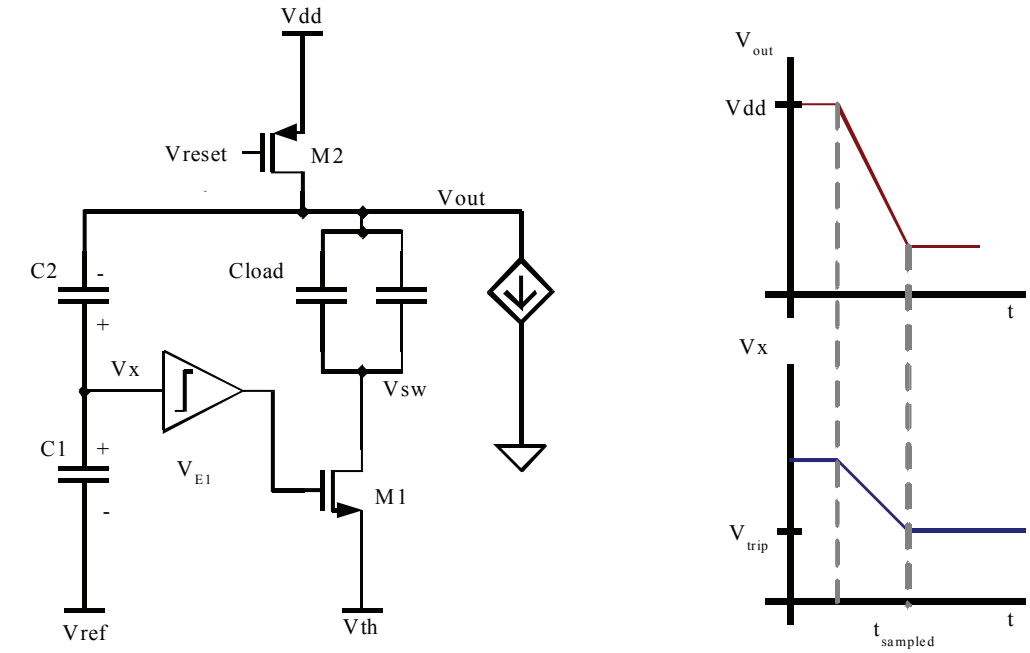

FIGURE 1: One stage of a single-slope CBSC-based pipelined ADC.

REFERENCES

[1] M. Gustavsson and N. N. Tan, "A global passive sampling technique for high-speed switched-capacitor time-interleaved ADCs," *Circuits and Systems II: Analog and Digital Signal Processing, IEEE Transactions on*, vol. 47, pp. 821-831, 2000.

[2] S. Gupta, M. Choi, M. Inerfield and Jingbo Wang, "A 1GS/s 11b Time-Interleaved ADC in 0.13/ spl mu/m CMOS," *Solid-State Circuits Conference, 2006. ISSCC 2006. Digest of Technical Papers. IEEE International*, pp. 2360-2369, 2006.

A High-performance Zero-Crossing-based Pipeline ADC

J. Chu, H.-S. Lee
Sponsorship: SRC/FRCP C2S2

In this work, we are designing a high-performance pipeline ADC using a Zero-crossing-based (ZCB) structure [1]. The ZCB circuits offer many advantages compared with the earlier comparator-based switch-capacitor circuits (CBSC) [2]. The focus of the project is to explore novel circuit structures for ZCB circuits to improve the speed and accuracy as well as the figure-of-merit (FOM). In particular, we employ a fully-differential structure to improve ZCBC's robustness against common mode noise. An additional benefit of differential design is the increase in the available signal range, which helps to improve SNR. We implemented a multi-bit MDAC to improve its power-efficiency and to help relax component accuracy requirements. In this design, we used dynamically biased current sources to achieve high linearity at high operating frequencies. Dynamic biasing can be used to compensate for nonlinearity of the ramp to improve the linearity of the system.

A prototype chip has been fabricated, and the preliminary measurement data indicates a better-than-10 effective number of bits (ENOBs) at 100MHz sampling rate. More tests are currently being performed.

In the next iteration of this project, time-interleaving will be used to achieve ultra-high sampling rates with very low power. In a time-interleaved structure, matching between the different channels will be very important to maintain the desired performance. Any mismatch in gain, offset, and timing can greatly degrade the performance. We plan to use sample-time adjustment to mitigate the timing errors. Careful design and layout will be needed to reduce the other mismatches.

REFERENCES

[1] L. Brooks and Hae-Seung Lee, "A Zero-Crossing-Based 8b 200MS/s Pipelined ADC," *Solid-State Circuits Conference, 2007. ISSCC 2007. Digest of Technical Papers. IEEE International,* pp. 460-615, 2007.

[2] T. Sepke, J. K. Fiorenza, C. G. Sodini, P. Holloway and Hae-Seung Lee, "Comparator-based switched-capacitor circuits for scaled CMOS technologies," *Solid-State Circuits Conference, 2006. ISSCC 2006. Digest of Technical Papers. IEEE International,* pp. 812-821, 2006.

[3] M. Gustavsson and N. N. Tan, "A global passive sampling technique for high-speed switched-capacitor time-interleaved ADCs," *Circuits and Systems II: Analog and Digital Signal Processing, IEEE Transactions on,* vol. 47, pp. 821-831, 2000.

[4] S. Gupta, M. Choi, M. Inerfield and Jingbo Wang, "A 1GS/s 11b Time-Interleaved ADC in 0.13/ spl mu/m CMOS," *Solid-State Circuits Conference, 2006. ISSCC 2006. Digest of Technical Papers. IEEE International,* pp. 2360-2369, 2006.

A Low-voltage Zero-crossing-based Delta-Sigma ADC

M. Guyton, H.-S. Lee
Sponsorship: CICS

Many analog signal-processing circuits use operational amplifiers (op-amps) in a negative feedback topology. The error in these feedback systems is inversely proportional to the gain of the op-amp. Because scaled CMOS technologies use shorter channel lengths and require lower power supply voltages, it becomes more difficult to implement high gain op-amps. Recently, a comparator-based switched-capacitor (CBSC) technique was proposed [1] that uses a comparator rather than an op-amp to implement switched-capacitor topologies. This technique was generalized to the use of zero-crossing detectors [2].

In this project, we investigate very-low-voltage delta-sigma converters. One of the biggest challenges of low-voltage circuits is the transmission gates that must pass the signal. If the signal is near the middle of the power supply range, neither the NMOS nor the PMOS transistor has sufficient gate drive to pass the signal properly. The switched-op-amp technique [3] was proposed to mitigate this problem. In this technique, the output of the op-amp is directly connected to the next sampling capacitor without a transmission gate to perform charge transfer.

During the charge-transfer phase, the op-amp is switched off, and the output is grounded.

Much like the standard switched-capacitor technique, zero-crossing-based (ZCB) circuits use two-phase clocking, having both sampling and charge-transfer clock phases. Unlike in a standard switched-capacitor circuit, in a ZCB circuit all current sources connected to the output node are off at the end of the charge-transfer phase. Therefore, there is no op-amp or current source to turn off to accommodate the charge transfer without a transmission gate. Thus, the ZCB technique is inherently better suited to low-voltage applications than are switched-op-amp circuit topologies. Figure 1 shows a fully-differential low-voltage ZCB integrator stage using the combined techniques. We have designed a fourth-order delta-sigma ADC for operation at 1-V power supply using this integrator stage for audio-band applications. A new output pre-sampling technique has been developed to dramatically reduce the linearity requirement of the ramp waveform. The chip fabricated in 130-nm CMOS is currently being tested.

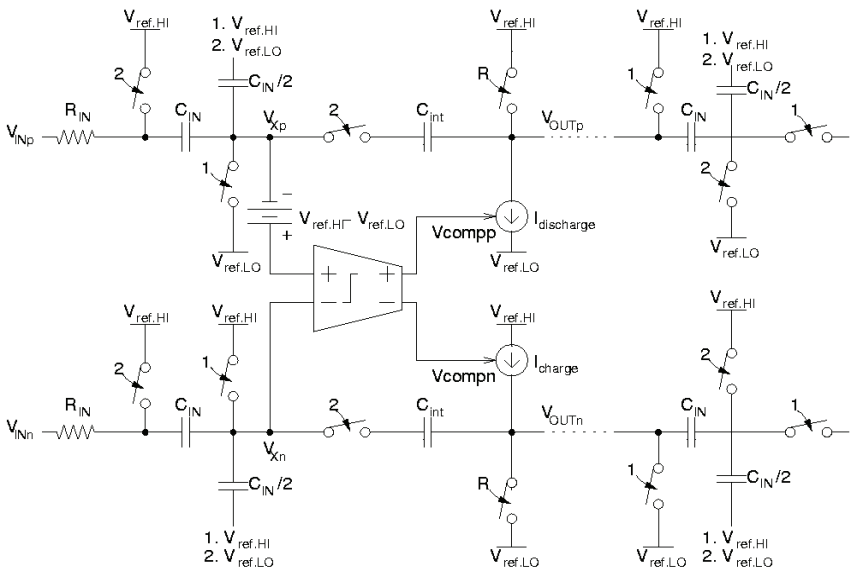

FIGURE 1: Fully-differential zero-crossing-based switched-capacitor integrator. The input of the next integrator stage is also shown.

REFERENCES

[1] T. Sepke, J.K. Fiorenza, C.G. Sodini, P. Holloway, and H.S. Lee, "Comparator-based switched-capacitor circuits for scaled CMOS technologies," *2006 IEEE ISSCC Digest of Technical Papers*, Feb. 2006, p. 220.

[2] L. Brooks and H.S. Lee, "A zero-crossing-based 8b, 200 MS/s pipelined ADC," *2007 IEEE ISSCC Digest of Technical Papers*, Feb. 2007, pp.460-461.

[3] J. Crols and M. Steyaert, "Switched-op-amp: an approach to realize full CMOS switched-capacitor circuits at very low power supply voltages," *IEEE Journal of Solid-State Circuits*, vol. 29, no. 8, pp. 936-942, Aug. 1994.

Design of a Reconfigurable Mixed-Signal System

P. Lajevardi, A. P. Chandrakasan, H.-S. Lee
Sponsorship: CICS

Switched-capacitor circuits can be used to implement many analog systems such as ADCs, DACs, filters, amplifiers, and integrators. In earlier phase of this research, a reconfigurable switched-capacitor system is proposed to implement different analog systems. A prototype system has been fabricated that shows basic reconfigurability to implement a pipe-lined ADC and a switched-capacitor filter. A second prototype system is being designed that utilizes highly reconfigurable blocks. Figure 1 shows the block diagram of such systems. The building blocks have the same functionality and can be programmed to implement a multiplier or an integrator with a reconfigurable coefficient. Such a system has many applications such as in software-defined radios and rapid prototyping of analog circuits.

The design of such systems has had limited success since many different op-amp topologies are required to cover a large performance and configuration space. Recently, [1] and [2] proposed zero-crossing based (ZCB) circuits to design ADCs. ZCB circuits replace the op-amp in traditional switched-capacitor circuits with a combination of a current source and a zero-crossing detector. ZCB circuits are well suited for highly reconfigurable system since their power consumption scales with the operating frequency and required SNR. In addition, ZCB circuits benefit from technology scaling due to their more digital circuit-like operation in contrast to conventional op-amp based circuits.

The system is designed to operate at different speeds while the power consumption is kept at the optimum level. A key challenge in design of this system is to keep the cost of programmability low in terms of additional power consumption and performance degradation. Several innovative circuit techniques have been used to address this issue.

REFERENCES

[1] T. Sepke, J. Fiorenza, C.G. Sodini, P. Holloway, and H.-S Lee, "Comparator-based switched-capacitor circuits for scaled CMOS technologies," in *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, San Francisco, CA, Feb. 2006, pp.220-221.

[2] L. Brooks and H.-S. Lee, "A zero-crossing-based 8b 200MS/S pipeline ADC," in *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, San Francisco, CA, Feb. 2007, pp. 460-461.

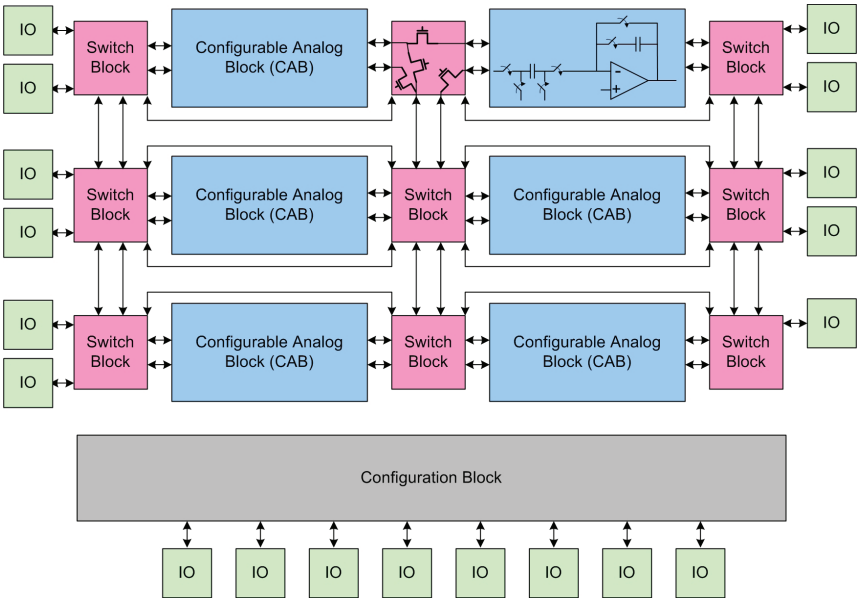

FIGURE 1: Block diagram of reconfigurable zero-crossing-based analog circuits. Each configurable analog block can be programmed to perform an integration or multiplication. The connection of blocks is also programmable.

Front-end Design for Portable Ultrasound Systems

S. Lee, A. P. Chandrakasan, H.-S. Lee
Sponsorship: Samsung Fellowship, CICS

Most current ultrasound imaging systems use piezoelectric materials for the ultrasound transducer. The recent development of micro-electromechanical systems (MEMS) allowed fabrication of capacitive micromachined ultrasound transducers (CMUTs). A CMUT is a micromachined capacitor whose value changes according to the DC bias voltage or external pressure due to the physical deformation of the top plate by electrostatic force or external pressure. The major advantages of this transducer technology are the potential for integration with supporting electronic circuits, ease of fabrication, higher resolution due to small transducer size, and improved bandwidth and sensitivity [1].

This project focuses on the front-end design of portable ultrasound systems using CMUTs. Figure 1 presents a conceptual block diagram of the system. Implementing an ADC at each channel input makes possible digital beam-forming in the receive (Rx) path, which enhances ultrasound image quality. To implement as many ADCs as the number of transducer channels, each ADC must consume as little power as possible, and each should be implemented in a small area. With the performance requirements at 10~20MHz sampling frequency with 8~10bits of resolution, successive approximation or zero-crossing-based (ZCB) ADC are good candidates [2]. We are investigating these topologies to determine the optimum topology for the application. We are also exploring the control of the fire timing and pulse shape of the transmit (Tx) elements, such that beam-forming in the Tx path is possible to give a larger Rx signal.

Recently, a few 2D imaging systems using CMUT as ultrasound transducers have been reported, but they do not use real-time imaging [1]. The digital image processing block will be considered in the system level for real-time imaging. After completing the 2D ultrasound image system using a 1D transducer, we will examine the feasibility of the 3D ultrasound image system using 2D transducers.

REFERENCES

[1] O. Oralkan, "Acoustical Imaging Using Capacitive Micromachined Ultrasonic Transducer Arrays: Devices, Circuits, and systems", Ph.D. thesis, Stanford University, 2004.

[2] L. Brooks and H.-S. Lee, "A Zero-Crossing-Based 8b 200MS/s Pipelined ADC," *Solid-State Circuits Conference, 2007. ISSCC 2007. Digest of Technical Papers. IEEE International*, pp. 460-615, 2007.

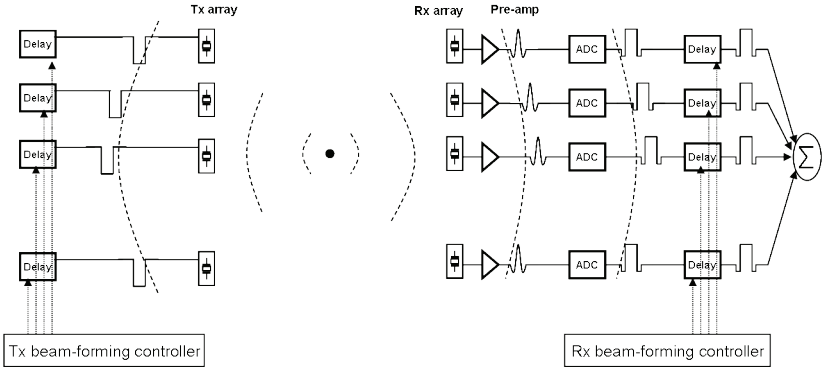

FIGURE 1: system block diagram

A High-accuracy, Zero-crossing-based Pipeline ADC

M. Markova, P. Holloway, H.-S. Lee
Sponsorship: CICS

Technology scaling poses challenges in designing analog circuits because of the decrease in intrinsic gain and reduced swing. An alternative to using high-gain amplifiers in the implementation of switched capacitor circuits has been proposed [1] that replaces the amplifier with a current source and a comparator. The new comparator-based switch capacitor (CBSC) and zero-crossing-based circuit (ZCBC) techniques have been implemented in two pipelined ADC architectures at 10MHz and 200MHz and 10-bit and 8-bit accuracy, respectively [1], [2].

The purpose of this project is to explore the use of the ZCBC technique for very-high-precision AD converters. The goal of the project is a 100MHz, 14-bit pipelined ADC. First, we are investigating dual-phase hybrid ZCBC operation to improve the power-linearity trade off of the A/D conversion [3] and to improve the power supply rejection. The first phase approximates the final output value, while the second phase allows the output to settle to its accurate value. Since the output is allowed to settle in the second phase, the currents through capacitors decay, permitting higher accuracy and power-supply rejection compared with standard ZCBCs. We are also developing linearization techniques for the ramp waveforms. Linear ramp waveforms require less correction in the second phase for given linearity, thus allowing faster operation. Innovative techniques for improving linearity beyond using a cascoded current source are explored, including output pre-sampling. In addition, overshoot reduction techniques will be used to improve the linearity requirements of the final phase. Alternatively, improvements in the zero-crossing detector design are pursued to lessen the linearity requirement of the ramp. Chopper stabilization will be used to reduce the effects of offsets and flicker noise.

REFERENCES

[1] T. Sepke, J.K. Fiorenza, C.G. Sodini, P. Holloway, and H.-S. Lee, "Comparator-based switched capacitor circuits for scaled CMOS technologies," *IEEE Journal of Solid State Circuits, IEEE International Solid State Circuits Conference Digest of Technical Papers*, February 2006, pp 220-221.

[2] L. Brooks, H.-S. Lee, "A zero-crossing based 8b 200MS/s pipelined ADC," *IEEE International Solid State Circuits Conference Digest of Technical Papers*, February 2007, pp 460-461.

[3] J.K. Fiorenza, "A comparator-based switched –capacitor pipelined analog-to-digital converter ," Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, 2007.

Machine Vision for an Intelligent Transportation System

Y. Fang, B. K. P. Horn, I. Masaki
Sponsorship: Intelligent Transportation Research Center, MTL

Environment-understanding technology is very vital for intelligent vehicles that are expected to respond automatically to fast- changing environments and dangerous situations. To obtain perceptual abilities, we should automatically detect static and dynamic obstacles and obtain related information such as their locations, speed, possible collision or occlusion, and other dynamic current or historic information. Conventional methods independently detect individual information, which is normally noisy and not very reliable. Instead we propose fusion-based and layered-based information-retrieval methodology as shown in Figure 1 to systematically detect obstacles and obtain their location/timing information for visible and infrared sequences. The proposed obstacle detection methodologies take advantage of the connections between different pieces of information and increase the computational accuracy of obstacle information estimation, thus improving environment understanding abilities and driving safety. Three examples are shown in Figures 2, 3 and 4.

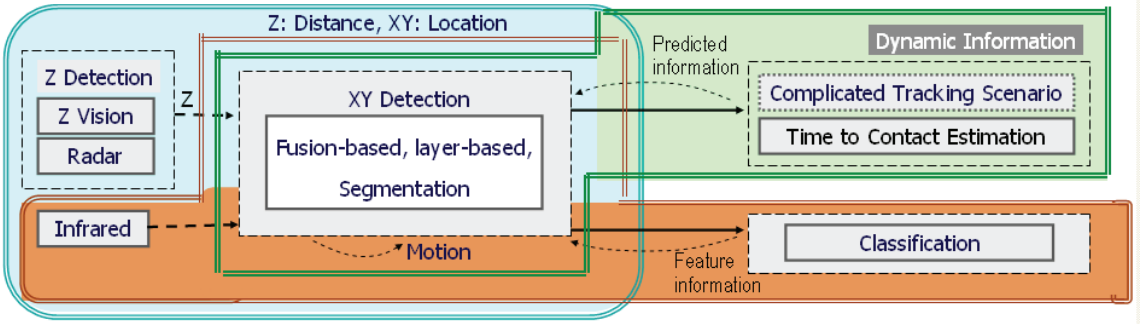

FIGURE 1: Segmentation result for urban day-time driving environment.

FIGURE 2: Segmentation result for urban day-time driving environment.

FIGURE 3: Segmentation result for night driving environment.

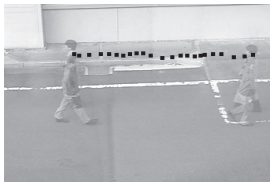

FIGURE 4: The segmentation/ tracking results before and after two people intersect.

REFERENCES

[1] Y. Fang, S. Yokomitsu, B. Horn, I. Masaki, "A Layered-based Fusion-based Approach to Detect and Track the Movements of Pedestrians through Partially Occluded Situations," to be presented at the *IEEE Intelligent Vehicles Symposium 2009 (IV2009)*.

[2] B. Horn, Y.Fang, I.Masaki, "Hierarchical framework for direct gradient-based time-to-contact estimation," to be presented at the *IEEE Intelligent Vehicles Symposium 2009 (IV2009)*.

[3] Y. Fang, B. K. P. Horn and I. Masaki, "Systematic information fusion methodology for static and dynamic obstacle detection in ITS," presented at *15th World Congress on Intelligent Transport Systems*.

[4] B. K. P. Horn, Yajun Fang and I. Masaki, "Time to Contact Relative to a Planar Surface," *Intelligent Vehicles Symposium, 2007 IEEE*, pp. 68-74, 2007.

[5] S. Negahdaripour and B. K. P. Horn, "Direct Passive Navigation," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. PAMI-9, pp. 168-176, 1987.

[6] B. K. P. Horn and E. J. Weldon, "Direct Methods for Recovering Motion," *International Journal of Computer Vision*, vol. 2, pp. 51-76, 06/01. 1988.

[7] Y. Fang, K. Yamada, Y. Ninomiya, B. Horn and I. Masaki, "Comparison between infrared-image-based and visible-image-based approaches for pedestrian detection," *Intelligent Vehicles Symposium, 2003. Proceedings. IEEE*, pp. 505-510, 2003.

[8] Y. Fang, K. Yamada, Y. Ninomiya, B. K. P. Horn and I. Masaki, "A shape-independent method for pedestrian detection with far-infrared images," *Vehicular Technology, IEEE Transactions on*, vol. 53, pp. 1679-1697, 2004.

Power-efficient Impedance-modulation Wireless Data Links for Biomedical Implants

S. Mandal, D. P. Kumar, R. Sarpeshkar

Low-power wireless links are important for the development of long-term implantable neural prostheses. Furthermore, in implanted systems with many neural recording electrodes, the data rate of the wireless link will need to be quite high since each electrode typically requires at least 5kHz of bandwidth. For low-power operation, inductively-coupled near-field wireless links have shown great promise and were used to develop a power-efficient data link for biomedical implants.

A bidirectional half-duplex wireless link that uses near-field inductive coupling between the implanted system and an external transceiver was designed in a 0.5-μm CMOS process. Our system minimizes power consumption in the implanted system by using impedance modulation to transmit high-bandwidth information in the uplink direction, i.e., from the implanted to the external system. We measured a data rate of 2.8Mbps at a bit error rate (BER) of $<10^{-5}$ (we could not measure error rates below 10^{-5}) and a data rate of 4.0Mbps at a BER of 10^{-3}. Experimental results also demonstrate data transfer rates up to 300kbps in the opposite, i.e., downlink direction. Theoretical analysis of the bit error rate performance was also carried out. This analysis allowed us to theoretically predict and experimentally verify an important effect regarding the asymmetry of rising and falling edges that is inherent in impedance modulation and that contributes to bit errors. The link dissipates 2.5mW in the external system and only 100μW in the implanted system, making it among the most power-efficient inductive data links reported. Our link is compatible with FCC regulations on radiated emissions.

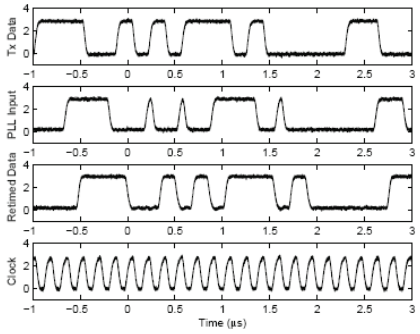

FIGURE 1: Uplink data transmission at 5.8Mbps with the coils 2cm apart.

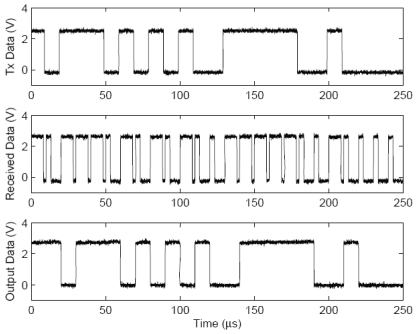

FIGURE 2: Downlink data transmission at 200kbps with the coils 2cm apart.

REFERENCES

[1] S. Mandal and R. Sarpeshkar, "Power-Efficient Impedance-Modulation Wireless Data Links for Biomedical Implants," *Biomedical Circuits and Systems, IEEE Transactions on,* vol. 2, pp. 301-315, 2008.

Low-power Circuits for Brain-machine Interfaces

R. Sarpeshkar, W. Wattanapanitch, S. Arfin, B. Rapoport, S. Mandal, M. Baker, M. Fee, S. Musallam, R. Andersen

Large-scale chronic multi-electrode neural recording and stimulation systems have emerged as an important experimental paradigm for investigating brain function. Experiments using such brain-machine interfaces (BMIs) have shown that it is possible to predict intended limb movements by analyzing simultaneous recordings from many neurons. These findings have suggested a potential approach for treating paralysis and other disorders and disabilities in humans. Chronic use of BMIs with large numbers of electrodes requires ultra-low-power operation so that the systems are miniature and implantable, heat dissipated in the brain is minimized, and repeated surgeries for battery replacement are unnecessary.

Currently, we are developing an ultra-low-power BMI system with applications for paralysis prosthetics, stroke, Parkinson's disease, epilepsy, prosthetics for the blind, and experimental neuroscience systems [1]. Our proposed system consists of a wireless hybrid analog-digital recording system that is capable of recording from neurons in the brain and a wireless neural stimulation system. Figure 1 shows an overall architecture of the recording system. It consists of low-power DAC-programmable analog circuits that are configured by external DSP. Depending on the user's choice, the system can be configured to report raw neural data from a selected set of electrodes, local field potential (LFP) data, or decoded motor parameters via a data telemetry uplink. Figure 2 shows the wireless neural stimulation system. The system consists of an external transmitter (not shown) controllable through a computer interface and a miniature, implantable wireless receiver-and-stimulator. The entire stimulation system weighs 0.6 g, occupies a footprint smaller than 1.5 cm2, and is capable of delivering biphasic current pulses to 4 addressable electrode sites at 32 selectable current levels ranging from 10 uA to 1 mA.

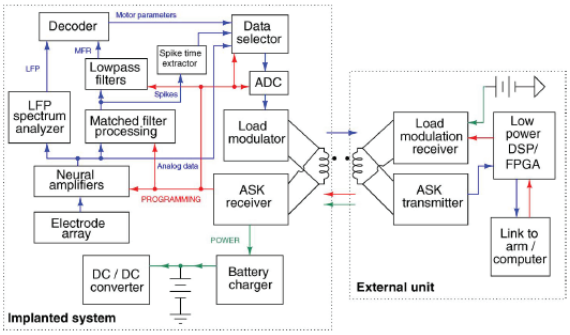

FIGURE 1: Block diagram of the hybrid analog-digital brain-machine interface system.

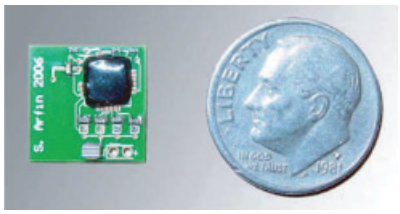

FIGURE 2: Protograph of the chip-on-board wireless neural stimulation system.

REFERENCES

[1] R. Sarpeshkar *et al*, "Low-Power Circuits for Brain-Machine Interfaces," *IEEE Transactions on Biomedical Circuits and Systems*, pp 173-183, Sept. 2008.

An Organic Thin-film Transistor Circuit for Large-area Temperature-sensing

D. He, I. Nausieda, K. Ryu, A. I. Akinwande, V. Bulović, C. G. Sodini
Sponsorship: SRC/FCRP C2S2, Hewlett-Packard, NSERC Fellowship

The organic thin-film transistor (OTFT) is a field-effect transistor technology that uses organic materials as the semiconductor. OTFTs have field-effect mobilities that are comparable to those of hydrogenated amorphous silicon TFTs, and OTFTs are compatible with large-area and mechanically-flexible substrates [1], [2]. The goal of this work is to demonstrate an integrated OTFT temperature-sensing circuit suitable for large-area and flexible substrates.

As shown in Figure 1, two important differences are observed between the OTFT's and the MOSFET's current-voltage characteristics when temperature is varied. First, the OTFT's current increases with temperature in both subthreshold and above-threshold regimes, whereas the MOSFET's above-threshold current decreases with temperature. Second, the OTFT's subthreshold slope is temperature independent over the measured range of -20 to 60°C, while the MOSFET's subthreshold slope is proportional-to-absolute-temperature (PTAT).

Because of these differences in temperature response, the OTFT temperature-sensing "ΔV_{BE} circuit" (Figure 2a) has a complementary-to-absolute-temperature (CTAT) response instead of an equivalent silicon circuit's PTAT response. The OTFT circuit is scaled to an array format to enable surface thermal sensing applications. As Figure 2b shows, the array consists of 3x3 temperature-sensing circuit cells of 1mm² each and is currently being characterized.

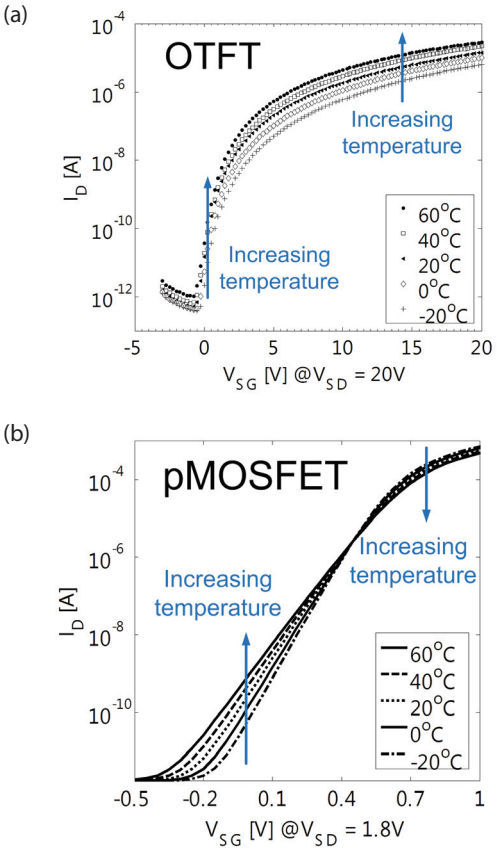

FIGURE 1: (a) The OTFT's (measured) and (b) pMOSFET's (BSIM3) current-voltage characteristics versus temperature.

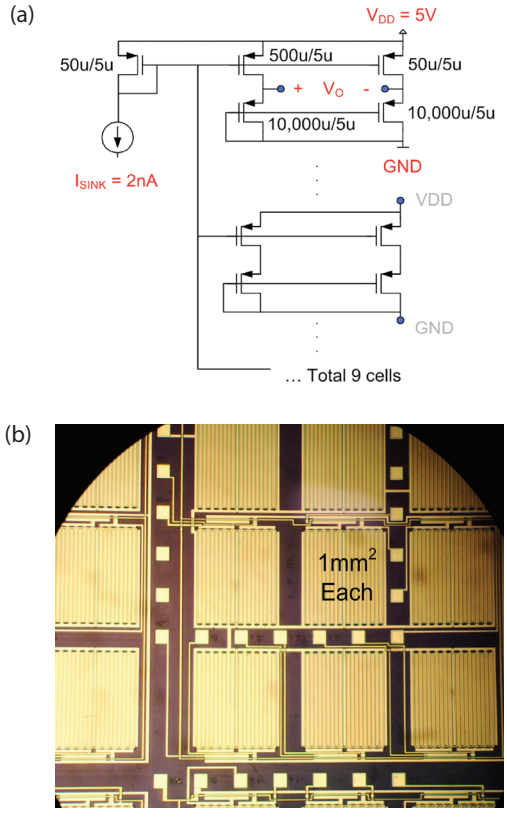

FIGURE 2: (a) The CTAT circuit array schematic and (b) die photo.

REFERENCES

[1] Y.Y. Lin, D.J. Gundlach, and T.N. Jackson, "High-mobility pentacene organic thin film transistors," *Device Research Conference Digest*, June 1996, pp. 80–81.

[2] I. Kymissis, C.G. Sodini, A.I. Akinwande, and V. Bulović, "An organic semiconductor based process for photodetecting applications," *IEEE International Electron Devices Meeting Technical Digest*, Dec. 2004, pp. 377–380.

Digital Phase-tightening for Millimeter-wave Imaging

K.M. Nguyen, C.G. Sodini
Sponsorship: SRC/FCRP C2S2

Millimeter-wave (MMW) imaging has potential applications such as collision-avoidance radar at 77GHz and concealed-weapons detection at 77GHz, 94GHz, and higher. This research investigates the challenges of designing an MMW imaging system. We envision an active imaging receiver that will consist of an array of 1000 antenna and per-antenna processor (PAP) units with an operating frequency of 77GHz [1]. Each PAP has digital logic that will estimate the phase and amplitude and reduce the data rate to the order of a kilohertz. A central processing unit (CPU) will perform digital beam-forming on the aggregated data from the array to achieve an expected frame rate of 10 fps. The 77-GHz input signal will be down-converted by a mixer with a 76-GHz local oscillator, generated by a PLL, to obtain an intermediate frequency (IF) of 1GHz. This signal is digitized by an analog-to-digital converter that is operating at 4.75GHz and is sent to the CPU.

Key blocks in the PAP being explored are the 76-GHz PLL and 4.75-GHz DLL. Since accurate beam-forming requires precise control of the phase over the array of elements, the PLL will be designed for minimal phase noise and power dissipation. The 76-GHz VCO used in the PLL is a cross-coupled LC tank design. The divider chain consists of nine divide-by-2 static frequency dividers created in emitter-coupled logic for the first six stages and CMOS logic for the lowest three stages. The highest frequency divider utilizes inductive peaking for increased operating frequency. A full PLL was designed in a 130-nm BiCMOS process. Test results showed a free-running oscillation frequency of 70GHz and gave improved understanding between modeling and experimental results. The next chip will contain a revised PLL and a phase tightening system, which will consist of a DLL, ADC, and digital logic to estimate the phase data. The ADC is clocked by the DLL, which outputs a clock signal at each of its delay cells. The phase-tightening system locks onto the IF signal's zero crossings by continually selecting the DLL output to clock the ADC.

REFERENCES

[1] A. Accardi, "Digital Compensation in Imaging Systems," Dissertation proposal, Massachusetts Institute of Technology, Cambridge, 2007.

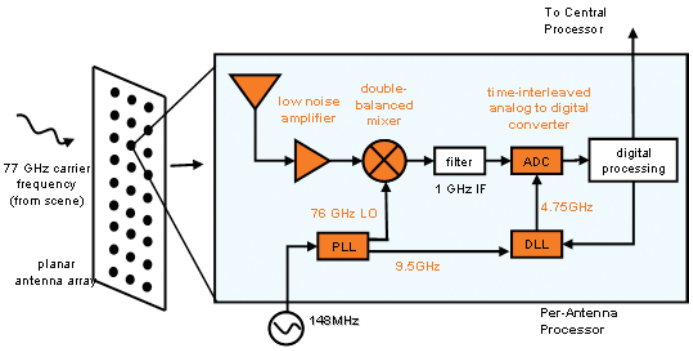

FIGURE 1: Block diagram of the MMW imaging system.

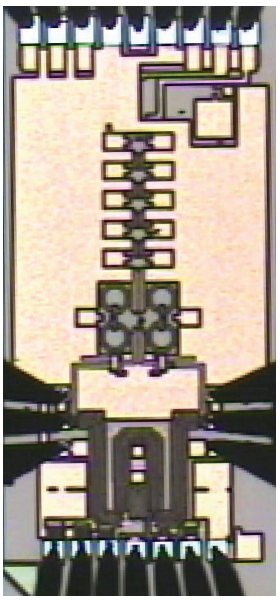

FIGURE 2: Die photo of 76-GHz PLL with on-wafer test probes.

Flip-chip Integrated Wideband Antennas for Millimeter-wave Passive Imaging

J. D. Powell, C. G. Sodini
Sponsorship: SRC/FCRP C2S2

The area of Millimeter-wave (MMW) system research and design has become increasingly popular in recent years, as advanced silicon processes have enabled integrated circuit operation in the MMW regime. Several applications exist for MMW design, including wireless communications at 60-GHz, collision-avoidance radar imaging at 77-GHz, and concealed-weapons detection imaging at 94-GHz and higher. This research focuses on a passive imager front-end that has been developed and tested for the application of concealed-weapons detection.

A key component of this research involves the design of a packaged antenna. A wideband Vivaldi-type design is used to achieve high gain and efficiency from 73-GHz to 105-GHz. The antenna was fabricated on a low-conductive dielectric constant material, RO4350B, at the MIT EML Laboratory. It is packaged onto the LNA bondpad terminals via gold solder bumps and silver epoxy and placed to maximize the distance from the die's ground plane. The antenna has a gain of approximately 8dB within the operating bandwidth and a typical efficiency of 80% [1],[2].

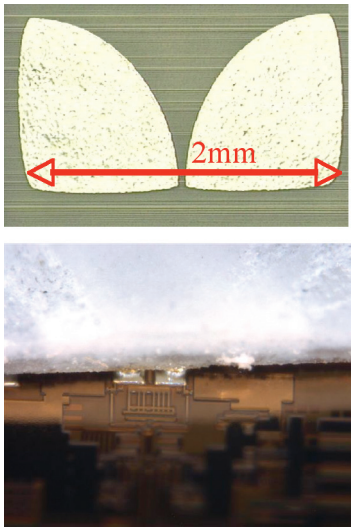

FIGURE 1: Photo of fabricated antenna and antenna mounted on LNA die using gold solder bumps.

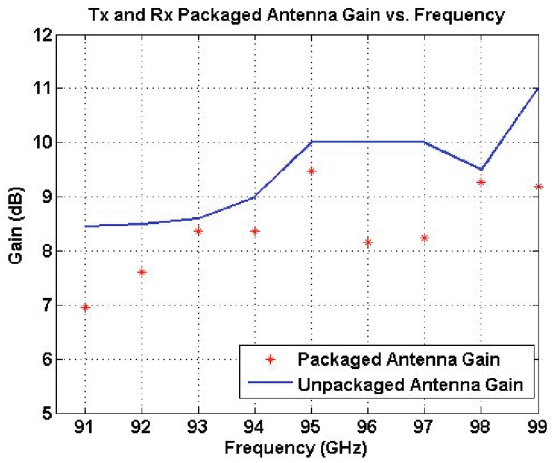

FIGURE 2: Measurement results for packaged versus unpackaged antenna gain.

REFERENCES

[1] J. D. Powell, "SiGe Receiver Front Ends and Flip-Chip Integrated Wideband Antennas for Millimeter-wave Passive Imaging," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, 2008.

[2] J. D. Powell, H. Kim, and C. G. Sodini, "SiGe Receiver Front Ends for Millimeter-wave Passive Imaging," IEEE Transactions on Circuit Systems-I, vol. 56, no. 3, pp. 541-553, March 2009.

A Leadless, Long-term ECG Monitor for In-home Use

E. Winokur, C. G. Sodini

Sponsorship: Texas Instruments

With the escalating costs of hospital visits, clinicians are opting to use at-home monitoring devices to diagnose patients. Current ECG Holter monitoring devices typically have 24 - 48 hour memory and battery capacity [1]. With many patients experiencing intermittent heart problems that can occur once every week – month, the Holter monitor is not a good solution and an event recorder or loop recorder is required [2]. However, each of these recorders can only save up to a few minutes of ECG recordings. This leads to the loss of most of the data, which could be very important in alerting the user for the onset of future episodes. Therefore, we have developed a Holter monitor prototype with the goal of battery and memory capacity of two weeks. Fig. 1 shows a block diagram of the system.

We based the long-term monitor prototype around a Texas Instruments MSP430 low-power microcontroller which enables high computing power with very low power consumption. The prototype monitor, which is currently being designed, will be mounted on standard 3M 2560 Red Dot electrodes and fabricated on a flexible PCB substrate. Mounting the PCB directly to the electrodes will improve the SNR by an estimated 40 dB compared to using wired leads [3]. The monitor will be 'L' shaped with rounded corners and placed on the patient's chest. The 'L' shape will enable several mounting sites to be placed on the board which will allow the doctor to choose which measurement he would like to record. The monitor will have 320 Mbytes of FLASH memory which is enough to store two weeks of data sampled at 250 Hz continuously. Total power consumption of the system is estimated to be less than 8mW.

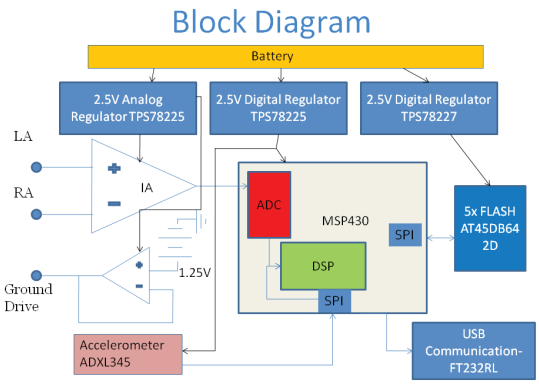

FIGURE 1: Block diagram of the ECG long-term Holter monitor system. The front end will use TI OPA333 and INA333 amplifiers and will have a bandwidth from 0.5 Hz to 125 Hz. 3 axis ADI accelerometer data will be sampled at 3Hz to help correlate activity level with the ECG recordings. The battery is a 3.7V 450 mAh Li-Pol cell from Cameron-Sino.

REFERENCES

[1] A. M. Abdulla. (2009, April 5, 2009). Holter monitor. *2009(April/21)*, pp. 1.

[2] M. A. Rockx, J. S. Hoch, G. J. Klein, R. Yee, A. C. Skanes, L. J. Gula and A. D. Krahn. (2005, Nov. 2005). Is ambulatory monitoring for "Community-acquired" syncope economically attractive? A cost-effective analysis of a randomized trial of external loop recorders versus holter monitoring. *AHJ 150(5)*, pp. 1065.e1-1065.e5.

[3] A. Searle and L. Kirkup, "A Direct Comparison of Wet, Dry and Insulating Bioelectric Recording Electrodes," *Physiol. Meas.*, vol. 21, pp. 271-283, 2000.

Energy Efficient On-Chip Equalized Interconnect

B. Kim, V. Stojanović

Sponsorship: SRC/FCRP IFC, Intel Corporation, CICS, SRC, Trusted Foundry

In recent high performance processor design, long distance interconnects became a serious bottleneck under tight power constraint [1]. Equalized on-chip interconnects have been presented significant improvement in performance without sacrificing too much power consumption [2,3].

This work presents further improvement of energy efficiency of equalized-interconnect by proposing two circuit techniques: 1) pre-distorted charge-injection (CI) feed-forward equalization (FFE); 2) trans-impedance (TIA) termination at receiver. Instead of using traditional analog subtraction, CI-FFE injects pre-computed the current value required for FFE into the channel while mitigating the nonlinearity of the driver. The non-linearity of the driver is statically compensated by pre-distorting FFE coefficients. The trans-impedance amplifier terminated at the receiver improves the bandwidth, signal amplitude, and reduces bias current.

A test-chip is fabricated in 90nm CMOS process and consumed about 0.4pJ/b running at 4Gb/s with vertical eye opening about 100mV and horizontal eye opening 50%UI differential peak-to-peak [4].

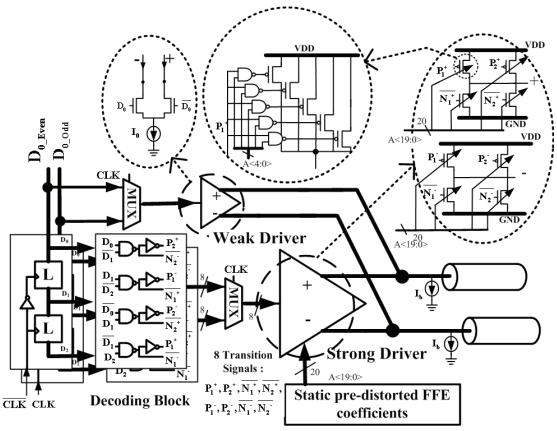

FIGURE 1: 3-tap Pre-distorted Charge-injection Feed-Forward Equalizing Transmitter

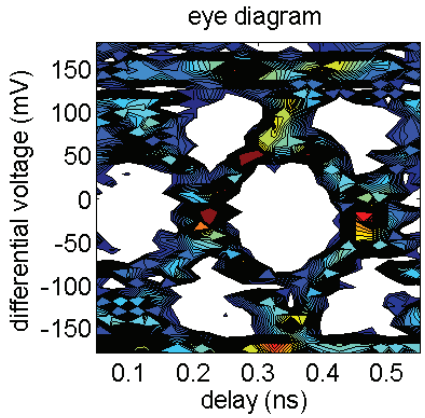

FIGURE 2: Measured eye diagram *in-situ*.

REFERENCES

[1] Kumar, R.; Zyuban, V.; Tullsen, D. M., "Interconnections in multi-core architectures: understanding mechanisms, overheads and scaling," in *Proc. 32nd International Symposium on Computer Architecture*, 2005.

[2] A.P. Jose, G. Patounakis and K.L. Shepard "Near speed-of-light on-chip interconnects using pulsed current-mode signaling," *VLSI Circuits, 2005. Digest of Technical Papers. 2005 Symposium on* no. SN -, pp. 108-111, 2005.

[3] Schinkel, D., Mensink, E., Klumperink, E.A., Tuijl, E. van and Nauta, B. "A 3-Gb/s/ch transceiver for 10-mm uninterrupted RC-limited global on-chip interconnects," *Solid-State Circuits, IEEE Journal of* vol. 41, no. 1 SN - 0018-9200, pp. 297-306, 2006.

[4] Kim, B. and Stojanović, V., "A 4Gb/s/ch 356fJ/b 10mm Equalized On-chip Interconnect with Nonlinear Charge-Injecting Transmit Filter and Transimpedance Receiver in 90nm CMOS Technology," in *Proc. International Solid-State Circuit Conference*, 2009.

A Fractionally Spaced Linear Receive Equalizer with Voltage-to-time Conversion

S. Song, B. Kim, V. Stojanović
Sponsorship: National Semiconductor, CICS

Based on voltage-to-time conversion technique [1-2], a pseudo-differential two-way-interleaved adaptive linear receive equalizer with two 2x-oversampled feed-forward taps has been designed in a 90-nm CMOS process. It integrates equalization and phase interpolation functions into one unit to simultaneously address inter-symbol-interference (ISI) cancellation and phase synchronization in a link receiver.

Due to the process speed limitation, the half-rate time interleaving technique is also applied (Figure 1). Four sampling phases (Φ_1- Φ_4) with 25% duty cycle are generated locally from Φ and Φ_- and another pair of quadrature clocks. A voltage-to-time (V2T) block converts the sampled signal into a delayed digital signal, transferring the sampled information into the time-domain. All four V2T converters are followed by a time-to-voltage (T2V) stage to realize summing, subtraction and multiplication. Equalizer tap weights are implemented as two programmable reference currents I_1 and I_2 biasing T2V blocks. Two slicers with tunable thresholds are added to sense the signs of the input signal and output error of the FSE, respectively, and to enable tap weight adaptation with external adaptive engine.

The design is fabricated in a 90-nm CMOS process. It operates at 4 Gbps with 8 mW power consumption and linearity of 4.3 effective bits at 1.2 V supply (Figure 2).

REFERENCES

[1] L. Brooks and H.-S. Lee, "A Zero-Crossing-Based 8b 200MS/s Pipelined ADC," *Solid-State Circuits Conference, 2007. ISSCC 2007. Digest of Technical Papers. IEEE International*, pp. 460-615, 2007.

[2] J. F. Bulzacchelli, M. Meghelli, S. V. Rylov, W. Rhee, A. V. Rylyakov, H. A. Ainspan, B. D. Parker, M. P. Beakes, Aichin Chung, T. J. Beukema, P. K. Pepeljugoski, L. Shan, Y. H. Kwark, S. Gowda and D. J. Friedman, "A 10-Gb/s 5-Tap DFE/4-Tap FFE Transceiver in 90-nm CMOS Technology," *Solid-State Circuits, IEEE Journal of*, vol. 41, pp. 2885-2900, 2006.

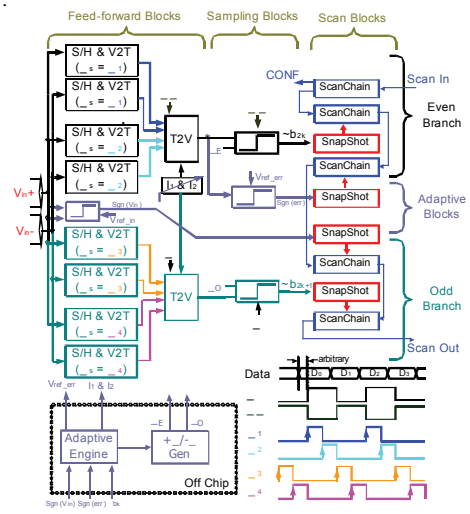

FIGURE 1: Block diagram of 2-tap 2xoversampled two-way time interleaved architecture. Scan-chain and snapshot are applied for ***in-situ*** link characterization.

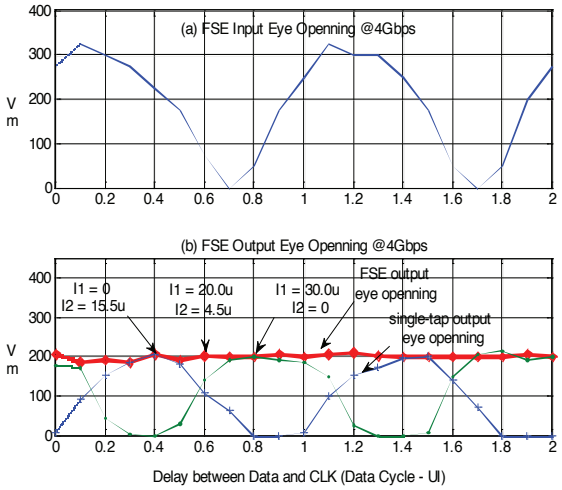

FIGURE 2: Input (a) and FSE output (b) eye opening vs. sampling phase.