1

Circuits & Systems

Table of Contents

Contact-resistance Variation in Advanced Technologies

K. Balakrishnan, D.S. Boning

Sponsorship: SRC/FCRP IFC, Samsung

Due to the continuous and aggressive scaling of CMOS technology, the parameters and characteristics that are most critical in determining the quality and robustness of a device are changing. One such trend is the impact of silicide-diffusion contact resistance, which is becoming an increasingly larger component of the total resistance in a MOSFET [1]. Consequently, variability in contact resistance must now be considered in order to accurately capture the robustness of a transistor. The primary goal of this work is to design and fabricate a test chip that will characterize the variability of contact resistances by measuring a large array of devices under test. Completing the characterization of contact resistances will allow for the generation of a compact model that incorporates sensitivities to various parameters into the determination of individual contact resistances. Finally, this methodology can be generalized to investigate contacts in novel technologies, such as carbon nanotube-based contacts.

A test chip has been designed and will be fabricated using an IBM 90-nm process to investigate this variability and gather statistics regarding its characteristics. Figure 1 shows a general setup with which the resistance of a contact can be measured using the equivalent of a four-point probing Kelvin-based structure. This scheme is multiplexed across multiple devices to gather large amounts of data for variability analysis (see Figure 2). In this test chip, the silicide-diffusion contact resistances are measured for a large array of transistors using a multiplexed current-force voltage-sense methodology. The design of experiments includes different combinations of values for parameters such as contact-to-gate distance, contact-to-diffusion edge distance, and the number of contacts on the sources and drains of devices. In addition, the test chip also contains structures that can gather contact resistance information as well as device characteristics to decouple variability information.

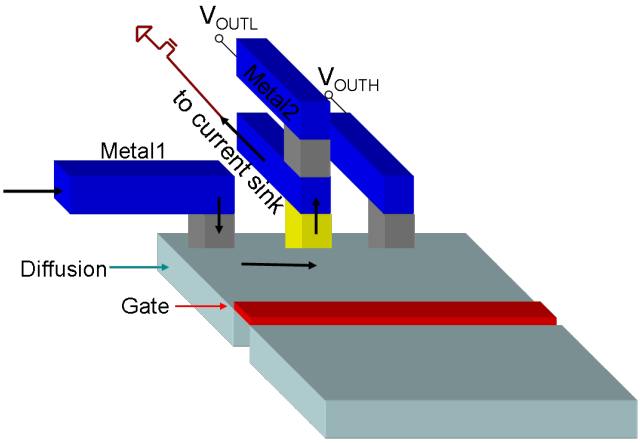

▲ Figure 1: Current-force voltage-sense methodology to determine resistance of contact highlighted in yellow. With the transistor turned off, the current forced through metal line will go through the contact, and then the resistance is directly proportional to $V_{OUTH} - V_{OUTL}$.

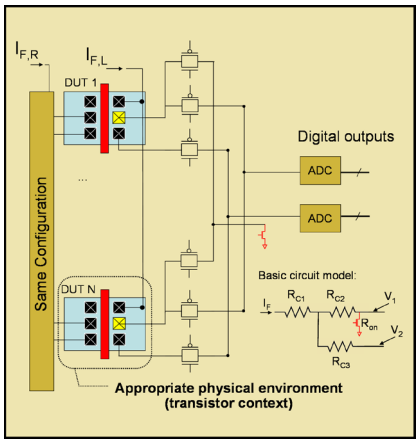

▲ Figure 2: Multiplexing strategy to characterize a large number of silicide-diffusion contacts. A sink device is used to carry all the input current. With off-chip ADCs, measurement accuracy and test chip area are both maximized, while the losses with respect to increased noise are minimal because of DC measurements.

Reference

[1] S.D. Kim, C.M. Park, and J.C.S. Woo, "Advanced model and analysis of series resistance on CMOS scaling into nanometer regime – Part 2: Quantitative analysis," *IEEE Transactions on Electron Devices*, vol. 49, pp. 467-472, Mar. 2002.

Test Circuits for $I_{D(sat)}$ Variation Characterization

A.H. Chang, D.S. Boning
Sponsorship: TSMC

The variation in process, device, and circuit is an increasingly difficult and critical concern for future integrated circuit design. Variation may have systematic components such as spatial, device-size, or pattern-density dependency, or it may have random stochastic components such as random doping fluctuations. Increasing effort is needed in the modeling and characterizing of device variations in order to design circuits robustly.

Figure 1 shows three common types of measuring schemes. Type *A* is the simplest approach, in which direct probing is used for full electrical characterization of a device. The Type A measuring scheme requires dedicated pads for probing; therefore, only a small number of probing terminals is available. The probing may also introduce extra stress on the device, which could change the device characteristic after measurement. Type B uses multiplexed circuitry to extract DC current-voltage (I-V) measurements [1, 3, 4]. This pad-sharing scheme is more efficient, enabling characterization of more devices than in Type A. Type C is a dedicated circuit structure focused on extracting a specific device parameter, such as V_T [2].

Figure 2 illustrates a systematic variation of saturation current across a single chip schematically. Several sources may contribute to observed $I_{D(sat)}$ variation, and each must be considered and understood. In deeply scaled technologies, variation in V_T is believed to be mostly due to random dopant fluctuation [3] and would not explain systematic dependencies or neighborhood pattern dependencies. Different STI pattern densities can be explored to see if the resulting device variations show a clear systematic layout density offset. This component is of particular interest, as it is conjectured that stress or thermal annealing process effects may contribute to $I_{D(sat)}$ variation. While previous research efforts have explored the variation and layout dependency of V_T and channel length individually, few have focused on the effect of stress or annealing related mobility variation [1-5]. In this project, our goal is to design test-circuit approaches to isolate the device variation parameters (V_T, L, and μ) and dependencies for future technologies, focusing especially on $I_{D(sat)}$ variation. A set of design rules and guidelines can then be formulated to minimize these variations.

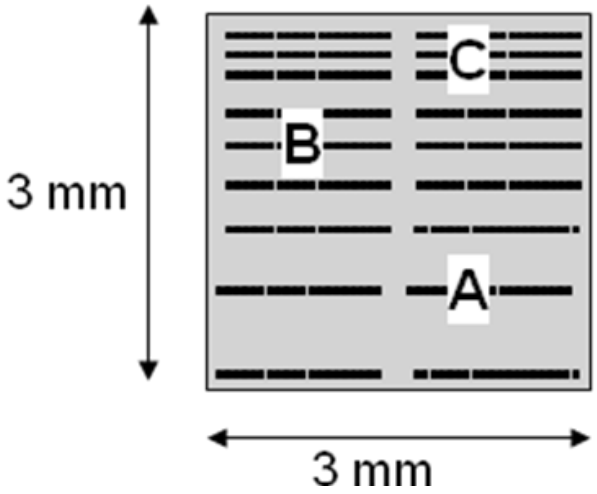

▲ Figure 1: Three common types of measuring schemes. Type A is direct probing; Type B is multiplexed I-V measurement structure; and Type C is customized circuits for specific parameter extraction.

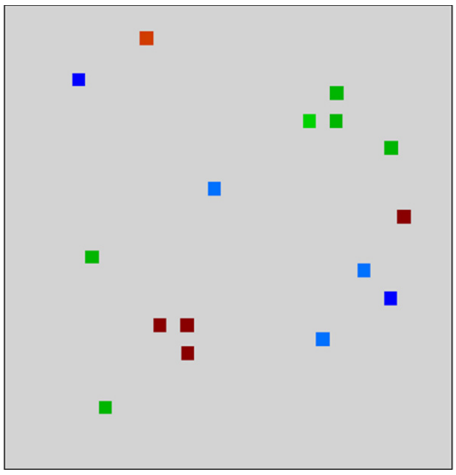

▲ Figure 2: Hypothesized clustering of saturation current across a die. The effect of local and neighborhood layout practices (transistor feature sizes, STI pattern density, etc.) on systematic deviations in $I_{D(sat)}$, and in variance of saturation current, need to be understood. Test structures and circuits to identify and separate sources of variation are being developed.

References

[1] K.G. Gettings, "Study of CMOS process variation by multiplexing analog characteristics," Ph.D. Thesis, MIT, Cambridge, June 2007.

[2] N. Drego, A. Chandrakasan and D. Boning, "A test-structure to efficiently study threshold-voltage variation in large MOSFET arrays," *IEEE International Symposium on Quality Electronic Design*, Mar. 2007, pp. 281-286.

[3] K. Agarwal, S. Nassif, F. Liu, J. Hayes, and K. Nowka, "Rapid characterization of threshold voltage fluctuation in MOS devices," *International Conference on Microelectronic Test Structures*, pp. 74-77, 2007.

[4] L.T. Pang and B. Nikolic, "Impact of layout on 90nm CMOS process parameter fluctuations," *Symposium on VLSI Circuits, Digest of Technical Papers*, 2006, pp. 69-70.

[5] P. Friedberg, Y. Cao, J. Cain, R. Wang, J. Rabaey, and C. Spanos, "Modeling within-die spatial correlation effects for process-design co-optimization," *Proc. of the Sixth International Symposium on Quality Electronic Design*, Mar. 2005.

Process Variation in High-speed RF Front-end Circuits

D. Lim, D.S. Boning

Sponsorship: IBM, SRC/FCRP C2S2

Modern RF circuit designs in deeply scaled CMOS technologies need efficient characterization of process-induced variation in performance variables and robust optimization methods to obtain high-yielding chips. A PLL front-end consists of a VCO and 2:1 frequency divider. Since the PLL front-end operates at the highest frequency in the system, speed, power consumption, and noise characteristics must be considered carefully during design. Furthermore, the tuning range of the VCO is significantly narrowed by the reduced ratio between varactor and parasitic capacitance in scaled devices. The operating range of the frequency divider is also limited since input signal power is reduced due to the loss in interconnects. Recently, 22.8% 3σ/μ variation in the self-oscillation frequency of dividers in 65-nm SOI technology has been reported [1]. The mismatch between the operating ranges of the VCO and the frequency divider due to process variation can cause a serious yield problem.

A 70-GHz mm-wave PLL front-end with an LC-VCO and 2:1 CML frequency divider has been implemented as in Figure 1. The variation in the VCO and divider performance has been measured and their correlation has been estimated. The variation in the frequency divider is critical for the functionality of the PLL front-end circuit, and the bias condition of the divider has been optimized for the highest functional yield while achieving the maximum output signal-to-noise ratio [2]. Additionally, the estimation of the variation in the maximum operating frequency of the divider is extremely time-consuming since input frequency must be swept by a small step for fine resolution in each Monte Carlo simulation run. We suggest a new time delay model for a CML differential buffer that is based on a quadratic function of physical delay components. The suggested model estimates the maximum operating frequency variation with roughly 1% error in 500 Monte Carlo runs, as Figure 2 shows.

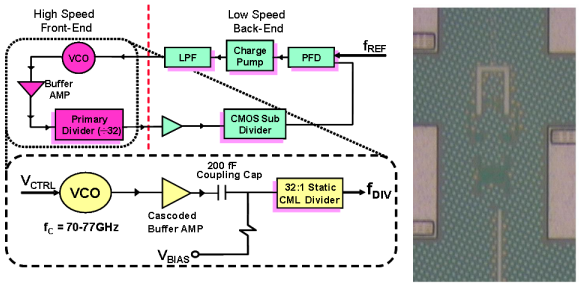

▲ Figure 1: A 70-GHz mm-wave PLL diagram and high-speed front-end building blocks, including a VCO and a frequency divider with external bias control. The circuit is fabricated in 65-nm SOI CMOS and the die photo is shown.

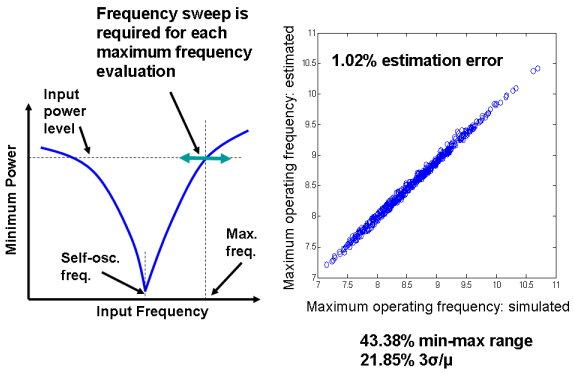

▲ Figure 2: An input sensitivity curve of a frequency divider and scattering plot of simulated and estimated maximum operating frequency of the divider. The maximum operating frequency simulation is extremely time-consuming for frequency and power sweeping; the suggested model provides around 1% estimation error.

Reference

[1] D. Lim, J. Kim, J.-O. Plouchart, C. Cho, D. Kim, R. Trzcinski, and D. Boning, "Performance variability of a 90-ghz static CML frequency divider in 65-nm SOI CMOS technology," *IEEE ISSCC Digest of Papers*, Feb. 2007, pp. 542-621.

[2] D. Lim, J. Kim, J.-O. Plouchart, C. Cho, D. Kim, and D. Boning, "Performance and yield optimization of mm-wave PLL front-end in 65-nm SOI CMOS," *IEEE RFIC Symposium Digest of Papers*, June 2007, pp. 525-528.

Spatial Variation in Advanced CMOS Devices and Circuits

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

Modern circuit designs in deeply scaled technologies need efficient characterization of process variation in order to obtain high-yielding chips. Circuit designers need accurate guidelines to prevent failures due to layout-induced variations. We address one specific portion of this need by characterizing and quantifying spatial variation trends in both device parameters and circuit performance.

At the device level, we have implemented a test-structure capable of efficiently measuring leakage currents of a large number of devices (~70K NMOS + ~70K PMOS). All devices are minimum length and range in width from minimum width to 3.0μm. By measuring leakage currents, we are able to extract mismatches in threshold voltage between two devices, ΔV_T [1]. Measured leakage currents and extracted ΔV_T from the test chip indicate no within-die spatial correlation and no inter-die correlation in the variation pattern. These results indicate that a truly random process, such as Random Dopant Fluctuation (RDF), is the dominant source of threshold voltage variation. Furthermore, this lack of correlation poses problems in low-power, low-voltage systems where the effect of V_T variation is dominant. In particular, sub-threshold circuits are most susceptible to completely random variation. Simulations show that even with strong correlation in channel length variation, as the operating voltage of a circuit decreases below twice the nominal V_T of the process, correlation in circuit performance decreases quickly.

To quantify this correlation in circuit performance, a test-chip has been implemented (Figure 1) containing adder delay paths in an oscillating configuration with simple asynchronous counters to measure frequency. A Phase-Frequency Detector and random sampling technique to quantify delay variation in individual bits of each adder are also included. Monitor ring-oscillator circuits are also laid out to characterize correlation between monitor circuits and actual critical paths of 64-bit Kogge-Stone adders. Preliminary results indicate only weak within-die spatial correlation between adjacent adders as well as between adders and adjacently placed monitor ring-oscillators. Figure 2 is a plot of the adder frequencies on a single-die showing no spatial trend. However, die-to-die correlation is high (~0.9) but degrades with decreasing V_{DD}, as predicted by the simulations performed previously. On-going data gathering will provide additional statistical confidence in these measurements.

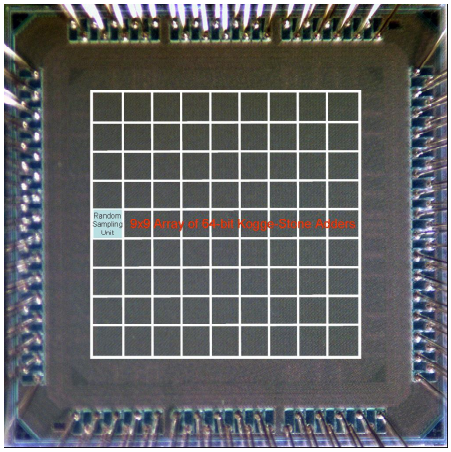

▲ Figure 1: Test-chip containing eighty 64-bit Kogge-Stone adders in an oscillating condition along with monitor ring-oscillators and frequency counters. A random-sampling technique is also implemented to measure delays between individual bits of the adders.

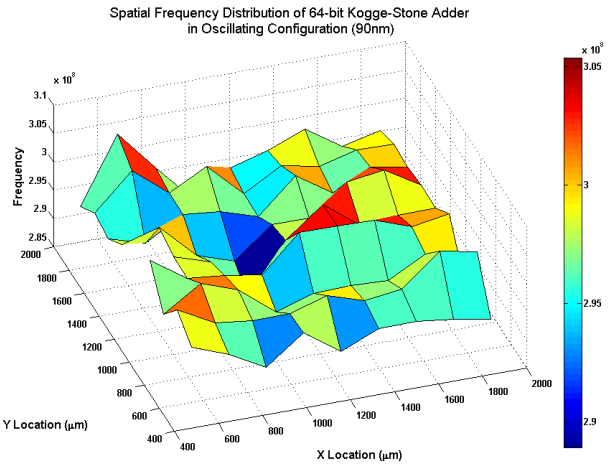

▲ Figure 2: Spatial distribution of 80 Kogge-Stone adders on a single die showing no spatial trends. Further statistical analysis reveals only weak within-die spatial correlation but strong die-to-die correlation, degrading with decreasing operating voltage

Reference

[1] N. Drego, A. Chandrakasan, and D. Boning, "A test structure to efficiently measure V_T-variation in large MOSFET arrays," *International Symposium on Quality Electronic Design*, Mar. 2007.

A Stochastic Integral Equation Solver for Efficient Variation-aware Interconnect Extraction

T. Moselhy, L. Daniel

Sponsorship: Cadence Design Systems, SRC/FCRP IFC, SRC

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 technological scaling, such manufacturing uncertainties are now beginning to play a major role in determining the electrical characteristics of the interconnect structures. Consequently, variation-aware interconnect extraction is becoming increasingly important.

In this research we have developed a new methodology 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 also can 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 at any time (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 parallizable and can therefore take advantage of the state of the art in processor design. 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 then 8 hours using Matlab on a standard 4-core machine and using only 121 MB RAM.

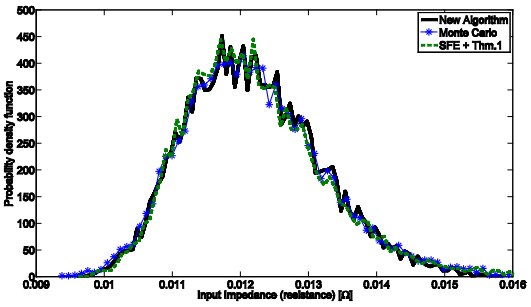

▲ Figure 1: Comparison between the probability density function of the microstrip line obtained from our new algorithm and the reference Monte Carlo simulation.

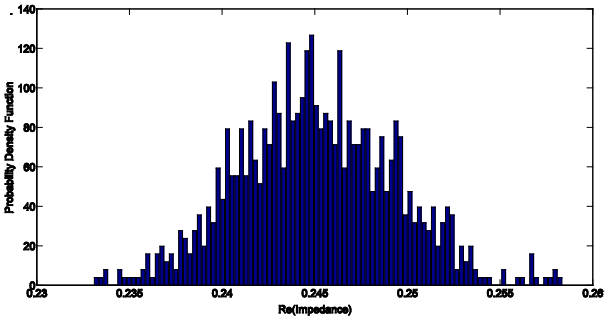

▲ Figure 2: Probability density function of the real part of the input impedance at 1GHz for correlation length Lc = 50μm. The resistance of the non-rough surface is 11.3% smaller than the mean of the obtained distribution.

References

[1] T. Moselhy and L. Daniel, "Stochastic Integral Equation Solver for Variation-Aware Interconnect Extraction," presented at *Design Automation Conference*, June 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*. Springer-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*, Nov. 2005, pp. 675-682.

A Piecewise-linear Moment-matching Approach to Parameterized Model-order Reduction for Highly Nonlinear Systems

B. Bond, L. Daniel

Sponsorship: SRC/FCRP, NSF, DARPA

The automatic extraction of parameterized macromodels for modern mixed-signal System-on-Chips is an extremely challenging task due to the presence of several nonlinear analog circuits and Micro-Electro-Mechanical (MEM) components. The ability to generate Parameterized Reduced Order Models (PROMs) of nonlinear dynamical systems could serve as a first step toward the automatic and accurate characterization of geometrically complex components and subcircuits, eventually enabling their synthesis and optimization.

Our approach to this problem combines elements of a non-parameterized trajectory piecewise-linear method [1] for nonlinear systems with a moment-matching parameterized technique [2] for linear systems. By building on these two existing methods, we have created an algorithm for generating PROMs for nonlinear systems. The algorithms were tested on three different systems: a MEM switch, shown in Figure 1, and two nonlinear analog circuits. All of the examples contain distributed strong nonlinearities and possess dependence on several geometric parameters.

In addition, we have proposed a model-construction procedure in which we approximate the system sensitivity to parameters of interest for the purpose of efficiently sampling important regions of the parameter space. Figure 2 shows the output of one PROM created for the example in Figure 1 and compared to the field solver output of the full nonlinear system at several parameter values. Typical PROMs constructed in this manner can be accurately reduced in size by a factor of 10, yielding a speedup of a factor of 10 in general. For further details on parameter-space accuracy and cost of the algorithms, see [3].

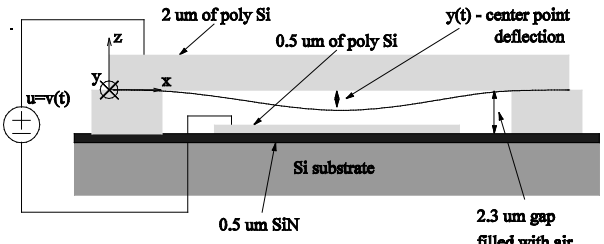

▲ Figure 1: Application example: MEM switch realized by a polysilicon beam fixed at both ends and suspended over a semiconducting pad and substrate expansion.

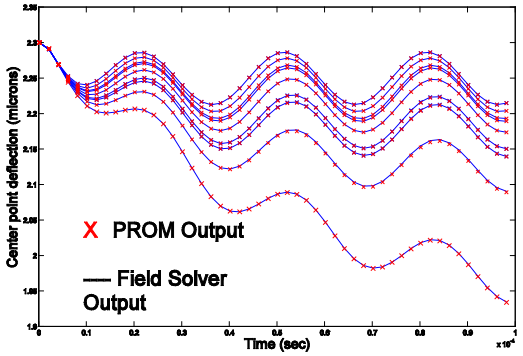

▲ Figure 2: Center point deflection predicted by our parameterized reduced model (crosses) at a series of parameter values, compared to a finite difference detailed simulation (solid lines).

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] L. Daniel, C.S. Ong, S.C. Low, K.H. Lee, and J.K. White, "A multiparameter moment-matching model reduction approach for generating geometrically parameterized interconnect performance models," *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, vol. 23, no. 5, pp. 678-693, May 2004.

[3] 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.

Stable Model-order Reduction for Highly Nonlinear Systems

B. Bond, L. Daniel
Sponsorship: SRC/FCRP, NSF, DARPA

The ability to generate accurate reduced-order models (ROMs) of nonlinear dynamical systems, such as analog circuits and micro-electromechanical 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 and set of weighting functions, which together allow us to extend some of these stability results to PWL systems composed of arbitrary unstructured matrices.

The stability of nonlinear systems is determined 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 a system's Jacobians all 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 and new set of interpolation functions to create a collection of stable nonlinear systems. The final ROM will switch between the various stable nonlinear ROMs. 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].

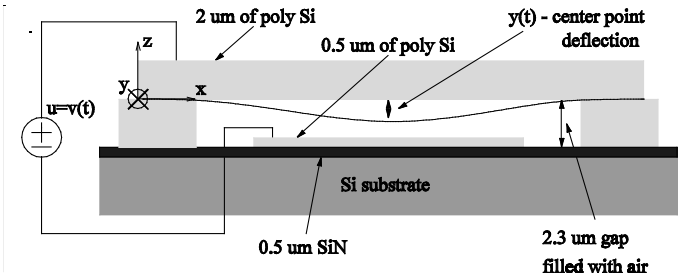

▲ Figure 1: Application example: MEM switch realized by a polysilicon beam fixed at both ends and suspended over a semiconducting pad and substrate expansion

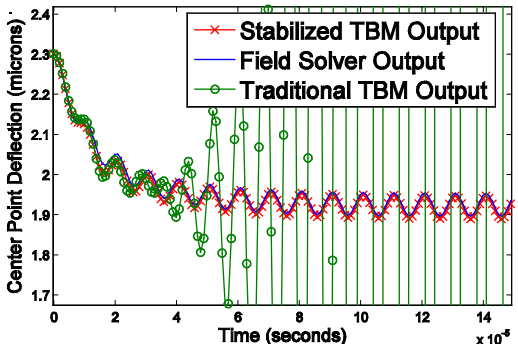

▲ Figure 2: Center point deflection predicted by our stabilized reduced model (red crosses), compared to a finite difference detailed simulation (solid blue lines) and the traditional TBM approach (green circles).

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.

Convex Relaxation Approach to the Identification of the Wiener-Hammerstein Model for Modeling of Non-Linear Analog Circuit Blocks

K.C. Sou, A. Megretski, L. Daniel
Sponsorship: DARPA, SRC/FCRP IFC

Analog and mixed/signal VLSI circuits exhibit an ever-increasing and pressing need for automatic and accurate characterization of their non-linear components and subcircuits, in order to enable synthesis and optimization. While non-linear model order reduction has already been attempted using several types of clever linearizations or parametric approximation, in this project we are instead attempting a completely orthogonal approach. Specifically, this work proposes an input/output system identification technique for the Wiener-Hammerstein model and its feedback extension. In the proposed framework, the identification of the nonlinearity is non-parametric. The identification problem can be formulated as a non-convex quadratic program (QP). A convex semi-definite programming (SDP) relaxation is then formulated and solved to obtain a sub-optimal solution to the original non-convex QP. The convex relaxation turns out to be tight in most cases. When the relaxation idea is combined with the use of local search, high-quality solutions to the Wiener-Hammerstein identification can frequently be found. We identify randomly generated Wiener-Hammerstein models as examples of the application. Furthermore, we are attempting to use our method to identify small analog circuit blocks such as operational amplifiers. This work has been accepted for publication and will be presented at the Conference on Decision and Control inDecember 2008. [1]

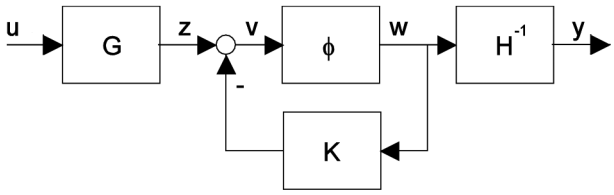

▲ Figure 1: Wiener-Hammerstein model to be identified.

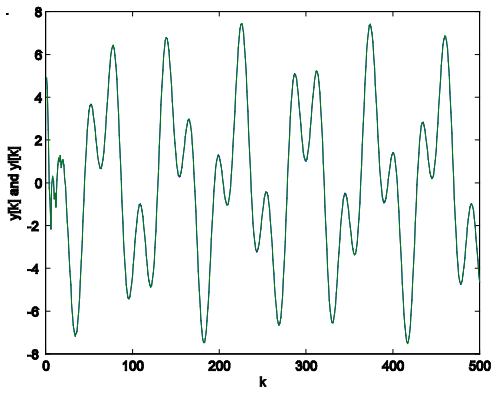

▲ Figure 2: Matching of outputs of the original system and the identified system.

Reference

[1] K.C. Sou, A. Megretski, and L. Daniel, "Convex relaxation approach to the identification of the Wiener-Hammerstein Model," to be presented at *IEEE Conference on Decision and Control*, Dec. 2008.

A Hierarchical, Equation-based Design Methodology for Optimizing Mixed-signal Systems

T. Khanna, W. Sanchez, J.L. Dawson

This work shows a hierarchical, equation-based optimization strategy suitable for system design, as Figure 1 shows. Because it is a hierarchical methodology, it scales gracefully to systems that are much larger than can be handled by known optimization methods. The desired system is broken into circuit blocks, whose performance spaces are quantitatively described as Pareto-optimal surfaces between system design variables. Exploiting the surfaces' gentle nature and amenability to low-order equation fits, the trade-off surfaces are abstracted to higher levels as representations of the circuit block. Thus, resources (such as power dissipation, noise budget, gain, etc.) are allocated at the system level very rapidly and very efficiently using familiar equation-based optimization strategies.

We have achieved excellent matching between flat and hierarchical optimizations in the discrete implementation of both a transmitter front-end and a 10-stage pipeline ADC in a 0.18-μm CMOS process. Using equation-based optimization, we have quantitatively described the Pareto surfaces of each pipeline stage and abstracted them to the system level. The surfaces were modeled with monomial fits, which all had less than a 10% relative error. We obtain a 4x and 25x running-time improvement in the receiver and ADC examples, respectively, when using a hierarchical optimization, a clear advantage for larger scale systems. The optimize times appear in Figure 2.

The value of the Pareto surfaces lies in the compact global perspective they provide to system designers. Because the Pareto surfaces are well behaved, we can use low-order functions in the abstraction, resulting in a low-complexity system optimization. The proposed methodology restores the tractability of system-level design problems and is a powerful aid to designers of large, mixed-signal systems.

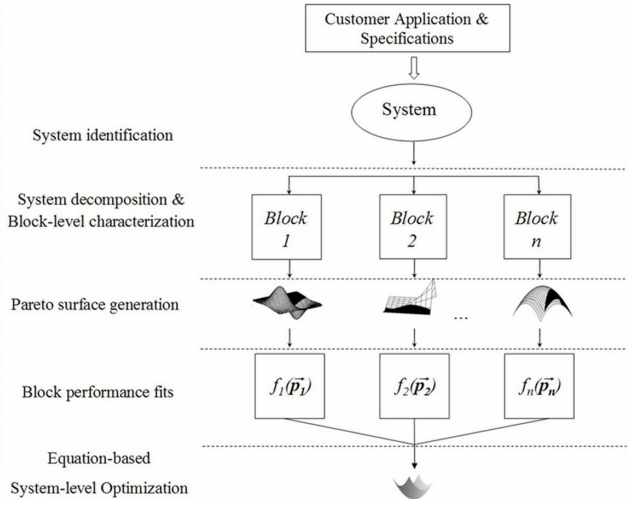

▲ Figure 1: Hierarchical bottom-up (H-BU) design methodology that shows the steps involved in the complete system design cycle.

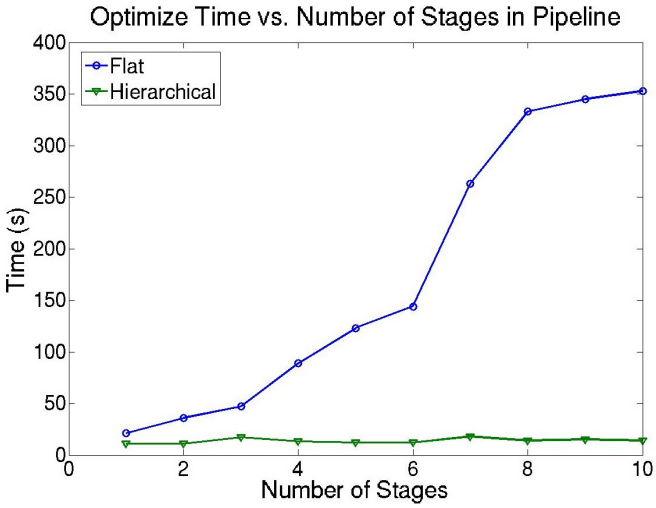

▲ Figure 2: Optimize times for the flat and hierarchical optimizations of pipeline ADC for varying number of stages. There is a 25x speed-up when using the hierarchy.

References

[1] M. Hershenson, S. Boyd, and T. Lee, "Efficient description of the design space of analog circuits," in *IEEE/Association for Computing Machinery Design Automation Conference*, June 2003.

[2] T. Eekeleart, T. McConaghy, and G. Gielen, "Efficient multiobjective synthesis of analog circuits using hierarchical Pareto-optimal performance hypersurfaces," *Design Automation and Test in Europe Conference Proc.*, pp. 1070-1075, June 2005.

[3] T. Eekeleart, R. Shoofs, G. Gielen, M. Steyeart, and W. Sansen, "Hierarchical bottom-up analog optimization methodology validated by a delta-sigma A/D converter design for the 802.11a/b/g standard," *IEEE/Association for Computing Machinery Design Automation Conference*, pp. 25-30, July 2006.

Iterative Robust Optimization of Analog Circuits

Y. Li, V. Stojanović

Sponsorship: SRC, CICS

As IC technologies scale down to the deep submicron region, process variation is becoming an increasingly severe issue for circuit designers. Designs are verified over process corners to improve the robustness of circuits and increase the manufacturing yield. However, corner-based robust design requires long design periods and often leads to overdesign. We are trying to develop new numerical algorithms fitted into an equation-based circuit optimization methodology [1], which incorporates the process variations, as well as provides yield estimation.

Inspired by the algorithm used in a robust taper design [2], we have developed and implemented the iterative robust optimization algorithm as shown in the left blocks in Figure 1. Rather than formulate the problem into a stochastic optimization problem as done in some previous work [3], we propose a more practical way. The optimization problem runs iteratively, with added robust constraints in each iteration. Thus, the optimization problem size grows, resulting in a more and more robust system. Relying on an existing optimization solver, we can solve the growing problem efficiently. As an example, a two-stage op-amp could be designed with robustness within minutes. The left flow in Figure 1 shows an outer loop around the iterative algorithm to generate a yield-aware design. The algorithm starts from a design with small process-variation range and we estimate the yield. The variation ranges keep growing until the yield reaches the desired value. This approach could enable rapid generation of trade-off surfaces for desired circuit blocks, parameterized by yield.

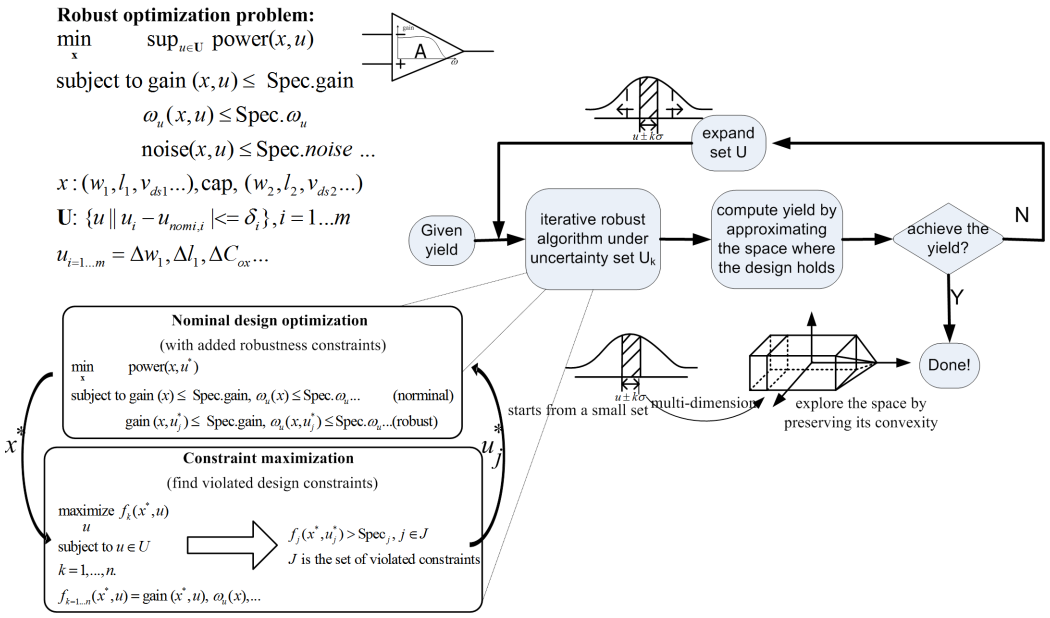

▲ Figure 1: Robust optimization problem formulation and the iterative robust optimization algorithm as shown on the left. Block diagrams on the right show the yield estimation flow.

References

[1] M. Hershenson, S.P. Boyd, and T.H. Lee, "Optimal design of a CMOS op-amp via geometric programming," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and System*, vol. 20, no. 1, pp. 1-21, Jan. 2001.

[2] A. Mutapcic and S. Boyd, "Robust design of slow-light tapers in periodic waveguides," to appear in *Engineering Optimization*, 2008.

[3] Y. Xu, X. Li, K. Hsiung, S. Boyd, and I. Nausieda, "OPERA: optimization with ellipsoidal uncertainty for robust analog IC design," in *Proc. of the 42nd Design Automation Conference*, pp. 632-637, June 2005.

On-chip, High-frequency Characterization of Carbon Nanotubes

F. Chen, A.P. Chandrakasan, V. Stojanović
Sponsorship: SRC/FCRP IFC

Measuring the high-frequency characteristics of nanoscale devices such as CNTs and nanowires is a critical step in determining their viability for semiconductor applications [1]. Previous efforts to measure high-frequency characteristics of CNTs have been limited by a handful of common problems. First, the traditional approach of using a network analyzer (VNA) to capture the frequency response is limited by the poor power transfer between the high impedance (> 10 kΩ) of the device and the 50 Ω test equipment termination). This impedance mismatch offsets the selective bandwidth of the VNA used to reduce the noise floor, resulting in a large variance of measured data due to signals being at or near the noise floor. Second, measurement parasitics from test probes and pads often dominate the reactance of the CNTs being measured, limiting both the accuracy of the results and the bandwidth of the measurement. Third, given the dimensions of CNTs, test setups are difficult to reproduce, limiting the range of lengths and number of CNTs that can be measured.

To address these issues, we have developed an on-chip test platform consisting of an array of 256 transceivers. Figure 1 shows a conceptual drawing of the CNT to CMOS test chip interface. Under each pad in the array is a transceiver that is independent of all others, allowing for measurement between any two pads in the array. Figure 2 shows the top level block diagram of two transceivers linked by a CNT "channel." Similar to [2] but with mostly on-chip components, the step response of the channel is captured by changing the threshold voltage of the sampler (VREF) and the relative phase of the receiver clock (RxClk) with respect to the transmit clock (TxClk). Each transceiver has an adjustable termination and employs a capacitance compensation technique to allow full-sized bond pads for device characterization at the chip interface while maintaining input drive bandwidths up to 1GHz for a 4 kΩ termination. A 20-bit counter accumulates samples at each point, to average out timing noise due to jitter and any dynamic voltage offsets in the sampler.

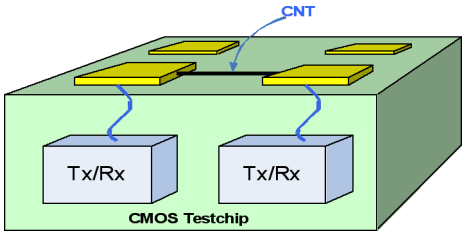

▲ Figure 1: Conceptual drawing of the CNT to CMOS test chip interface.

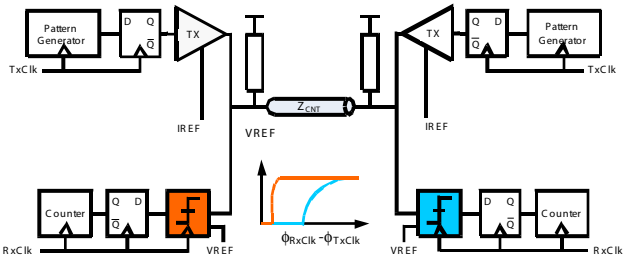

▲ Figure 2: Block diagram of 2 transceivers linked by a CNT and conceptual waveforms captured by shifting VREF and RxClk.

References

[1] P.J. Burke, et al., "Nanotube technology for microwave applications," *Microwave Symposium Digest, 2005 IEEE MTT-S International*, June 2005, pp. 12-17.

[2] K. Soumyanath, et al., "Accurate on-chip interconnect evaluation: a time-domain technique," *IEEE Journal of Solid-State Circuits*, vol. 34, no. 5, pp. 623-631, May 1999.

Organic Thin-film Transistor Integrated Circuits

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 an organic material as the semiconductor. Electronically, OTFTs have field-effect mobilities that are comparable to those of hydrogenated amorphous silicon TFTs [1]. Mechanically, organic materials can be processed at room temperature, thus permitting substrates that are low-cost, large-area, and mechanically flexible [2]. We investigate the OTFT technology through device characterization and circuit design.

The OTFTs used in our work are lithographically processed at temperatures less than 95 ºC to produce integrated circuits compatible with mechanically-flexible substrates [3]. Figure 1a shows the OTFT device cross-section, with typical output characteristics in Figure 1b. **Device parameters such as threshold voltage, charge mobility,** subthreshold slope, and contact resistance are characterized and studied.

We design OTFT circuits to achieve two goals: to aid the understanding of device physics and to evaluate the feasibilities of various OTFT applications. One such circuit is the complementary-to-absolute-temperature (CTAT) circuit shown in Figure 2. The CTAT circuit promotes the study of thermally activated device mechanisms. At the same time, we use the circuit to explore highly-linear and low-power temperature-sensing applications. Another OTFT circuit is the 4×4 active-matrix imager, as described in [4]. Additional OTFT circuits include digital logic gates, ring oscillators, display pixel drivers, comparators, and static random-access memory cells.

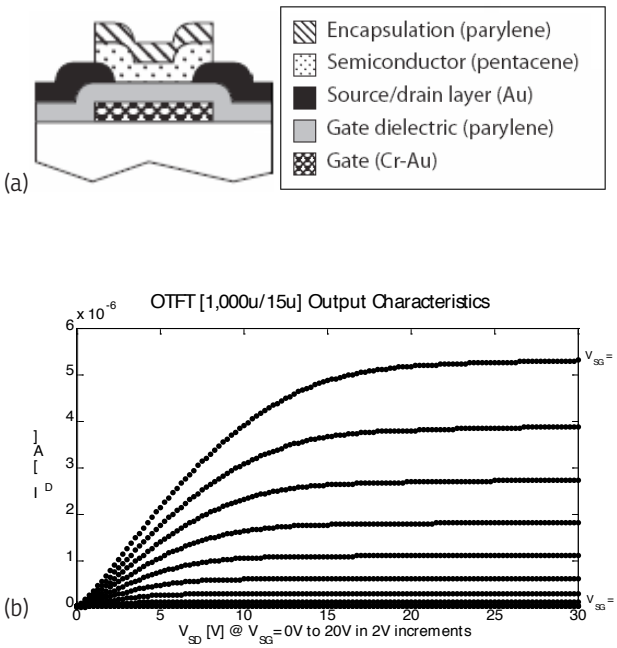

▲ Figure 1: (a) OTFT device cross-section and (b) typical OTFT output characteristics.

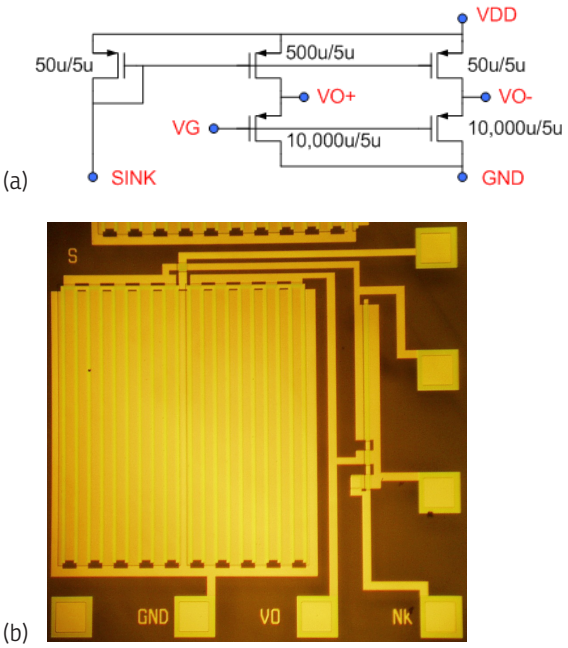

▲ Figure 2: (a) CTAT circuit schematic and (b) its die micrograph (area of 1mm^2).

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] M. Chason, P.W. Brazis, J. Zhang, K. Kalyanasundaram, and D.R. Gamota, "Printed Organic Semiconducting Devices," *Proc. of the IEEE*, July 2005, pp. 1348–1356.
[3] I. Kymissis, C.G. Sodini, A.I. Akinwande, and V. Bulovic, "An organic semiconductor based process for photodetecting applications," *IEEE International Electron Devices Meeting Technical Digest*, Dec. 2004, pp. 377–380.
[4] I. Nausieda, K. Ryu, I. Kymissis, A.I. Akinwande, V. Bulovic, and C.G. Sodini, "An organic imager for flexible large area electronics," *Digest of IEEE International Solid-State Circuits Conference*, San Francisco, Feb. 2007, pp. 72–73.

Design and Integration of Complementary CNT-FETs

K.-J. Lee, J. Kong, A.P. Chandrakasan
Sponsorship: SRC/FCRP IFC, Intel

The high mobility and nanometer-scale dimensions of carbon nanotubes (CNTs) make them attractive for many electronic applications. Experimental work has shown near-ballistic transport properties of CNT field-effect transistors (FET) [1] and demonstrated rudimentary circuit structures [2, 3]. While these efforts highlight the potential of CNT-FETs, improving device reliability and large-scale integration schemes remains a big challenge. This project investigates the fabrication process of complementary CNT-FETs. Similar to [3, 4], CNT-FETs are fabricated by first growing high-density, aligned CNTs via chemical vapor deposition. As-grown CNTs have a mixture of metallic and semiconducting CNTs, which significantly degrades the on-off current ratio. This mixture can be improved by electrically burning metallic CNTs. Then, device parameters are optimally tuned by high-K gate insulator deposition and metal-gate engineering [2]. This project examines device reliability, design trade-offs, and the integrity of the process flow.

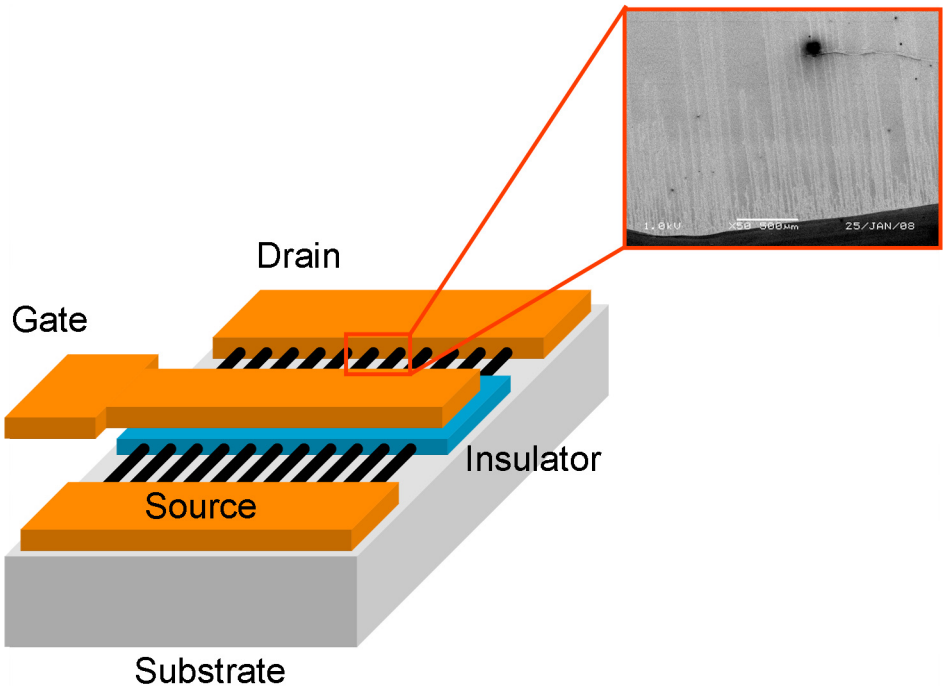

▲ Figure 1: Schematic of a CNT-FET with metal-gate on top of a high-K gate insulator. A single device consists of many CNTs. An SEM micrograph shows high-density, aligned CNTs grown from a thick catalyst region at the bottom of the image.

References

[1] A. Javey, J. Guo, M. Paulsson, Q. Wang, D. Mann, M. Lundstrom, and H. Dai, "High-field quasi-ballistic transport in short carbon nanotubes," *Physics Review Letter*, vol. 92, p. 106804, Mar. 2004.

[2] Z. Chen, J. Appenzeller, Y.-M. Lin, J. Sippel-Oakley, A.G. Rinzler, J. Tang, S.J. Wind, P.M. Solomon, and P. Avouris, "An integrated logic circuit assembled on a single carbon nanotube," *Science*, vol. 311, no. 5768, p. 1735, Mar. 2006.

[3] P.-W. Chiu and C.-H. Chen, "High-performance carbon nanotube network transistors for logic application," *Applied Physics Letters*, vol. 92, p. 063511, Feb. 2008.

[4] S.J. Kang, C. Kocabas, T. Ozel, M. Shim, N. Pimparkar, M.A. Alam, S.V. Rotkin, and J.A. Rogers, "High-performance electronics using dense, perfectly aligned arrays of single-walled carbon nanotubes," *Nature Nanotechnology*, vol. 2, pp. 230-236, Mar. 2007.

Optimization of Transistors for VHF Resonant Power Converters

A. Sagneri, D.J. Perreault

Sponsorship: National Semiconductor Corporation, NSF, CICS

At current switching frequencies (about 1-10 MHz), the required energy storage in a typical dc-dc converter yields passive component dimensions that are large with respect to integrated processes. A continuing effort to switch in the VHF regime (30-300 MHz) has relaxed this restriction, eliminating the need for magnetic materials and reducing component sizes to the point where integration or co-packaging is feasible [1-4]. A prototype boost converter operating at 110 MHz, with an 8-18V input and a 33V, 23W output achieved better than 87% efficiency using an off-the-shelf RF power LDMOSFET [4]. The need for lower cost and integrated switches prompted experiments with LDMOSFETs fabricated using an integrated power process. These prototypes operate at 50 MHz over similar voltage and power ranges, but with roughly 12% lower efficiency [2].

The present work shows that optimization of device layout, without changes to the underlying process or its design rules, leads to significant improvement in integrated LDMOSFET performance. Layout optimization is accomplished by creating a device model in Matlab that relates device parasitic elements and on-state resistance to layout geometry, as in Figure 1. A search of the design space bounded by layout rules yields the best performing device. It was possible, for example, to reduce gate resistance by a factor of three using this technique.

The ability to operate these devices above their typical breakdown voltage limits allows further performance gains. For LDMOSFETs, hot-carrier effects usually set the breakdown voltage [5]. In VHF power converters, soft-switching provides minimal voltage and current overlap, as Figure 2 shows, thereby avoiding the hot-carrier regime. Therefore, operation closer to the avalanche breakdown limit is possible. This allows for a shorter drain extension at a given breakdown voltage, conferring smaller parasitic capacitance and lower specific on resistance. Testing has validated the use of a device usually rated for 20 V at up to 35 V with appropriate topologies.

Prototype converters built using these optimized devices achieve efficiencies on par with those using commercial RF LDMOS devices, or roughly 88% drain efficiency for an overall converter efficiency of 85-86%.

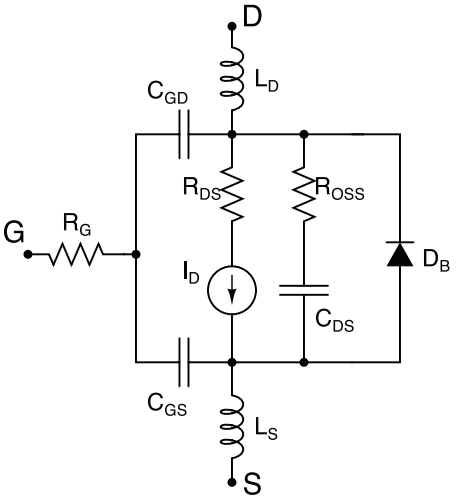

▲ Figure 1: MOSFET model detailing loss mechanisms important at VHF. In particular, losses in R_{OSS} and R_G represent displacement loss and gating loss, both frequency-dependent mechanisms that can be influenced favorably by optimizing the layout geometry.

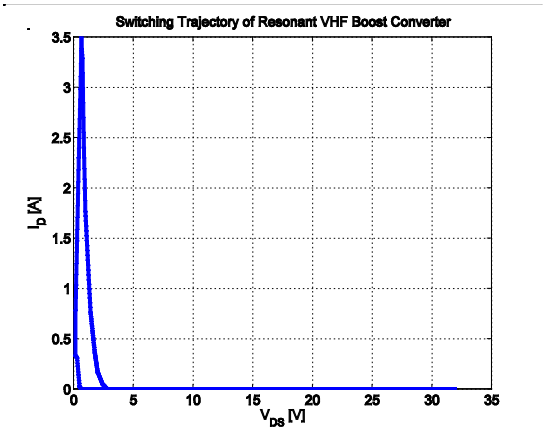

▲ Figure 2: Switching trajectory of a resonant dc-dc boost converter. Voltage and current are never simultaneously large. This allows hot-carrier effects to be largely neglected when determining maximum allowable switch operating voltage.

References

[1] J.M. Rivas, R.S. Wahby, J.S. Shafran, and D.J. Perreault, "New architectures for radio-frequency dc-dc power conversion," in *Proc. 35th Annual IEEE Power Electronics Specialist Conference*, Aachen, Germany, June 2004, pp. 4074-4084.

[2] A.D. Sagneri, "The design of a VHF dc-dc boost converter," Master's thesis, Massachusetts Institute of Technology, Cambridge, 2007.

[3] J.M. Rivas, "Radio frequency dc-dc power conversion," Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, 2006.

[4] R. Pilawa, A.D. Sagneri, J.M. Rivas, D. Anderson, and D.J. Perreault, "Very high-frequency resonant boost converters," in *Proc. 38th IEEE Power Electronics Specialist Conference*, Orlando, FL, June 2007.

[5] A. Hastings, *The Art of Analog Layout*. New Jersey: Prentice-Hall, Inc., 2001.

Two-point Wideband Linearization of RF Power Amplifiers for Memory Compensation

H. Boo, S. Chung, J.L. Dawson

Linearizing RF power amplifiers (PAs) has become a very challenging issue in RF design. The PAs are known for their nonlinear behavior and distortion emissions. High data rates and envelope variations in communication systems impose severe wideband linearity requirements on the PA. Also, at high power levels, memory effect, frequency-dependent distortion, comes into effect. Previous work on Cartesian Feedback (CFB) offers high linearization while also suppressing memory effect. However, the loop dynamics inherently limit its application to low bandwidth transmission. Digital predistortion (DPD) relies on baseband symbol mapping using a look-up table to predistort symbols; the mapping implements the inverse of the PA nonlinearity. Although the technique allows wide bandwidth transmission, it cannot compensate for memory effect since the look-up table does not model the constantly changing behavior of the PA.

We consider a two-point linearization system in which a CFB and DPD are combined to linearize the PA for high bandwidth transmission with memory-effect compensation. The CFB path takes care of the low bandwidth portion while suppressing the memory effect, and the DPD covers the high frequency range while suppressing the memoryless nonlinearity. The two-point architecture offers linearization without any complex PA modeling or the use of power-hungry DSP.

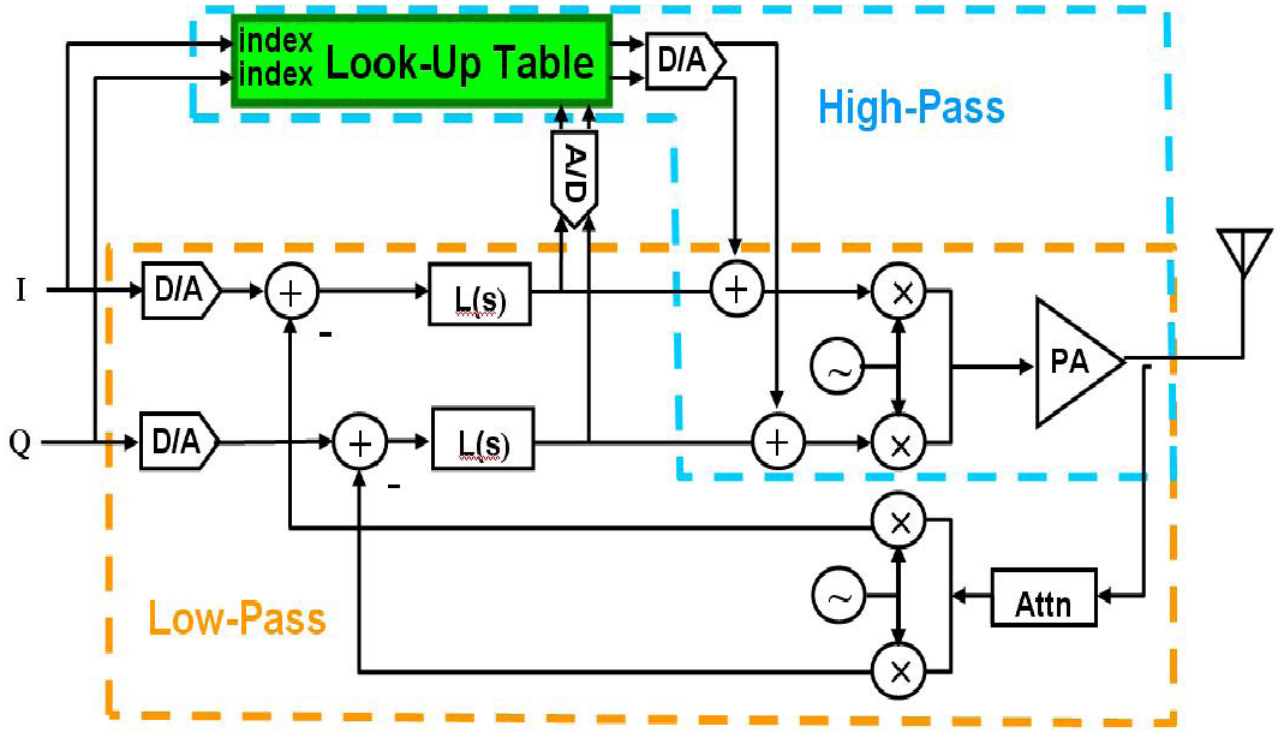

▲ Figure 1: Two-point architecture that combines CFB and DPD for PA linearization. The architecture enables high bandwidth transmission, in addition to suppressing memory effect.

References

[1] S. Chung, J.W. Holloway, and J.L. Dawson, "Open-loop digital predistortion using Cartesian feedback for adaptive RF power amplifier linearization," *IEEE MTT-S International Microwave Symposium*, Honolulu, HI, June 2007.

[2] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg, *Atom-photon Intereactions*. New York: John Wiley & Sons, Inc., 1992.

Energy-efficient Pulsed-UWB Transceiver for Insect Flight Control

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

Due to continued process improvement and miniaturization of electronics, complex highly integrated systems can now be realized on the micro-scale. Micro-scale electronics and micro-electromechanical systems (MEMS) are particularly relevant for biomedical and biological applications, in which devices are often implanted within or on organisms. An emerging application for micro-scale electronics is hybrid-insect flight control, where MEMS devices are placed on and within insects to alter flight direction. Such a hybrid insect system would take the best qualities of biology — energy storage, efficient flight control, highly adapted sensing — and combine them with the best qualities of MEMS — low weight, small size, deterministic control, and interfacing with computation. 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, 2]. Power, weight and volume are all highly constrained, necessitating a highly integrated solution with minimal off-chip components. To reduce power requirements and system complexity, communication is unidirectional from the base station to the moth. Data is transmitted by PPM modulation in one of three 500-MHz channels in the 3-5-GHz band. The highly digital base station transmitter consumes zero static bias current and achieves an energy efficiency of 113-to-19pJ/pulse at data rates from 100kbps-to-15.6Mbps [3]. The moth receiver consists of a packaged chip along with an off-chip crystal resonator to provide a stable clock. Figure 2 presents a block-diagram of the system-on-chip for the hybrid-insect wireless receiver. The non-coherent receiver amplifies, squares, and integrates received pulses to measure the amount of energy received in a given time period. A highly parallelized demodulator detects packets and rapidly synchronizes and receives payload data. The moth stimulator supplies a digital, pulse-width modulated signal to eight channels to control flight direction.

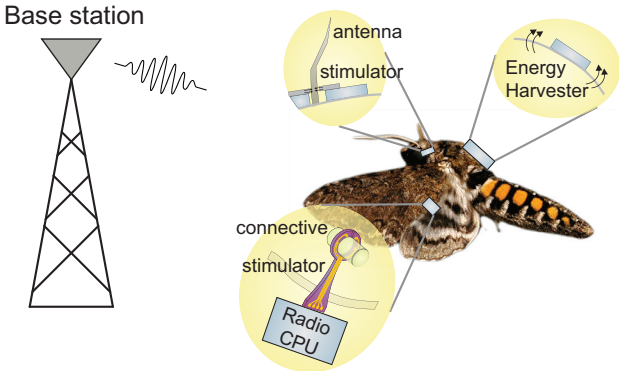

▲ Figure 1: Overview of hybrid-insect flight control system (Moth image courtesy of Armin Hinterworth).

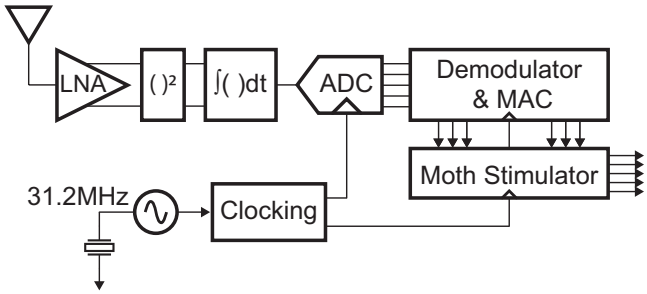

▲ Figure 2: Block diagram of the wireless receiver for the moth.

References

[1] F.S. Lee and A.P. Chandrakasan, "A 2.5nJ/b 0.65V 3-to-5GHz sub-banded UWB receiver in 90nm CMOS," in *Proc. IEEE International Solid-State Circuits Conference*, San Francisco, CA, Feb. 2007, pp. 116-117.

[2] D.D. Wentzloff and A.P. Chandrakasan, "A 47pJ/pulse 3.1-to-5GHz all-digital UWB transmitter in 90nm CMOS," in *Proc. IEEE International Solid-State Circuits Conference*, San Francisco, CA, Feb. 2007, pp. 118-119.

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

Wideband Polar Transmitter for Multi-standard Communications

S. Chung, P. Godoy, S. Rayanakorn, D.J. Perreault, J.L. Dawson
Sponsorship: Deshpande Center for Technological Innovation

This work focuses on implementing an RF transmitter that is suitable for multi-standard transmission without the need for an array of bulky, impossible-to-integrate filters. Polar transmitters have long been recognized as having a fundamental advantage in this regard [1]. However, many challenges must be overcome to put polar transmitters into use; these challenges divide into three main parts: (1) a wideband, highly efficient amplitude modulator, (2) a wideband phase modulator, and (3) the predistortion/time alignment circuitry.

To achieve high-efficiency amplitude modulation, a switching-mode DC/DC converter whose switching frequency is several times the signal bandwidth, is typically used. However, the converter efficiency is severely degraded for wideband signals due to increased switching losses [2]. To overcome this difficulty, we use only a few discrete amplitude values in the DC/DC converter. We can do this by leveraging the fact that most modulation schemes are digital and require only discrete signal amplitudes (e.g., for QAM-16, only three amplitude levels are possible.). A low-pass filter shapes the resulting signal spectrum to meet spectral mask requirements.

A two-point modulator performs the wideband phase modulation [3]. It is a phase-locked loop (PLL) with two input paths. The first input path is low-pass-filtered to the output by the closed-loop transfer function of the PLL, whereas the second input path is high-pass-filtered to the output. In theory, the bandwidth of the two-point modulator is unbounded. However, nonlinearity in the voltage-controlled oscillator introduces significant phase errors in the second data path. These phase errors are corrected with an adaptive digital predistortion circuit, similar to that applied to power amplifiers [4].

The wideband predistortion [4] is necessary to linearize a highly power-efficient but nonlinear switching power amplifier. The time-delay alignment compensates for the delay difference between the amplitude path and the phase path.

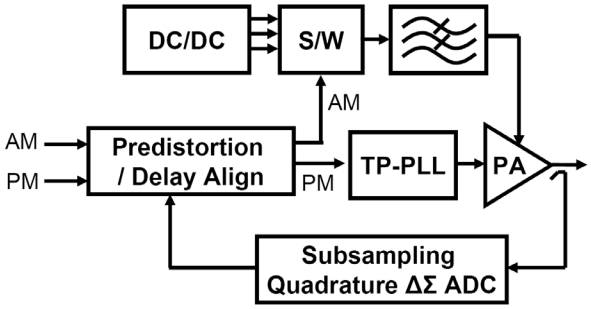

▲ Figure 1: Polar transmitter architecture.

References

[1] E. McCune, "Advanced architectures for high-efficiency, multi-mode, multi-band terminal power amplifiers," *IEEE Radio and Wireless Conference*, pp. 167-170, Sept. 2004.

[2] F. Wang, D.F. Kimball, D.Y. Lie, P.M. Asbeck, and L.E. Larson, "A monolithic high-efficiency 2.4-GHz 20-dBm SiGe BiCMOS envelope-tracking OFDM power amplifier," *IEEE Journal of Solid-State Circuits*, vol. 42, no. 6, pp. 1271-1281, June 2007.

[3] K.C. Peng, C.H. Huang, C.J. Li, and T.S. Horng, "High-performance frequency-hopping transmitters using two-point delta-sigma modulation," *Microwave Theory and Techniques, IEEE Transactions on*, vol. 52, no. 11, pp. 2529-2535, Nov. 2004.

[4] S. Chung, J.W. Holloway, and J.L. Dawson, "Open-loop digital predistortion using Cartesian feedback for adaptive RF power amplifier linearization," *IEEE MTT-S International Microwave Symposium Digest*, pp. 1449-1452, June 2007.

A Low-jitter Programmable Clock Multiplier Based on a Pulse Injection-locked Oscillator

B. Helal, M.H. Perrott

Sponsorship: NSF, CICS

Recently, there has been an increased interest in CMOS circuits that leverage injection-locked oscillators for applications such as frequency division and clock multiplication. Clock multiplication using subharmonic injection-locked oscillators [1] has the potential advantage of suppressing the phase noise of the oscillator at a significantly higher bandwidth than integer-N Phase-Locked Loops (PLLs), which are typically used for clock multiplication. However, the lack of continuous tuning of the injected oscillator necessitates a large injection power to ensure locking over sufficient bandwidth. Even when an Injection-Locked Phase-Locked Loop (ILPLL) provides continuous tuning [2], increased frequency spurs can result due to the mismatch between the injection and the PLL paths and the typical analog nature of the tuning path. A similar issue occurs with realigned PLLs as described in [3].

This research aims to develop a subharmonic injection-locked architecture that is continuously tuned with low frequency-spurs and significantly lower output jitter than a typical PLL. We propose a Pulse Injection-Locked Oscillator (PILO)that is suitable for continuous tuning using a highly-digital tuning technique that we recently introduced [4]. In addition, the operation of the PILO can be more intuitively understood than typical injection-locked techniques and a linearized context can be used to model its phase noise using [3].

Figure 1 shows the proposed PILO circuit in which the LC tank of the oscillator is frequency locked by periodically shorting the tank with a switch that is driven by a train of narrow pulses. The pulses are generated from a reference source, whose frequency is at a sub-multiple of the desired output frequency. To achieve low output jitter, the reference source must also have low jitter since its noise will be mostly passed on to the oscillator output [3].

Figure 2 shows the prototype of the proposed highly-digital PILO architecture that offers continuous frequency tuning of the injection-locked oscillator and eliminates sources of analog mismatches in the tuning path. The cornerstones of the tuning technique are a digital correlation technique and a scrambling time-to-digital converter (TDC). The architecture eliminates mismatches and offsets in analog detection and integration blocks by using a single detection path and replacing the charge pump and analog integrator with a digital accumulator [4]. A test chip was fabricated using a CMOS 0.13 µm process, and the prototype used a 50 MHz reference input to generate an output at 3.2 GHz (and up to 4 GHz). For the 3.2 GHz output, measured results demonstrated overall reference spurs of -63.4 dBc and estimated random and deterministic jitter of 134 fs (rms) and 211 fs (peak-to-peak), respectively [5].

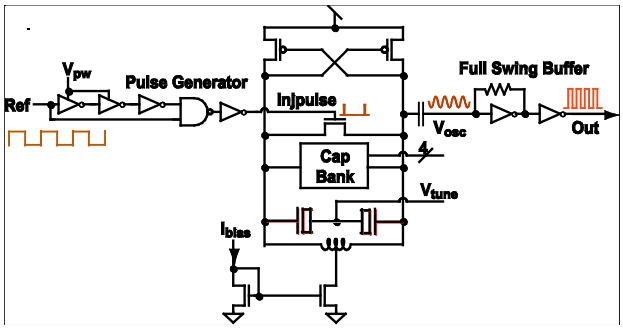

▲ Figure 1: Proposed PILO structure.

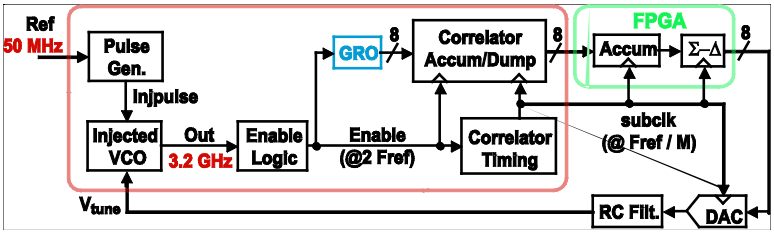

▲ Figure 2: PILO prototype and its highly-digital tuning circuit.

References

[1] H. Ahmed, C. DeVries, and R. Mason, "A digitally tuned 1.1 GHz subharmonic injection-locked VCO in 0.18 um CMOS," in *Proc. European Solid State Circuits Conference (ESSCIRC)*, Sept. 2003, pp. 81- 84.

[2] S. Kudszus, M. Neumann, T. Berceli, and W. Haydl, "Fully integrated 94-GHz subharmonic injection-locked PLL circuit," *IEEE Microwave Guided Wave Letters*, vol. 10, no. 2, pp. 70-72, Feb. 2000.

[3] S. Ye, L. Jansson, and I. Galton, "A multiple-crystal interface PLL with VCO realignment to reduce phase noise," *IEEE Journal of Solid-State Circuits*, vol. 37, no. 12, pp. 1795–1803, Dec. 2002.

[4] B. Helal, M. Straayer, G. Wei, and M. Perrott, "A highly-digital MDLL-based clock multiplier that leverages a self-scrambling time-to-digital converter to achieve sub-picosecond jitter performance," *IEEE Journal of Solid-State Circuits*, vol. 43, no. 4, pp. 855 – 863, Apr. 2008.

[5] B. Helal, C.-M. Hsu, K. Johnson, and M. Perrott, "A Low-noise Programmable Clock Multiplier based on a Pulse Injection-locked Oscillator with a Highly-digital Tuning Loop," *IEEE Radio Frequency Integrated Circuits (RFIC) Symposium Digest of Papers*, Jun. 2008, pp. 423-426.

A 3.6-GHz Low-noise, 500-kHz BW Digital-frequency Synthesizer with Digital Noise Cancellation

C.-M. Hsu, M.H. Perrott

Sponsorship: SRC/FCRP C2S2

A digital fractional-N frequency synthesizer is presented that leverages a noise-shaping time-to-digital converter (TDC) and a simple quantization noise cancellation technique to achieve low phase noise with a wide PLL bandwidth of 500 kHz. In contrast to previous cancellation techniques, the proposed structure requires no analog components and is straightforward to implement with standard cell digital logic. With the cancellation technique enabled, the synthesizer achieves phase noise of -132 dBc/Hz at 3 MHz offset, and an integrated phase noise from 1 kHz to 40 MHz of 204 fs rms at 3.67 GHz.

Figure 1 shows a block diagram of the proposed synthesizer. High-resolution digital phase detection is performed with an improved version of the gated ring oscillator (GRO) time-to-digital converter presented in [1]. Another interesting component of the architecture is an asynchronous frequency divider that avoids divide-value delay variation at its output. In addition, in contrast to previous digital PLL implementations [2], the digitally-controlled oscillator (DCO) is implemented as a conventional LC voltage-controlled oscillator (VCO) with coarse and fine varactors, which are controlled by two passive 10-bit, 50-MHz digital-to-analog converter (DAC) structures. An additional four-bit MIM capacitor bank is included in the VCO to improve its tuning range. To control both the coarse and fine varactors in the VCO, the loop filter consists of two paths.

The chip is implemented in a 0.13-μm CMOS process and has an active area of 0.95 mm². The prototype consumes 26mA from a 1.5V supply, excluding the VCO output buffer that consumes 7mA from a 1.1V supply. Figure 2 shows the measured phase noise at 3.67GHz from an Agilent Signal Source Analyzer E5052A, and the results are shown with and without cancellation of the quantization noise. As the figure reveals, greater than 15 dB noise cancellation is achieved such that the VCO dominates out-of-band noise. With noise cancellation enabled, the in-band noise is -108dBc/Hz at 400 kHz offset, and out-of-band noise is -132dBc and -150dBc at 3 MHz and 20 MHz offsets, respectively. The reference spur was measured with an Agilent Spectrum Analyzer 8595E to be -65dBc. Fractional spurs were first measured at carrier frequencies spanning from 3.620 GHz to 3.670 GHz in increments of 1MHz. With this setup, worst case spurs were measured to be -53dBc at carrier frequencies of 3.649 and 3.651GHz and -64dBc at carrier frequencies of 3.648 and 3.652GHz;, the spurs were less than -65 dBc at all the other carrier frequencies. When VCO frequency was set closer to 3.65GHz, the worst case spur was observed to be -42dBc when carrier frequency was at 3.6504GHz.

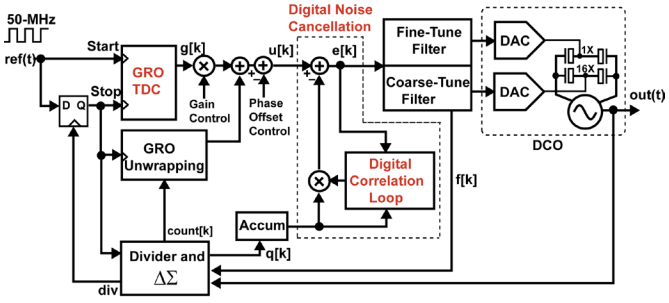

▲ Figure 1: Block diagram of proposed synthesizer.

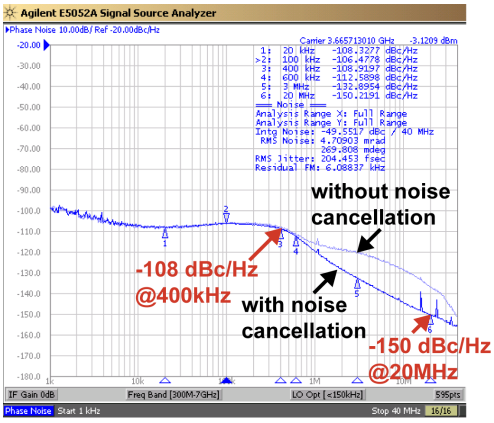

▲ Figure 2: Measured phase noise performance.

References

[1] B. Helal, M. Straayer, G. Wei, and M. Perrott, "A Low-jitter 1.6-GHz Multiplying DLL Utilizing a Scrambling Time-to-digital Converter and Digital Correlation," *VLSI Symposium Digest of Technical Papers*, pp. 166-167, June 2007.

[2] R.B. Staszewski et al., "All-digital PLL and GSM/EDGE Transmitter in 90-nm CMOS," *ISSCC Digest of Technical Papers*, pp. 316-317, Feb. 2005.

An All-digital UWB Transmitter with Dual Capacitively-coupled Power Amplifiers

P.P. Mercier, D.C. Daly, A.P. Chandrakasan
Sponsorship: DARPA, NSERC, STMicroelectronics

Applications like sensor networks, medical monitoring, and asset tracking have led to a demand for energy-efficient and low-cost wireless transceivers. These types of applications typically require low effective data rates, thus providing an opportunity to employ simple modulation schemes and aggressive duty-cycling. Due to their inherently duty-cycled nature, pulsed-UWB systems have been shown to be amenable to low-power operation [1, 2]. Furthermore, the use of non-coherent signaling greatly simplifies both transmitter and receiver implementations, offering substantial energy savings [3].

This work presents an all-digital transmitter designed for a non-coherent pulsed-UWB system [4]. By exploiting the fact that center frequency tolerances are relaxed in wideband non-coherent communication, the transmitter can synthesize UWB pulses from an energy-efficient, single-ended digital ring oscillator. To generate phase modulated pulses (which are required for spectral scrambling purposes), the oscillator output is fed to two banks of parallel tri-state inverters, shown in Figure 1. Maintaining opposite common modes at the output of these inverters during idle mode (i.e., when no pulses are being transmitted) eliminates low-frequency turn-on and turn-off transients typically associated with single-ended digital circuits driving single-ended antennas. Thus, no area-expensive balun is required to generate BPSK-modulated pulses. The parallel inverter banks permit digital pulse-shaping, resulting in on-chip FCC-compliant operation, as Figure 2 shows. The transmitter was fabricated in 90-nm CMOS, consumes zero static bias current, and achieves an energy efficiency of 113-to-19pJ/pulse at data rates from 100kbps-to-15.6Mbps.

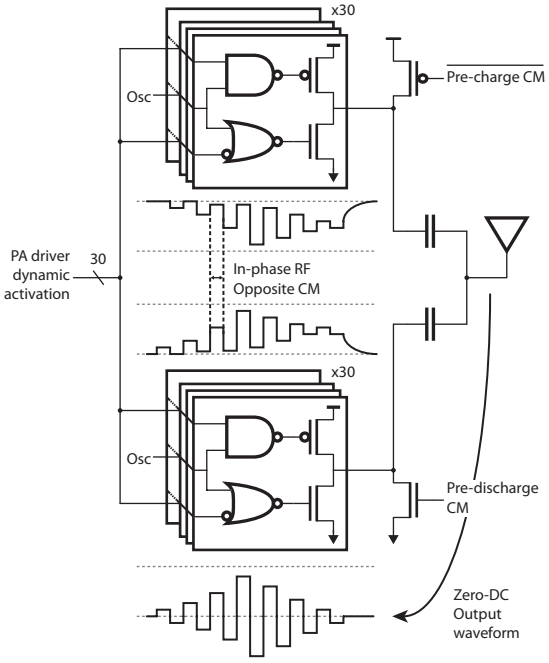

▲ Figure 1: Dual-digital power amplifiers create a bi-polar (zero-DC) output pulse by combining paths that are in-phase at RF yet have counter-phase common-mode components that are cancelled.

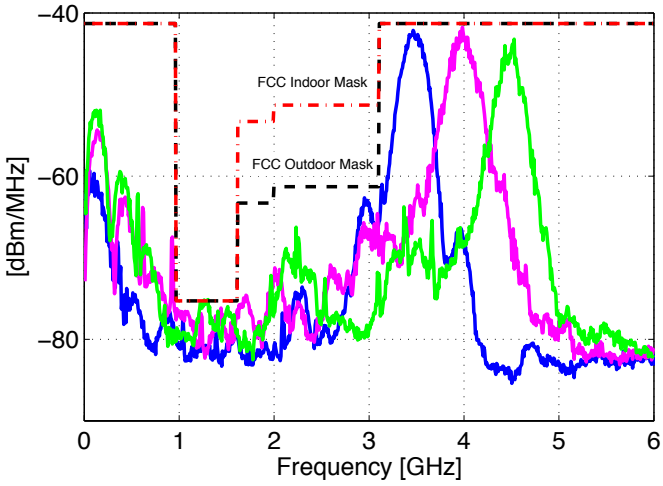

▲ Figure 2: Output spectral densities in three channels, from 3.5-to-4.5 GHz, illustrating on-chip FCC compliance.

References

[1] F.S. Lee and A.P. Chandrakasan, "A 2.5nJ/b 0.65V 3-to-5GHz Sub-banded UWB Receiver in 90-nm CMOS," in *Proc. IEEE International Solid-State Circuits Conference*, Feb. 2007, pp. 116-117.

[2] D.D. Wentzloff and A.P. Chandrakasan, "A 47pJ/pulse 3.1-to-5GHz All-digital UWB Transmitter in 90-nm CMOS," in *Proc. IEEE International Solid-State Circuits Conference*, Feb. 2007, pp. 118-119.

[3] L. Stoica, A. Rabbachin, H. Repo, S. Tiuraniemi, and I. Oppermann, "An ultra-wideband system architecture for tag-based wireless sensor networks," *IEEE Transactions on Vehicular Technology*, vol. 54, pp. 1632-1645, Sept. 2005.

[4] P.P. Mercier, D.C. Daly, and A.P. Chandrakasan, "A 19pJ/pulse UWB Transmitter with Dual Capacitively-coupled Digital Power Amplifiers," in *Proc. IEEE Radio Frequency Integrated Circuits Symposium*, Jun. 2008.

Phase-locked Loop Design for Millimeter-wave Imaging

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

Millimeter-wave imaging has potential applications such as collision-avoidance radar at 77 GHz and concealed weapons detection at 77 GHz, 94 GHz, and higher. This research investigates the challenges of designing a phase-locked loop (PLL) that could be used in a millimeter-wave (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 77 GHz [1]. A central processor will perform digital beamforming on the aggregated data from the array to achieve an expected frame rate of 10 fps. The 77-GHz input signal will be downconverted by a mixer with a 76-GHz local oscillator, generated by the PLL, to obtain an intermediate frequency (IF) of 1-GHz. This signal is digitized by an analog-to-digital converter that is operating at 4.75 GHz and is sent to the CPU.

Since accurate beamforming requires tight control of the phase over the array of elements to prevent blurring between pixels, the PLL will be designed for minimal phase noise and power dissipation. The PLL will provide a pure 76-GHz tone for the imaging system. Figure 1 shows a block diagram of the PLL that was designed in 130-nm silicon-germanium (SiGe) BiCMOS process. A low phase noise 150-MHz crystal oscillator will be used as the reference. The PLL bandwidth was chosen to be 5 MHz to balance the phase noise caused by the voltage controlled oscillator (VCO) and the charge pump. The 76-GHz VCO is based on the cross-coupled design used in the passive imager [2]. The divider chain consists of nine divide-by-2 static frequency dividers. The first six are created in emitter coupled logic and the last three are designed in CMOS. The highest frequency divider utilizes inductive peaking for increased operating frequency. In simulation, the core PLL consumes 130mW and the buffers consume 45mW from a 2.5V supply. Figure 2 shows the layout of the test chip containing the full PLL that was submitted for fabrication in December 2007.

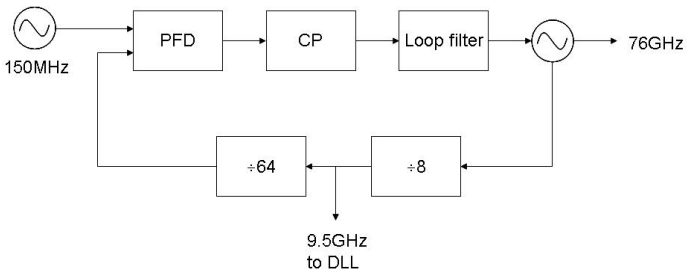

▲ Figure 1: Block diagram of the 76-GHz PLL.

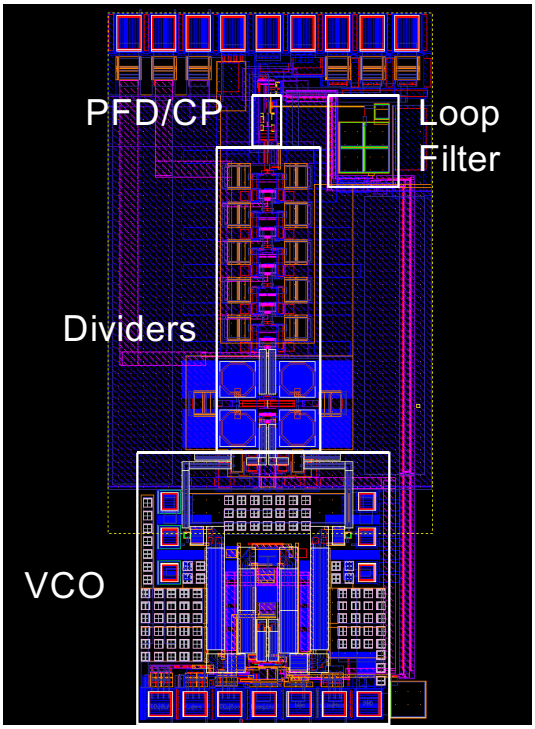

▲ Figure 2: Layout of 76-GHz PLL that was submitted in December 2007 in a 130-nm SiGe BiCMOS process. The die dimension is 2mm x 1mm.

References

[1] A. Accardi, "Digital Compensation in Imaging Systems," Dissertation proposal, Massachusetts Institute of Technology, Cambridge, 2007.
[2] J. Powell and C.G. Sodini, "A 77-GHz Receiver Front-end for Passive Imaging," presented at the *Radio Frequency Integrated Circuits Symposium*, June 2007.

Digital Phase Tightening for Improved Spatial Resolution in Mm-wave Imaging Systems

K. Lu, C.G. Sodini

Sponsorship: Lincoln Laboratory

Due to advances in silicon and digital processing technology, millimeter-wave (MMW) imaging solutions with high antenna array density are now viable at a low cost [1]. Millimeter resolution is desirable for many applications, such as automotive collision avoidance and concealed weapons detection.

An MMW imaging system utilizes an antenna array to capture the image. Each array node records magnitude and phase information, which then can be combined via a beamforming process to estimate the radiation incident from a particular direction, according to a beam pattern [2]. The beam can be steered electronically to acquire different regions of the target image. Unconstrained by the physics of a lens, digital image formation provides a greater level of computational flexibility. Unfortunately, this places a very stringent specification on the amount of timing jitter in the system. Using digital phase tightening can alleviate this problem.

Figure 1 shows the system used for digital phase tightening of the antenna array. The system extracts the phase and magnitude information from the input signal, which is approximately a 1 GHz sinusoid. After many cycles, this system locks onto the maxima, the minima and the zero crossings of the input. The outputs from the logic carry the phase information, while the outputs from the ADC carry the magnitude information. These outputs, gathered from each antennae array node, are used to calculate the image.

We are designing a system that is easy to manufacture and scale, and for which the control of phase is totally in the digital domain. Figure 2 shows a screenshot of the layout of a prototype system. The design is currently being fabricated in a 90-nm CMOS process.

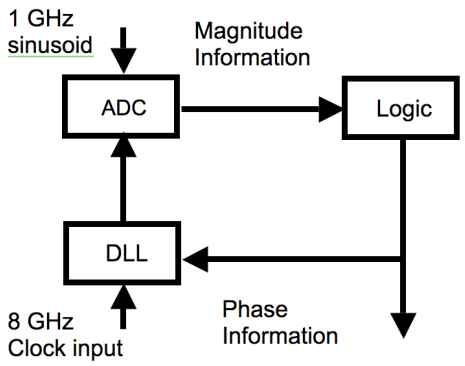

▲ Figure 1: Block diagram of overall system.

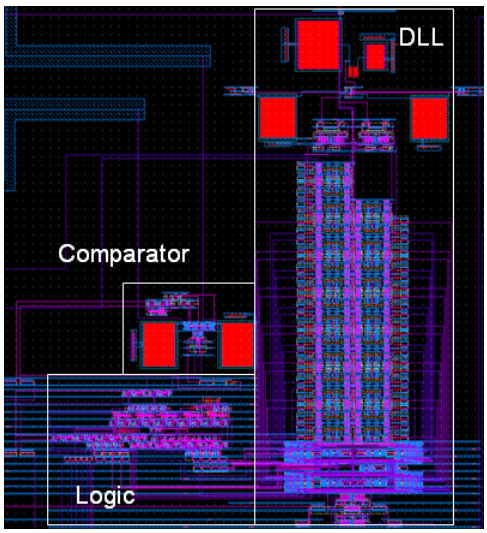

▲ Figure 2: Screenshot of the chip in layout with key blocks highlighted.

References

[1] W. Simburger et al., "Silicon-based RF ICs up to 100 GHz: Research trends and applications," in *Proc. Solid-State and Integrated Circuits Technology*, 2004, pp. 1:12-19.

[2] A. Accardi, "Digital Compensation in Imaging Systems," Ph.D. thesis proposal, Massachusetts Institute of Technology, Cambridge, 2007.

77- and 94-GHz Front-end Receivers with Flip-chip Antenna for MM-wave Imaging

J. Powell, H. Kim, C.G. Sodini
Sponsorship: SRC/FCRP C2S2, Lincoln Laboratory

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 at 77-GHz and higher. Significant advances have been made in these areas using SiGe technology [14]. This research focuses on passive imaging front-end receivers that have been developed and tested for the application of concealed weapons detection.

In this research, 77- and 94-GHz front-end receivers with antennas have been designed. Compared with current MMW research, these systems are wideband receivers that are fully differential. The RF front-ends are composed of low-noise amplifiers (LNA) and double-balanced mixers that downconvert the RF frequency to the intermediate frequency (IF) range of 1-9 GHz. The 77-GHz front-end also incorporates an on-chip cross-coupled voltage controlled oscillator (VCO). The antenna design is a Vivaldi-type antenna, designed for wideband performance. The antenna will be connected to the RF front-end via a flip-chip method. Gold solder bumps will be used to fuse the antenna to the receiver chips. Figure 1 shows a conceptual design of the MMW antenna.

Figure 2 shows the integrated 77-GHz RF front-end receiver implemented in 0.13 um SiGe. It achieves a particularly impressive maximum conversion gain of 46 dB. The full conversion gain ranges from 38-46 dB, and NF ranges from 6.5-10 dB within the IF frequency band of 1-9 GHz. The P1dB is approximately -38 dBm at 76 GHz RF.

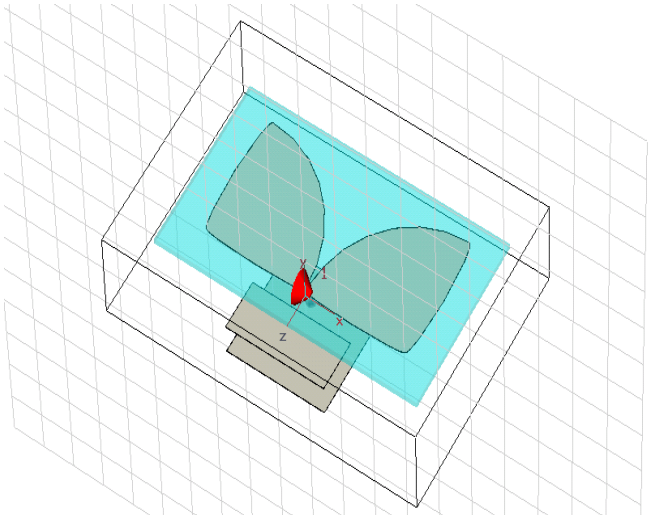

▲ Figure 1: Vivaldi-type MMW antenna, designed in CST Microwave Studio.

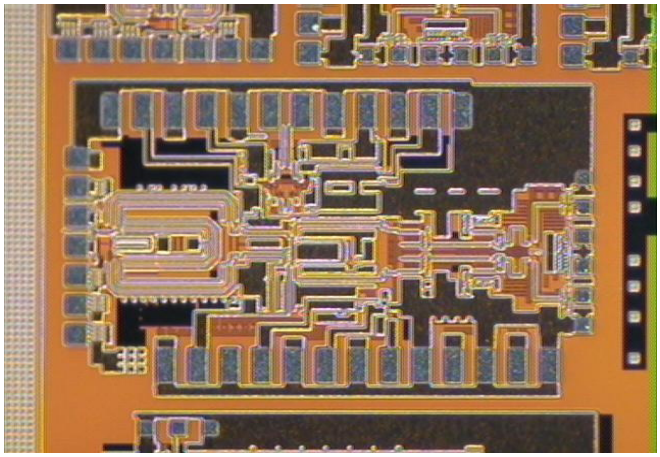

▲ Figure 2: Photo of 77-GHz Front-end receiver die.

References

[1] J. Powell, H. Kim, and C. Sodini, "A 77-GHz front-end receiver for passive imaging," *Proc. RFIC*, June 2007, Honolulu, HI, pp. 145-148.
[2] B.A. Floyd, S.K. Reynolds, U. Pfeiffer, and T. Beukema, "SiGe bipolar transceiver circuits operating at 60 GHz," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 1, pp. 156-167, Jan. 2005.
[3] S. Reynolds, et al., "Progress toward a low-cost millimeter-wave silicon radio" (Invited), *Proc. IEEE CICC*, Sept. 2005, pp. 563-570.
[4] A. Natarajan, et al., "A silicon-based 77-GHz transmitter with local LO-path phase shifting," *Proc. ISSCC*, Feb. 2006, pp. 182-183.

Noise-shaping Gated-Ring Oscillator Time-to-digital Conversion

M. Straayer, M.H. Perrott
Sponsorship: Lincoln Laboratory

High-performance oversampling and feedback applications, such as phase-locked loops, delay-locked loops, and clock-and-data recovery loops, utilize time-to-digital converters (TDC) as fundamental building blocks that bridge the gap between analog and digital signal domains. In this work, we fully utilize the inherent low-pass filtering action of these systems by providing first-order noise-shaping within a gated-ring oscillator time-to-digital converter.

As Figure 1 shows, the gated-ring oscillator TDC is enabled during the measurement interval, and the transitions that occur for that measurement are counted and summed. At the end of each measurement interval, the oscillator is disabled, and the state of the oscillator is held until the next input sample. The preservation of oscillator state transfers the quantization error from one sample to the next, which results in a first-order difference equation on the quantization error, or equivalently first-order noise-shaping in the frequency domain [1].

A prototype 11-bit GRO-TDC was designed and fabricated in 0.13-um CMOS [2]. The 1x1mm chip with 20 pads has an active area of only 157x258um (0.04mm2). A 1.5-V supply used for measurements results in a raw TDC resolution equal to 6ps, and operation was verified from 1.0-1.6V. Likewise, the sample rate is set to 50-Msps and verified over 100-Msps. Current is a linear function of the average measurement time, and it ranges from 1.5-14mA (2.2-21mW). Figure 2 shows a measurement of the TDC output with a 1.2ps(ppk) input signal at 26kHz. The rms timing jitter (measurement error) in a 1MHz bandwidth is below 100fs.

The authors wish to acknowledge MIT Lincoln Laboratory for research support through the Lincoln Scholars Program.

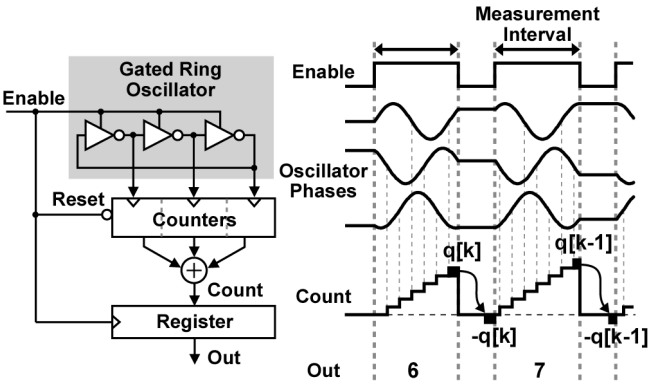

▲ Figure 1: Gated ring oscillator concept.

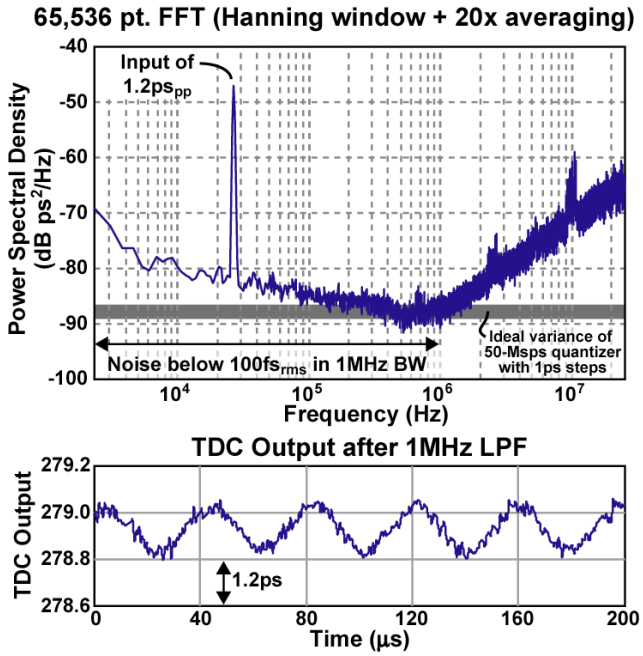

▲ Figure 2: Measured output for 1.2ps(ppk) 26-kHz input.

References

[1] B. Helal, M. Straayer, and M.H. Perrott, "A Low-jitter 1.6-GHz multiplying DLL utilizing a scrambling time-to-digital converter and digital correlation," *IEEE VLSI Symposium Digest of Technical Papers*, 2007.

[2] C.M. Hsu, M. Straayer, and M.H. Perrott, "A low-noise, wide bandwidth, 3.6-GHz digital SD fractional-N frequency synthesizer with quantization noise cancellation," *IEEE ISSCC Digest of Technical Papers, 2008.*

An Ultra-Low Power CMOS RF Transceiver for Medical Implants

J. Bohorquez, J.L. Dawson, A.P. Chandrakasan
Sponsorship: SRC/FCRP C2S2

Until recently, few medical implantable devices existed and fewer still provided the capability for wireless transmission of information. Most devices capable of data transmission did so through inductive coupling, which requires physical contact with the base-station and allows for only low data rates. In 1999, the FCC created the Medical Implant Communications Service (MICS) band in the range of 402–405 MHz specifically for medical telemetry [2]. The MICS band plan allows for RF communication between a medical implant and a base-station that is up to two meters away. This research seeks to design a transceiver specifically optimized for low-power, short-distance data transmission in a temperature-regulated environment, i.e., the human body. We do this by pushing as much complexity as possible out of the implant and into the base-station, taking advantage of the attributes of the environment, such as temperature control and slow transients, and incorporating the antenna into the oscillator for reduced power and improved performance. By optimizing the transceiver for reduced volume and power, we hope to extend the battery lifetime and functionality of medical implants for greater comfort and benefits to patients.

We propose a simple, almost all-digital transceiver (Figure 1) comprising a direct modulation frequency-shift keying transmitter and a super-regenerative receiver. The transmitter is composed of a digitally-controlled oscillator (DCO) and simple digital logic; the receiver uses the same DCO, a quench oscillator, an envelope detector, a comparator and simple digital blocks. Data is transmitted by directly modulating the frequency of the DCO through a capacitor bank. The frequency deviation constant is digitally set through a serial-to-parallel interface such that digital data composed of ones and zeros shift the DCO frequency by a desired amount. We linearize the digital-to-frequency relationship of the DCO by pre-distorting the capacitor banks used to tune the frequency and the frequency deviation constant of the transmitter. Instead of driving the antenna with a matched power amplifier, we exploit the low radiation power requirement of MICS to incorporate a loop antenna into the DCO. The inherently high Q of the antenna leads to improved noise performance for a given amount of power.

The SRR receives on-off keying (OOK) data and determines whether a one or a zero was sent by measuring the amount of time required for the envelope of the DCO output to reach a threshold. Input signals with large amplitudes and strong frequency content near the DCO's resonant frequency result in faster startup times. A digital counter determines the startup time and compares it to a threshold. Figure 2 shows normalized time domain signals in the receiver chain.

Our design was implemented in IBM 90-nm CMOS. The transmitter consumes under 350μW and meets MICS mask specifications with data rates up to 120kbps. The receiver consumes under 400μW and achieves a sensitivity of -99dBm or better at data rates of 40kbps, or -93dBm at data rates of 120kbps. The DCO tunes 24MHz in frequency steps smaller than 2kHz.

The authors acknowledge the support of the Focus Center for Circuit & System Solutions (C2S2), one of five research centers funded under the Focus Center Research Program, a Semiconductor Research Corporation Program.

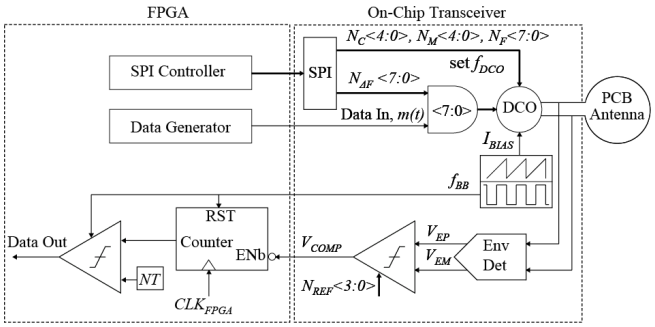

▲ Figure 1: Direct FSK transmitter and OOK super-regenerative receiver with on-chip quench oscillator.

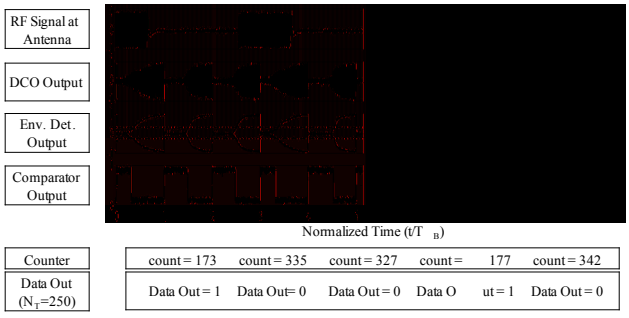

▲ Figure 2: Normalized time domain signals in the receiver chain.

Reference

[1] FCC Rules and Regulations, "MICS Band Plan," *Part 95*, Jan. 2003.

Impedance Measurement Utilizing an Integrated High-frequency Lock-in Amplifier for the Assessment of Neuromuscular Disease

O.T. Ogunnika, R.C. Cooper, J.L. Dawson
Sponsorship: CIMIT

Electrical Impedance Myography (EIM) is a non-invasive, painless clinical technique for the diagnosis and monitoring of a variety of neuromuscular diseases including generalized disorders, such as amyotrophic lateral sclerosis and muscular dystrophies, and localized conditions, such as focal nerve injuries [1]. It involves the application of a low-intensity alternating current to a muscle group and the measurement of the consequent surface voltage patterns [2]. This data is then used to compute the impedance of the muscle group under investigation.

The EIM measurement system (Figure 1) will be designed to take measurements from 10kHz to 10MHz. With the exception of the analog-to-digital converter (ADC) and direct digital frequency synthesizer (DDFS), all circuit blocks will be integrated onto a single IC. The system will operate in two modes selected by the multiplexer: high precision (using lock-in amplifier) and high speed (using bandpass filter). Note that in both modes of operation, this measured signal will be digitized by the ADC for further processing on a portable computer. The signal to be used in the interrogation of the muscle tissue is first created by this portable computer and then downloaded to the DDFS. The signal conditioner converts the signal generated by the DDFS from single-ended to differential and amplifies it to a level suitable for use with the crosspoint switches and reconfigurable electrode array in contact with the patient's skin. An instrumentation amplifier senses and amplifies the voltage that is induced across the muscle. Its output is bandpass-filtered before processing by the lock-in amplifier.

Lock-in amplifiers are capable of extracting a signal from noise in situations where linear filtering is not sufficient [3]. It is essentially a phase-sensitive detector that measures the amplitude and phase of a noisy signal and then reduces the effective noise bandwidth by means of synchronous detection, which requires knowledge of the desired signal's frequency. The architecture to be used in this work is the dual-phase lock-in amplifier shown in Figure 2. The input signal is multiplied by a reference signal at the same frequency, producing an output with a DC component proportional to the product of their amplitudes. The AC components in the output of the lock-in amplifier are filtered out by a lowpass filter. Good noise isolation requires the output lowpass filter to have a very small bandwidth (<1Hz). Typically, the implementation of filters with very small bandwidths require the use of passive components that produce large time constants leading to undesirably long signal settling times. Thus, the design of the output filter requires a tradeoff between measurement time and bandwidth. The division of the total signal amplification between the AC coupled input amplifiers and the DC coupled output amplifiers will be optimized. Too much amplification by the AC amplifiers could saturate the mixer while too much amplification by the DC amplifiers will lead to increased signal drift with time and temperature. Auto-scaling of the input signal drive will also be implemented to ensure that the EIM system can adapt to variations in load impedance without the need for time-consuming calibration.

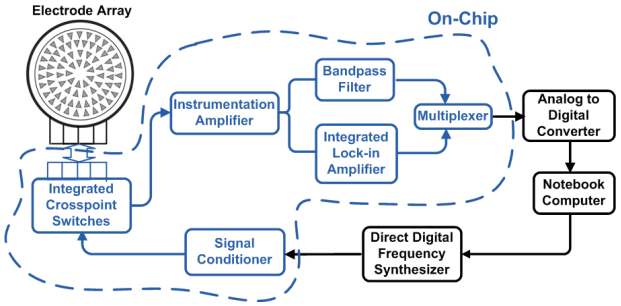

▲ Figure 1: System diagram of multi-frequency EIM measurement system.

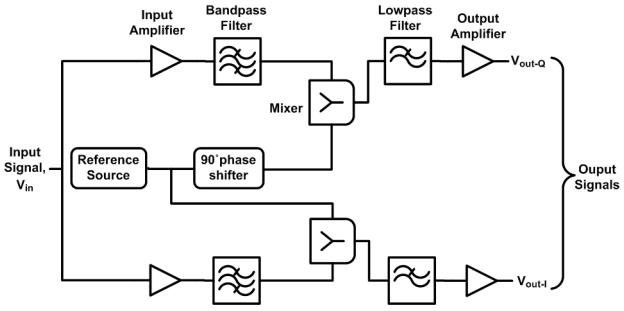

▲ Figure 2: Concept diagram of dual-phase lock-in amplifier.

References

[1] S.B. Rutkove, R. Aaron, and C.A. Shiffman, "Localized bio-impedance analysis in the evaluation of neuromuscular disease," *Muscle & Nerve*, pp. 390-397, Mar. 2002.

[2] G.J. Esper, C.A. Shiffman, R. Aaron, K.S. Lee, and S.B. Rutkove, "Assessing neuromuscular disease with multifrequency electrical impedance myography," *Muscle & Nerve*, vol. 34, pp. 595-602, July 2006.

[3] G. Ferri, P. De Laurentiis, A. D'Amico, and C. Di Natale, "A low-voltage integrated CMOS analog lock-in amplifier prototype for LAPS applications," *Sensors and Actuators*, A 92, pp. 263-272, 2001.

A Low-power, Blocking-capacitor-free, Charge-balanced Electrode-stimulator Chip with Less than 6-nA DC Error for 1-mA Full-scale Stimulation

J. Sit, R. Sarpeshkar

Large DC-blocking-capacitors are a bottleneck in reducing the size and cost of neural implants. We describe an electrode-stimulator chip that removes the need for large DC-blocking capacitors in neural implants by achieving precise charge-balanced stimulation with <6 nA of DC error. For cochlear implant patients, this is well below the industry's safety limit of 25 nA. Charge balance is achieved by dynamic current balancing to reduce the mismatch between the positive and negative phases of current to 0.4%, followed by a shorting phase of at least 1 ms between current pulses to further reduce the charge error. On +6- and -9-V rails in a 0.7–μm AMI high voltage process, the power consumption of a single channel of this chip is 47 μW when biasing power is shared by 16 channels.

An Energy-efficient Micropower Neural Recording Amplifier

W. Wattanapanitch, M. Fee, R. Sarpeshkar

This paper describes an ultra-low-power neural recording amplifier. The amplifier appears to be the lowest power and most energy-efficient neural recording amplifier reported to date. We describe low-noise design techniques that help the neural amplifier achieve input-referred noise that is near the theoretical limit of any amplifier using a differential pair as an input stage. Since neural amplifiers must include differential input pairs in practice to allow robust rejection of common-mode and power supply noise, our design appears to be near the optimum allowed by theory. The bandwidth of the amplifier can be adjusted for recording either neural spikes or local field potentials (LFPs). When configured for recording neural spikes, the amplifier yielded a midband gain of 40.8-dB- and a -3-dB-bandwidth from 45 Hz to 5.32 kHz; the amplifier's input-referred noise was measured to be 3.06 μV_{rms} while consuming 7.56 µW of power from a 2.8-V supply, corresponding to a noise efficiency factor (NEF) of 2.67 with the theoretical limit being 2.02. When configured for recording LFPs, the amplifier achieved a midband gain of 40.9 dB and a -3-dB bandwidth from 392 mHz to 295 Hz; the input-referred noise was 1.66 μV_{rms} while consuming 2.08 µW from a 2.8-V supply corresponding to an NEF of 3.21. The amplifier was fabricated in AMI's 0.5-µm CMOS process and occupies 0.16 mm^2 of chip area. We obtained successful recordings of action potentials from the robust nucleus of the arcopallium (RA) of an anesthesized zebra finch brain with the amplifier. Our experimental measurements of the amplifier's performance including its noise were in good accord with theory and circuit simulations.

A 12-b, 100-MS/s, Fully Differential Zero-crossing-based ADC

L. Brooks, H.-S. Lee

Sponsorship: CICS, NDSEG, DARPA

In an effort to improve on the resolution and production-worthiness of zero-crossing-based circuits (ZCBC) [1], a 12-b, 100-MS/s, fully differential pipelined ADC has been designed with the goals of improved power supply noise rejection, automatic offset compensation, and output signal range enhancement.

Figure 1 shows the implementation of two adjacent pipeline stages. The differential input (vip,vim) is sampled onto capacitors C1 and C2 during the sampling phase, when p1 is high. Then during the transfer phase, when p2 is high, current sources I2, I3, and I4 turn on to create a differential ramp on the output nodes (vop, vom). When the zero-crossing detector (ZCD) detects that the virtual ground condition has been realized, the sampling switch M3 turns off to lock the desired charge onto capacitors C3 and C4, which are the sampling capacitors of the next stage.

Figure 2 shows the schematic of the ZCD implementation. A pre-amplifier that performs a differential to single-ended conversion precedes a dynamic threshold-detecting latch that is similar to the dynamic zero-crossing detector introduced in [1]. The pre-amplifier consists of a differential pair (M1, M2) , a current bias source (M5), and a current mirror (M3, M4). For increased power savings, device M6 shuts off the bias current after the ZCD switches. Devices Ma and Mb function as enable switches to vary the gain of the current mirror to provide a digital method of adjusting the offset of the ZCD. For automatic offset compensation, a chopping technique similar to [2, 3] has been developed. The offset is estimated from the chopped output signal prior to demodulation. This offset estimate is then fed back and nulled in the ZCD to reduce the flicker noise and improve the signal range.

A fully differential implementation has been used to provide better power supply rejection. Continuous time common mode feedback is not necessary for several reasons. One is that unlike an op-amp based implementation, this implementation has very little common mode gain from the input to the output of each stage. Another is that using switch M3 as configured in Figure 1 as the only sampling switch ensures that despite mismatch between the current sources, the positive and negative sampling capacitors receive the same amount of charge and thus receive a common mode reset.

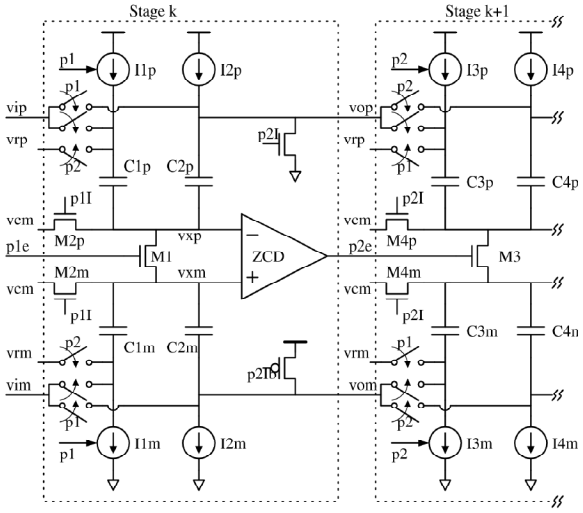

▲ Figure 1: Schematic of adjacent fully differential zero-crossing-based pipelined stages.

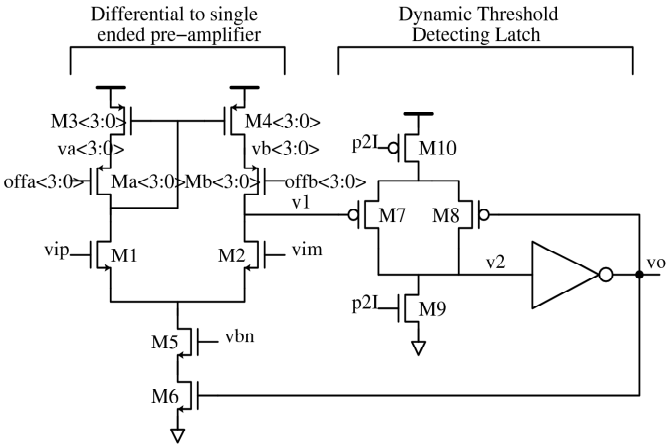

▲ Figure 2: Schematic of zero-crossing detector with digitally programmable offset.

References

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

[2] H. van der Ploeg, G. Hoogzaad, H.A. Termeer, M. Vertregt, and R.L. Roovers, "A 2.5-v 12-b 54-msample/s 0.25-µm CMOS ADC in 1-mm2 with mixed-signal chopping and calibration," *IEEE Journal of Solid-State Circuits*, vol. 36, pp. 1859–1867, Dec. 2001.

[3] S.M. Jamal, D. Fu, N. Chang, P. Hurst, and S.H. Lewis, "A 10-b 120-msample/s time-interleaved analog-to-digital converter with digital background calibration," *IEEE Journal of Solid-State Circuits*, vol. 37, no. 12, pp. 1618–1627, Dec. 2002.

Ultra-high Speed A/D Converters Using ZCBC Topology

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. In particular, this work will investigate ADC's with sampling rate up to 10 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. Therefore channel matching is an important design consideration for time-interleaved ADCs.

Although digital calibration can mitigate many of these non-idealities, timing mismatches constitute a non-linear error that is more difficult to remove. At sampling rates up to 10Gs/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], parasitic capacitance limits its effectiveness. We have developed a double-global sampling technique to remove the effect of parasitic capacitance on 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 bootstrapping still enables multi-GHz input bandwidth.

Power optimization is a major design consideration when implementing a time-interleaved ADC. We will lower total power consumption by exploring innovative technologies for implementing the individual ADCs in the channel, such as the zero-crossing-based circuit (ZCBC) topologies previously presented. In particular, this work is investigating a fast, single-slope architecture (see Figure 1). The primary emphasis is the development of high-speed, highly power-efficient single-slope ZCBC architecture. Since the single-slope architecture is more sensitive to non-idealities such as ramp nonlinearity, we are carefully studying the sources of non-idealities and developing clever techniques to address the accuracy issues.

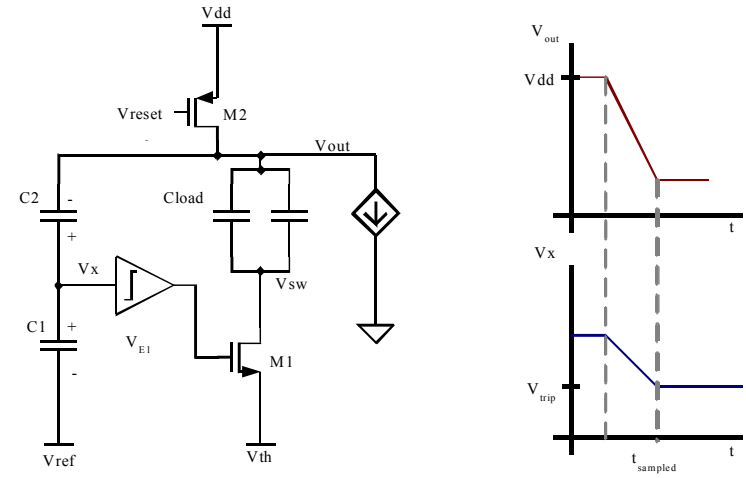

▲ Figure 1: One stage of a single-slope CBSC-based pipelined ADC.

References

[1] M. Gustavsson, "A global passive sampling technique for high-speed switched-capacitor time-interleaved ADCs," *IEEE Trans. on Circuits and Systems II*, pp. 821-831, 2000.

[2] S. Gupta, et al, "A 1GS/s 11b time-Interleaved ADC in 0.13µm CMOS," *Digest of the International Solid State Circuits Conference*, 2006, pp. 576-5771.

A High-bandwidth Zero-crossing-based Sigma-delta ADC

J. Feng, H.-S. Lee

Sponsorship: BAE Systems, CICS

Recently, comparator-based switched-capacitor (CBSC) circuits and zero-crossing-based circuits (ZCBC) were introduced [1, 2] as a viable alternative to op-amp-based circuits. The use of op-amps in analog signal processing circuits is becoming more difficult due to the decreased intrinsic device gains and reduced signal swings obtained in scaled CMOS technologies. Op-amps rely upon high gain in the negative feedback mode in sampled data systems because the gain determines the accuracy of the output value. CBSC and ZCBC-based circuits replace the op-amp using a comparator and a current source and therefore do not require high gain and stability simultaneously, as in op-amp-based circuits. In this work, we explore the use of zero-crossing-based circuits to replace op-amps in high-bandwidth ,discrete time sigma-delta ADCs.

Sigma-delta ADCs are advantageous for realizing high resolutions with relaxed matching requirements between integrated circuit components. Traditionally, sigma-delta ADCs have been limited to voice band signals because of the inherent oversampling in the architecture. However, because of the demand for higher conversion bandwidths such as in broadband internet, sigma-delta ADCs have made significant progress in improving the effective conversion rate while maintaining high resolutions [3].

This research focuses on the development of innovative circuits and sampling architectures to design a zero-crossing-based sigma-delta ADC with a signal bandwidth of >20 MHz, sampling rate of 500 MHz, and 13 bits of accuracy. An example of a possible application appears in VDSL systems. We are investigating chopper-stabilization techniques to remove errors due to offsets, overshoot, and flicker noise in zero-crossing detectors. The use of zero-crossing circuits will ultimately result in large power savings over op-amp-based sigma-delta ADCs.

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 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 pipelined ADC," *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, San Francisco, CA, Feb. 2007, pp. 15-17.

[3] P. Balmelli and Q. Huang, "A 25-MS/s, 14-b, 200-mW ΣΔ modulator in 0.18-µm CMOS," *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, San Francisco, CA, Feb. 2004, pp. 74-75.

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

M. Guyton, H.-S. Lee

Sponsorship: CICS, DARPA

Many analog signal processing circuits use operational amplifiers (op-amps) in a negative feedback topology. 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 switched-op-amp circuit topologies. Figure 1 shows a fully-differential low-voltage ZCB integrator stage using the combined techniques. We are designing 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.

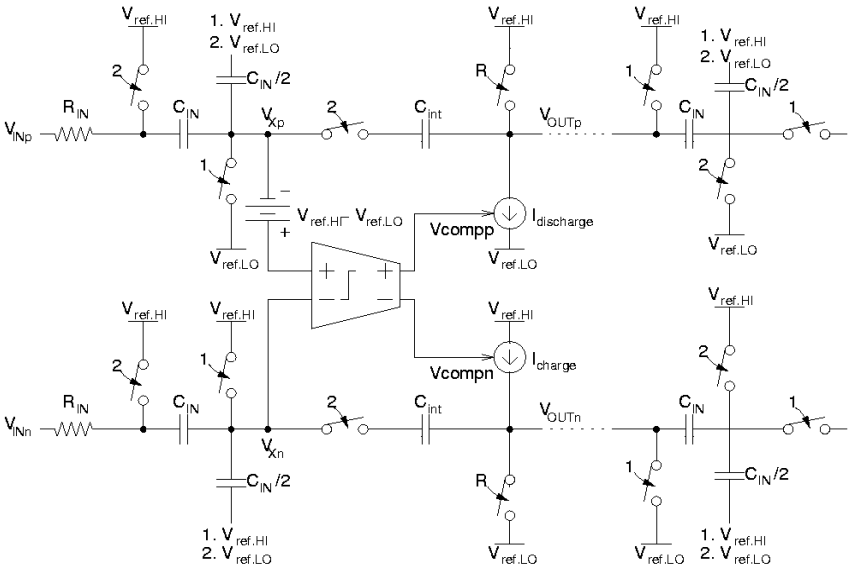

▲ Figure 1: Fully-differential zero-crossing-based switched-capacitor integrator. The input of the next integrator stage is also shown.

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.

A High-performance Zero-crossing-based Pipeline ADC

J. Chu, H.-S. Lee

Sponsorship: SRC/FCRP C2S2

In this work, we are designing a high-performance pipeline ADC using a zero-crossing-based (ZCB) structure [1], which is an extension of the comparator-based switch-capacitor circuit (CBSC) design methodology [2]. The focus of the project is to explore novel circuit structures based on ZCB circuits (ZCBCs) to improve the figure-of-merit (FOM) of ZCBCs. The FOM is generally defined as the power consumption divided by the signal bandwidth and the signal-to-ratio, and it is a measure of the power efficiency of a circuit. The state-of-the-art FOM for a 12-bit ADC with sampling rates of 100MS/s and up is approximately 1pJ/step. In this project, we are investigating pipeline ADCs with 12-bit resolution and 200 Ms/s sampling rate. Our target for the FOM is 10 fJ/step or better, representing 2 orders of magnitude improvement from the state-of-the-art. In particular, we are investigating the use of a differential structure to improve ZCBCs' 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 plan to implement a multi-bit MDAC to improve its power efficiency and to help relax component accuracy requirements. In this design, we plan to use dynamically biased current sources to achieve high linearity at high operating frequencies. Dynamic biasing can be used to compensate for finite output impedance in current sources, which improves the linearity of the system.

Future iterations of this project will use time-interleaving to achieve ultra-high sampling rates (>1GS/s) 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 non-idealities such as gain error, offset, or timing errors can greatly degrade the performance. We plan to use a global sampling technique, which mitigates the timing errors [3-4]. Careful design and layout will be needed to reduce the other mismatches.

References

[1] L. Brooks and H.-S. Lee, "A zero-crossing-based 8b 200MS/s pipelined ADC," *IEEE International. Solid-State Circuits Conference Digest of Technical Papers*, 2007, pp. 15-17.

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

[3] M. Gustavsson, "A global passive sampling technique for high-speed switched-capacitor time-interleaved ADCs," *IEEE Transactions on Circuits and Systems II*, pp. 821-831, 2000.

[4] S. Gupta, et al., "A 1Gs/s 11b time-interleaved ADC in 0.13-µm CMOS," *Digest of International Solid-State Circuits Conference*, 2006, pp. 576-577.

Highly Reconfigurable Zero-crossing-based Analog Circuits

P. Lajevardi, A.P. Chandrakasan, H.-S. Lee

Sponsorship: CICS, DARPA

Switched-capacitor circuits can be used to implement many analog systems, such as ADCs, DACs, filters, amplifiers, and integrators. In an earlier phase of this research, we proposed a reconfigurable switched-capacitor system to implement different analog systems. A prototype system is now being 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 the system. The building blocks have the same functionality and can implement an integrator or a multiplier with a reconfigurable coefficient. There is a reconfigurable connection between blocks. 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 circuits to design ADCs. Zero-crossing circuits can replace the op-amp in the traditional switched-capacitor design with a combination of a current source and a zero-crossing detector. The power consumption of zero-crossing-based analog circuits scales according to the operating frequency and required SNR. In addition, while new technology nodes provide transistors with higher f_t, the design of op-amps is becoming more challenging as the supply voltage and intrinsic gain of transistors are decreasing. Zero-crossing circuits can benefit more from technology scaling due to their more digital-circuit-like operation. In addition, due to the lack of high static gain and stability requirements, the topology of zero-crossing detectors is far simpler than that of op-amps. Therefore, a single topology is suitable for a very wide range of performance space. Zero-crossing circuits are used to implement the reconfigurable analog blocks in this research. The system can operate at different speeds and SNR requirements while the power consumption is kept at the optimum level.

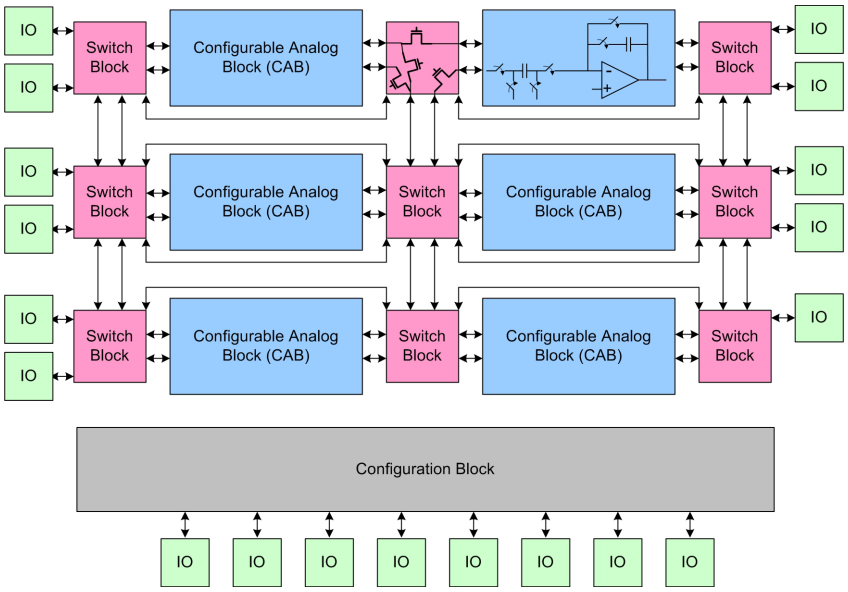

▲ Figure 1: Block diagram of reconfigurable zero-crossing-based analog circuits. Each configurable analog block can be programmed to perform an integration or multiplication. The connection of blocks is also programmable.

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 8-b 200MS/S pipeline ADC," in *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, San Francisco, CA, Feb. 2007, pp. 460-461.

Zero-crossing-based Switched-capacitor Filters

S. Lee, H.-S. Lee

Sponsorship: Korea Foundation for Advanced Studies Scholarship

Design of switched-capacitor circuits in scaled CMOS technologies is becoming difficult because of low device intrinsic gain and reduced power supply voltage. To solve these problems, comparator-based switched-capacitor (CBSC) circuits and zero-crossing-based circuits (ZCBC) were suggested as possible solutions to the op-amp based circuits. In this project, we explore zero-crossing-based circuits (ZCBC) in high-order differential switched-capacitor filters to replace area-inefficient op-amps-based continuous-time base-band filters.

If high-order filters are implemented by cascading second-order filters, the frequency response will be very sensitive to component variations. To implement less sensitive high-order filters, this project will use a low-pass ladder filter network as the prototype for the zero-crossing-based high-order switched-capacitor filter. After several transforming steps, this low-pass ladder filter prototype has been transformed into a differential high-order switched-capacitor filter, which is shown in Figure 1. This architecture has been verified by SWITCAP, giving the frequency response with low-pass characteristics.

Differential op-amps in a high-order switched-capacitor filter shown in Figure 1 are replaced by zero-crossing-based circuits (ZCBC). Figure 2 shows a differential zero-crossing-based circuits (ZCBC) integrator, which is the key component in a high-order switched-capacitor filter, which consists of a comparator and current sources.

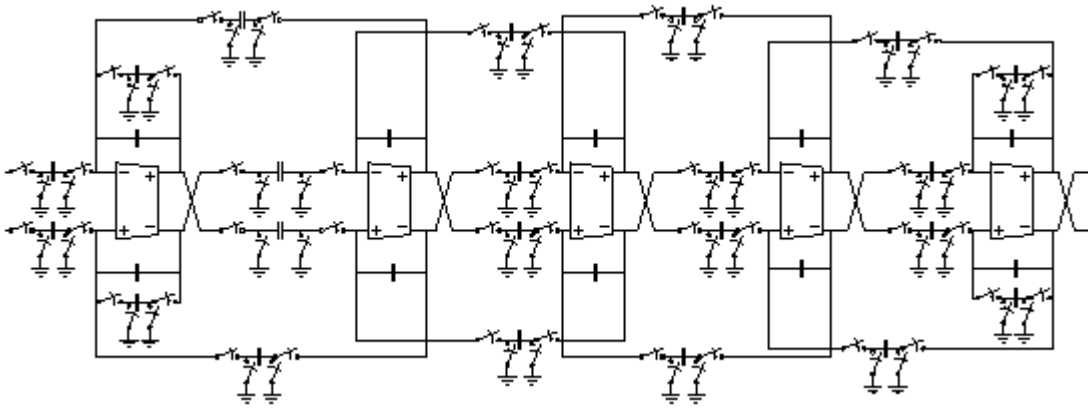

▲ Figure 1: A 5th-order differential-switched-capacitor ladder filter.

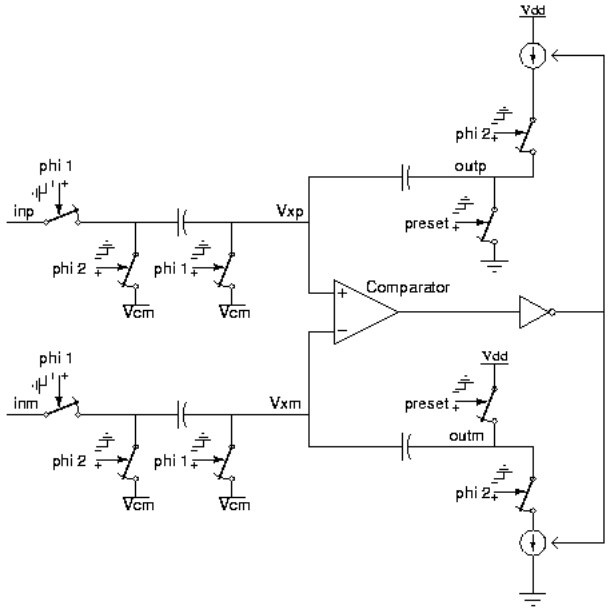

▲ Figure 2: Circuit diagram for a differential ZCBC integrator.

High-accuracy Pipelined A/D Converter Based on Zero-crossing Switched-capacitor Circuits

M. Markova, P. Holloway, H.-S. Lee
Sponsorship: EECS Fellowship, 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) technique has 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 CBSC technique for very high-precision AD converters. The goal of the project is a 100MHz 16-bit pipelined ADC. First, we are investigating multiphase CBSC operation to improve the power-linearity trade off of the A/D conversion [3]. We are also developing linearization techniques for the ramp waveforms. Linear ramp waveforms require fewer phases, thus allowing faster operation. Techniques for improving linearity beyond using a cascoded current source are explored. A linear ramp generator, which decouples the current source from the output ramp through a Miller capacitor, is proposed to improve the linearity of the ramp waveform in all phases. This ramp generator improves the range by improving linearity through compensation of the gate-to-source voltage of the current source without the use of a cascode. In addition it lends itself to a symmetric differential implementation for the final phase to ensure adequate noise rejection. At the target resolution of 16 bits, power supply and substrate noise coupling can limit the performance. We are studying their effects in CBSC circuits. For reduced sensitivity to power supply and substrate noise, we are developing a differential CBSC architecture. Other techniques that we are presently developing include power-efficient offset cancellation in comparators and exploiting *a priori* information from previous stages in the pipeline structure to increase linearity and speed.

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*, vol. 41, no. 12, pp. 2658-2668, Dec. 2006.

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

On-chip Voltage-scalable Switched-capacitor DC-DC Converter

Y.K. Ramadass, A.P. Chandrakasan

Sponsorship: DARPA

Minimizing the energy consumption of battery-powered systems is a key focus in integrated circuit design. Dynamic voltage scaling (DVS) [1] and sub-threshold operation are popular methods to achieve energy efficiency in systems that have widely variant performance demands. However, to realize the full energy benefits of voltage scaling, an efficient voltage-scalable DC-DC converter is of great importance. DVS systems also often require multiple on-chip voltage domains with each domain having specific power requirements. A switched-capacitor (SC) DC-DC converter is a good choice for such battery operated systems because it can minimize the number of off-chip components and does not require any inductors, thereby reducing the overall DC-DC converter volume and cost.

Figure 1 shows the architecture of a switched-capacitor DC-DC converter [2] implemented in 0.18-µm CMOS. The converter employs Pulse Frequency Modulation (PFM) to achieve voltage regulation. In this mode of control, the converter stays idle till the load voltage (V_L) falls below a user-defined reference voltage (V_{REF}), at which point the comparator enables the switch matrix to transfer one charge packet to the load. The switched capacitor DC-DC converter employs on-chip charge-transfer capacitors and can provide scalable load voltages from 0.3V to 1.1V. In order to maintain efficiency over this voltage range, the converter employs 5 different gain settings (G<0:4>) which help in minimizing conduction loss [2]. The automatic frequency-scaler block helps to adjust the frequency of operation and the size of the switches (enW<0:1>) within the switch matrix, with changes in load power. This helps in scaling the switching losses as the load power varies. A divide-by-3 switching scheme [2] was employed in the converter to reduce the parasitic bottom-plate losses and improve efficiency. Also, the all-digital control circuitry used in the converter consumes no static power. The voltage-scalable SC DC-DC converter with integrated on-chip charge transfer capacitors was implemented in National Semiconductor's 0.18-µm CMOS process and consumed an active area of 0.57 mm². The converter achieved above 70% efficiency over a wide range of load powers from 5µW to 1mW, while delivering load voltages from 300mV to 1.1V. Figure 2 shows the efficiency of the converter with change in load voltage while delivering 100µW to the load.

A modified version of this design was implemented as part of a 65-nm sub-threshold microcontroller system [3]. The DC-DC converter in the microcontroller occupied just 0.12 mm² in area and was able to achieve 75% efficiency at a load voltage of 500mV with the microcontroller as the load.

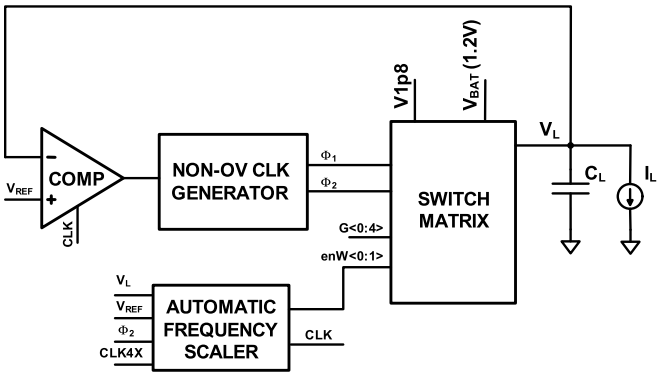

▲ Figure 1: Architecture of the switched-capacitor DC-DC converter employing the PFM mode of control.

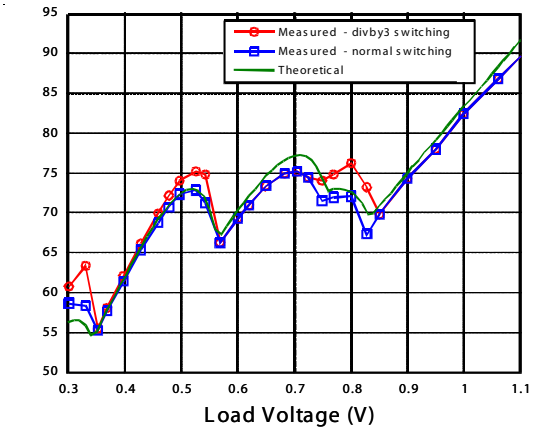

▲ Figure 2: Efficiency plot of the DC-DC converter with change in load voltage.

References

[1] B.H. Calhoun and A.P. Chandrakasan, "Ultra-dynamic voltage scaling using sub-threshold operation and local voltage dithering in 90nm CMOS," *ISSCC Digest of Technical Papers*, pp. 300-301, Feb. 2005.

[2] Y.K. Ramadass and A. Chandrakasan, "Voltage scalable switched capacitor DC-DC converter for ultra-low-power on-chip applications," *IEEE Power Electronics Specialists Conference*, pp. 2353-2359, June 2007.

[3] J. Kwong et al., "A 65nm sub-V_t microcontroller with integrated SRAM and switched-capacitor DC-DC converter," *ISSCC Digest of Technical Papers*, pp. 318-319, Feb. 2008.

A Highly Digital, Resolution- and Voltage-scalable SAR ADC

M. Yip, A.P. Chandrakasan
Sponsorship: DARPA, NSERC

In energy constrained systems such as wireless sensor nodes or mobile electronics, it is desirable to have analog-to-digital converters (ADC) that can operate in the sub-threshold regime to minimize energy during long standby periods but can also dynamically elevate their performance to meet system demands. In this work, a highly digital, energy-efficient successive approximation register (SAR) ADC with scalable resolution from 5 to 10-bits is being designed. Recent SAR ADCs have achieved energy efficiencies on the order of a few fJs per conversion step, but only at a fixed resolution [1]. The main challenge in designing a scalable ADC is maintaining its energy-efficiency across all resolutions.

This SAR ADC consists of a comparator, a digital-to-analog converter (DAC) and digital control logic to implement a binary search algorithm. Figure 1 shows the system level block diagram. Clock gating is used to enable sample rate scaling from 1-MS/s down to 0. Special techniques are used to deactivate extraneous circuitry as resolution is reduced to maintain energy efficiency. Often, this scalability involves using analog switches at critical nodes and must therefore be carefully designed, especially at low supply voltages.

Lastly, voltage scaling is used to maintain a constant energy efficiency as the resolution is reduced. Recently, SAR ADCs operating on a 500-mV supply have been reported [2]. This ADC will be designed to operate at 10-bits at 1-V, down to 5-bits at 400-mV. Figure 2 shows how voltage scaling allows the ADC to maintain constant energy efficiency over all resolutions.

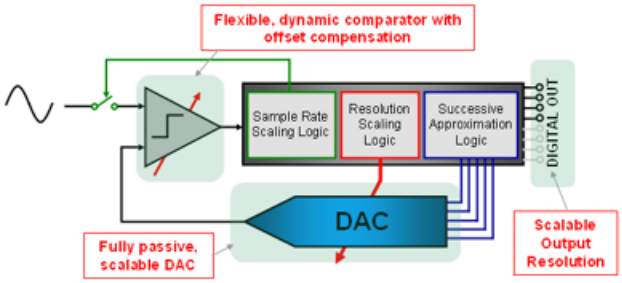

▲ Figure 1: System-level block diagram of scalable ADC. Digital logic is used to control the resolution, the sample rate and the comparator mode of operation.

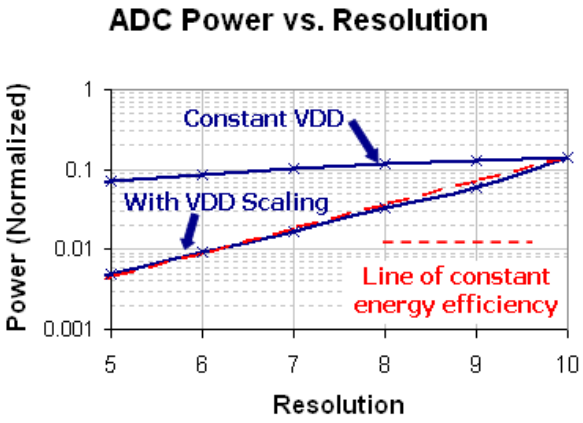

▲ Figure 2: ADC power vs. resolution. Voltage scaling from 1-V down to 400-mV is used to maintain energy efficiency across all resolutions.

References

[1] 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," in *Proc. IEEE International Solid-State Circuits Conference*, San Francisco, CA, Feb. 2008, pp. 244-245.

[2] S. Gambini and J. Rabaey, "Low-power successive approximation converter with 0.5 V supply in 90 nm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 42, no. 11, pp. 2348-2356, Nov. 2007.

An AES-based Energy-efficient Encryption Processor with Resistance to Differential Power Analysis 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 processor. 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., "SubBytes," "ShiftRows," "MixColumns," and "AddRoundKey"(Figure 1). For the design of an energy-efficient processor with performance requirement, two architectural approaches can be used with voltage scaling. Since a total of 32 bytes is manipulated in one round of transformation, parallel operation of each byte can compensate for the reduced speed that results from the supply voltage scaling. Pipelining also helps maintain the throughput even with the reduced supply voltage.

The DPA attack is based on the asymmetry in power dissipation, depending on the input data. Dynamic differential logic consumes the same dynamic power regardless of any input data [2], since during the precharge phase, one node is always precharged to VDD, and during the evaluation phase, one node is always discharged to ground. Resistance to DPA attack with dynamic differential logic, however, comes at the expense of increased power dissipation. Therefore, the trade-off between security and power dissipation should be examined carefully and optimized for the specified design purposes.

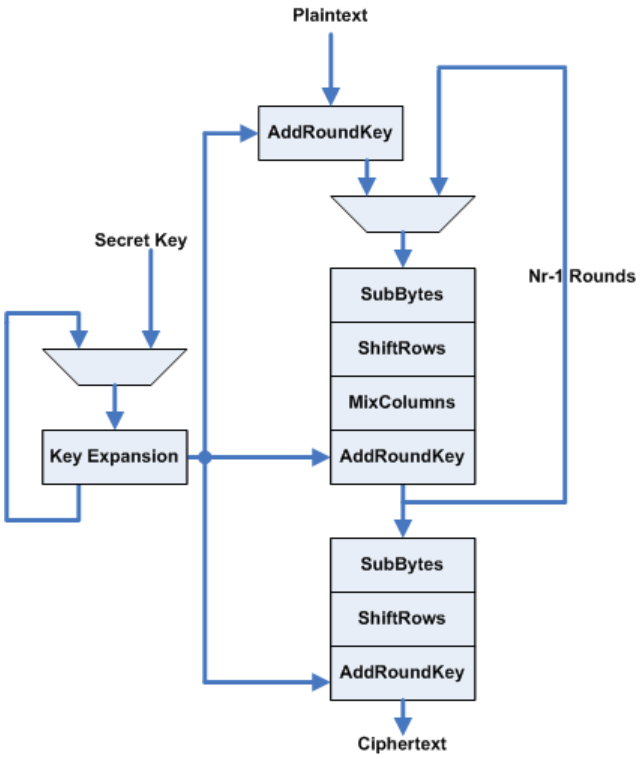

▲ Figure 1: AES algorithm with a repeated round transformation

References

[1] NIST, "Federal information processing standards PUB 197 advanced encryption standard" (Nov. 2001). [Online]. Available: http://www.nist.go/aes/

[2] I. Verbauwhede, P. Schaumont, and H. Kuo, "Design and performance testing of a 2.29-GB/s Rijndael processor," *IEEE Journal of Solid-State Circuits*, vol. 38, no. 3, pp. 569-572, Mar. 2003.

A Micropower DSP for Sensor Applications

N. Ickes, A.P. Chandrakasan

Sponsorship: DARPA, Texas Instruments

Distributed microsensor networks consist of hundreds or thousands of miniature sensor nodes. Each node individually monitors the environment and collects data as directed by the user, and the network collaborates as a whole to deliver high-quality observations to a central base station. The large number of nodes in a microsensor network enables high-resolution, multi-dimensional observations and fault-tolerance that are superior to more traditional sensing systems. However, the small size and highly distributed arrangement of the individual sensor nodes make aggressive power management a necessity.

The aim of our project is to develop a micropower DSP platform optimized for medium-bandwidth (up to 100 kHz) microsensor applications, such as acoustic sensing and tracking. These applications require significant signal processing capability at each node within a sensor network, while maintaining a roughly 100 µW average power consumption to enable self-powered (energy scavenging) operation. As illustrated in Figure 1, our DSP includes a general-purpose processor core with an energy efficient instruction set, as well as coprocessors for accelerating Fourier transforms and FIR filtering. Power consumption in the large (60 kB) on-chip memory is reduced by dividing the memory into banks (to reduce access energy) and by power-gating inactive banks (to reduce leakage energy). The CPU, FIR, and FFT cores are also power-gated. The DSP was fabricated in 90-nm CMOS by ST Microelectronics.

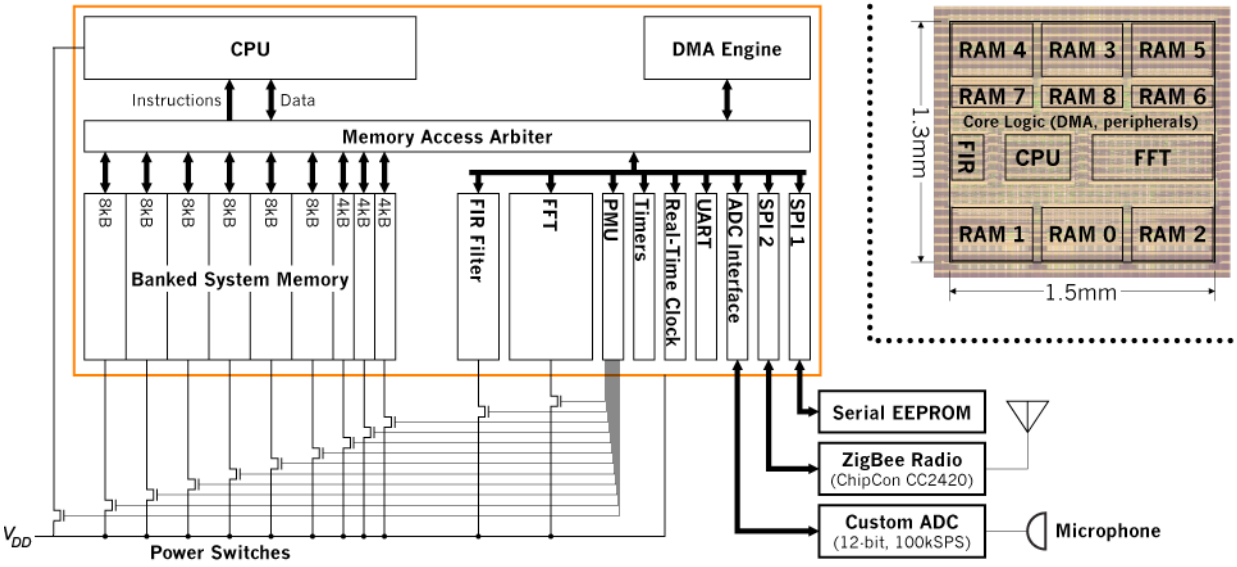

▲ Figure 1: DSP architecture, illustrating the twelve independent power domains, controlled by off-chip power switches. When combined with an external nonvolatile memory (for program storage), radio, and ADC, the DSP becomes a complete microsensor node. Inset: Die photo of the DSP, implemented in 90-nm CMOS.

A 65-nm Sub-V_t Microcontroller with Integrated SRAM and DC-DC Converter

J. Kwong, Y. Ramadass, N. Verma, A.P. Chandrakasan

Sponsorship: DARPA, Texas Instruments Graduate Woman's Fellowship (J. Kwong), Intel Foundation Ph.D. Fellowship (N. Verma)

Aggressive scaling of the power supply to below the device threshold voltage (V_t) is a compelling approach for energy minimization in digital circuits [1]. Although circuits exhibit slower speeds at low supply voltages, the trade-off remains attractive for energy-constrained systems with relaxed throughput constraints. However, effects of process variation become more prominent at low supply voltages, impacting functionality. A 65-nm sub-V_t microcontroller (see Figure 1) demonstrates several approaches to enable operation down to 300mV [3]. In sub-V_t, logic gates no longer exhibit rail-to-rail voltage swings due to variation and reduced ratio of on-to-off currents. A standard cell library design methodology increases device sizes appropriately to mitigate these effects [4]. Moreover, circuit delays exhibit order-of-magnitude higher variability at low voltages. Conventional timing analysis approaches that treat delays as deterministic are insufficient. Instead, a variation-aware methodology combining simulation and analysis was developed to verify hold-time constraints.

The SRAM represents a dominant portion of power and area in this system. Therefore, energy and leakage reduction through voltage scaling is highly desirable. In conventional 6-T SRAMs, V_t variation causes severely degraded read-current and increased cell instability, limiting the minimum functional voltage. The SRAM in this chip employs an 8-T bit-cell to address these limitations [5]. Further, peripheral circuit assists enforce the relative device strengths needed for read and write functionality, even in the presence of significant variation. A fully integrated, switched-capacitor DC-DC converter provides highly efficient power delivery at the low voltage and power levels required by energy-constrained systems [6]. Featuring multiple gain settings and efficient control circuitry, the DC-DC converter can deliver variable load voltages and achieves above 75% efficiency while supplying 500mV in the range of 10µW to 250µW. The microcontroller chip (see Figure 2) was fabricated by Texas Instruments.

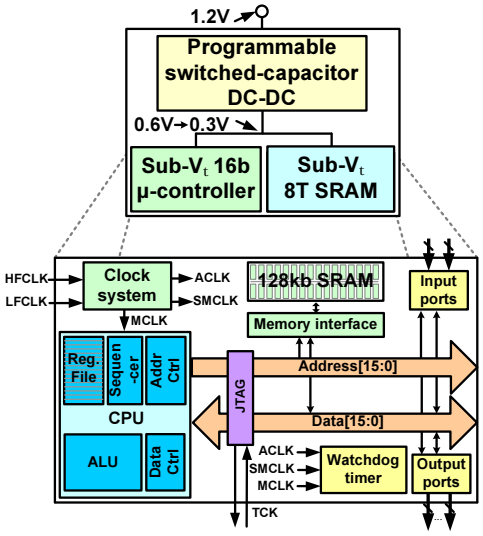

▲ Figure 1: Block diagram of 65-nm sub-V_t system-on-a-chip and detailed diagram of microcontroller core.

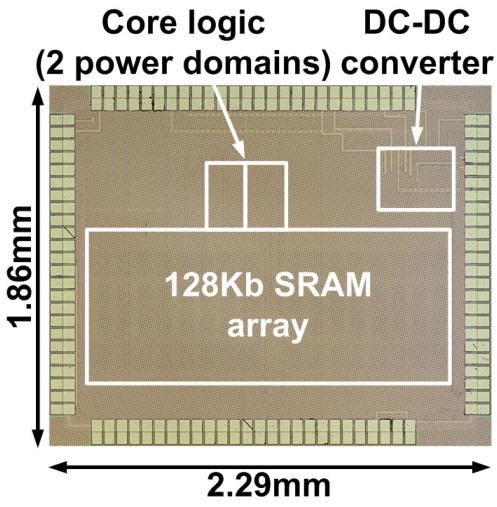

▲ Figure 2: Die micrograph of prototype test-chip.

References

[1] A. Wang, A. Chandrakasan, and S. Kosonocky, "Optimal supply and threshold scaling for sub-threshold CMOS circuits," *Proc. IEEE Computer Society Annual Symposium on VLSI*, Apr. 2002, pp. 7-11.

[2] B. Zhai, L. Nazhandali, J. Olson, et al., "A 2.60pJ/Inst sub-threshold sensor processor for optimal energy efficiency," *IEEE Symposium on VLSI Circuits Dig. Tech. Papers*, June 2006, pp. 154-155.

[3] J. Kwong, Y. Ramadass, N. Verma, et al., "A 65nm sub-V_t microcontroller with integrated SRAM and switched-capacitor DC-DC converter," *IEEE International Solid-State Circuits Conference Dig. Tech. Papers*, Feb. 2008, pp. 318-319.

[4] J. Kwong and A.P. Chandrakasan, "Variation-driven device sizing for minimum energy sub-threshold circuits," *Proc. IEEE International Symposium on Low Power Electronics and Design*, Oct. 2006, pp. 8-13.

[5] N. Verma and A.P. Chandrakasan, "A 65nm 8T sub-V_t SRAM employing sense-amplifier redundancy," *IEEE International Solid-State Circuits Conference Dig. Tech. Papers*, Feb. 2007, pp. 328-329.

[6] Y.K. Ramadass and A.P. Chandrakasan, "Voltage-scalable switched capacitor DC-DC converter for ultra-low-power on-chip applications," *IEEE Power Electronics Specialists Conference*, pp. 2353-2359, June 2007.

Multi-stage Converter Architectures for Microprocessor Power Delivery

R. Pilawa-Podgurski, D.J. Perreault
Sponsorship: SRC/FCRP IFC

Today's low processor operating voltages and high currents result in increased ohmic losses in the interconnect between the voltage regulator and the processor. To reduce these losses, it is desirable to place a large step-down voltage regulator in close proximity to the processor. This research investigates power converter architectures that make use of available on-die device characteristics to provide large voltage step-down and high bandwidth regulation.

In a given CMOS process, we typically have access to slow devices with a moderate blocking voltage and fast devices with low blocking voltage. Figure 1 shows a proposed multi-stage converter architecture that makes use of these different device characteristics. A switched-capacitor transformation stage utilizing slow, high-voltage switches provides efficient unregulated voltage step-down. This stage is followed by a high-bandwidth magnetic regulation stage, which utilizes fast low-voltage devices. Excellent efficiency and power density can be achieved through this technique, as can large voltage step-down and fast transient response.

Coupling the fast magnetic regulating stage to the switched-capacitor stage in a novel manner can eliminate the current spikes associated with conventional switched–capacitor converters. This "merged two-stage" architecture also improves overall efficiency and drastically reduces the total required capacitance. Figure 2 shows a photograph of a prototype implemented with discrete components [1]. Experimental measurements demonstrate the benefits of the merged two-stage architecture, and this proof-of-concept converter verifies the validity of the approach.

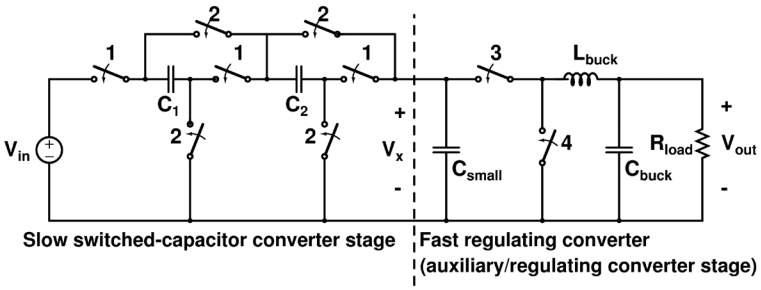

▲ Figure 1: Schematic of proposed merged two-stage converter.

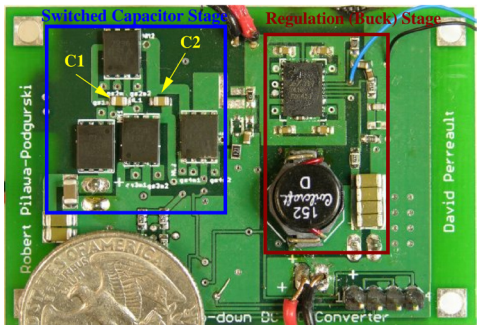

▲ Figure 2: Photograph of merged two-stage converter prototype implemented with discrete components.

Reference

[1] R. Pilawa-Podgurski, D. Giuliano, and D. Perreault, "Merged two-stage power converter architecture with soft charging switched-capacitor energy transfer," *IEEE Power Electronics Specialists Conference*, Rhodes, Greece, June 2008.

Enhanced Low-voltage 8T SRAM with Body Bias Adjustment

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

Body bias presents an attractive—albeit low-frequency—mechanism to correct for mean deviations in device parameters because the electrical manipulation in a global fashion of the body terminals of bulk CMOS devices incurs minimal area overhead over standard body contacts. See Figure 2 for the illustration of minimal area penalty for 8T SRAM. Conventional work on body biasing in SRAM [1] focuses on improving yield by optimizing the strength ratio between pfet and nfet for static functionality of the 6T cell. This work focuses on the compelling advantages of 8T SRAM (cell schematic shown in Figure 1) to enable lower VDD operation, akin to prior low-VDD body biasing work on static CMOS digital circuits [2].

The key challenges relate to detecting the optimum body bias. Not only do accelerated leakage currents and junction currents at forward bias and gate oxide reliability degradation at very negative biases set maximum limits for the p-well and n-well voltages, but active retention failures have to be delicately balanced with active write errors. A technique to detect and set the optimum body bias is being developed and implemented in IBM 65-nm bulk CMOS technology.

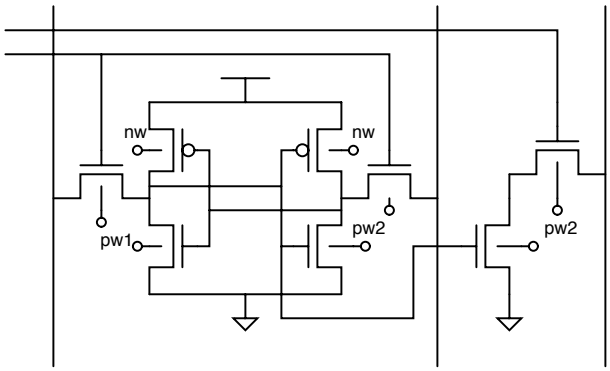

▶ Figure 1: Shown to the right is the schematic for one 8T cell with the body terminals explicitly marked for each device.

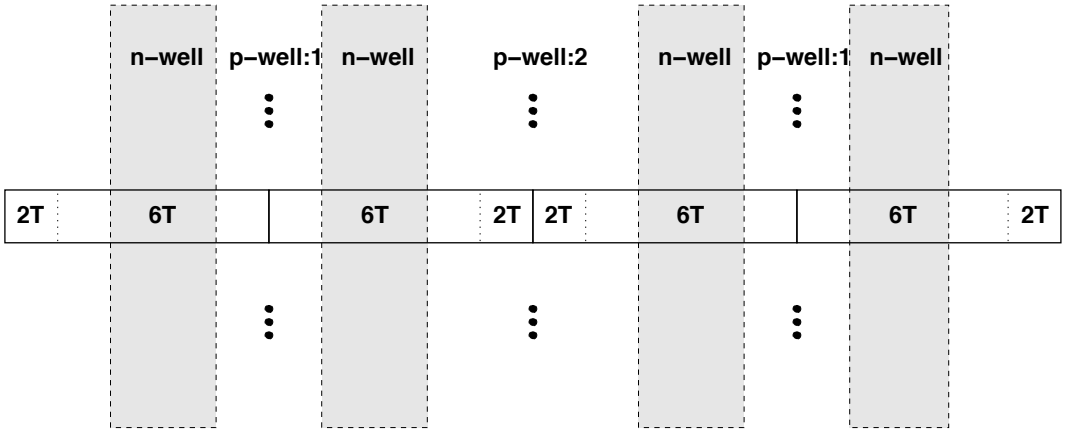

▲ Figure 2: Shown above is the physical layout of the 8T cell array. The nwell and pwell run vertically in alternating strips.

References

[1] S. Mukhopadhyay, K. Kim, H. Mahmoodi and K. Roy, "Design of a process variation tolerant self-repairing SRAM for yield enhancement in nanoscaled CMOS," *IEEE Journal of Solid-State Circuits*, vol. 42, pp. 1370-1382, 2007.

[2] G. Ono and M. Miyazaki, "Threshold-voltage balance for minimum supply operation," *IEEE Journal of Solid-State Circuits*, vol. 38, pp. 830-833, 2007.

SRAM Design for Ultra-low-power Systems

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

Low-power circuit design has been an important research area because of the ever-increasing need for longer battery life in energy-starved applications. Most digital systems require on-chip memory blocks that dominate not only area but also power consumption. Hence, lowering SRAM voltage is critical. First, this reduces the active energy, which is given by CV_{DD}^2. Figure 1a shows energy/access vs. V_{DD} plot for a 64-kbit SRAM block. The total energy makes a minimum around 400mV because of the opposing trends in the active energy and leakage energy components. Second, leakage-currents also decrease at lower supply voltages, resulting in reduced leakage-power. Figure 1b shows the leakage power vs. V_{DD} plot. The leakage power decreases by ~50X over the range. This significant decrease is due not only to scaling in V_{DD} but also to alleviation of drain-induced barrier lowering (DIBL).

Designing functional SRAMs at lower supply voltages is, however, extremely challenging due to the increasing effects of local variation with device scaling. To maximize density, the bit-cell is designed to be very small, aggravating its variability. Hence, as the supply voltage decreases, read and write failures begin to be highly prominent. Peripheral assists to address these and new bit-cell topologies have been proposed recently in [1] and [2]. However, architectural innovations are also required.

Operating voltage of an SRAM is often bounded by the performance requirement of the system. In such a scenario, the operating voltage can be brought down only if the performance of the SRAM is improved. Figure 2 shows the normalized read access time vs. V_{DD} with different offset voltages for the sense amplifier shown. Since the sense amplifier is in the critical path, its offset directly determines the discharge time, and hence a larger offset translates into a longer access period. In order to acquire the same performance at a lower V_{DD}, sense-amplifier offset voltage should be analyzed. [3] analyzes a very commonly used sense amplifier structure but does not provide a model for the offset. An offset compensation scheme that would trim the reference voltage value to minimize the offset can be implemented. This scheme should be implemented with a reasonable area overhead in order to be applicable to SRAMs.

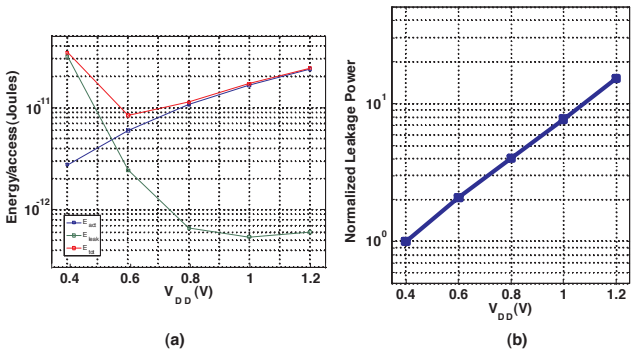

▲ Figure 1: Simulated energy/access vs. V_{DD} plot (a) and leakage power vs. V_{DD} (b) plot for a 64-kbit SRAM in 65-nm technology.

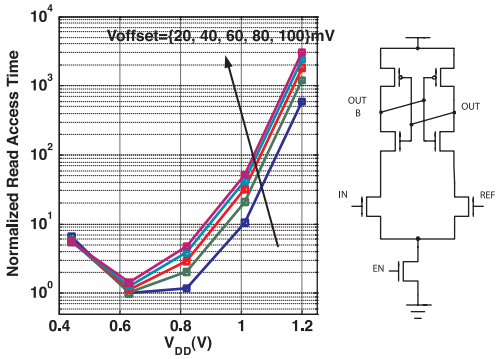

▲ Figure 2: Normalized read access time vs. V_{DD} plot with different offset voltages for the sense amplifier shown.

References

[1] N. Verma and A.P. Chandrakasan, "A 65-nm 8-T Sub-V_t SRAM employing sense-amplifier redundancy," *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, Feb. 2007, pp. 328-329.

[2] T.-H. Kim, J. Liu, J. Keane, and C.H. Kim, "A high-density sub-threshold SRAM with data-independent bitline leakage and virtual ground replica scheme,"*IEEE Int. Solid-State Circuits Conf. (ISSCC) Digest of Technical Papers*, Feb. 2007, pp. 330-331.

[3] B. Wicht, T. Nirschl, and D. S-Landsiedel, "Yield and speed optimization of a latch-type voltage sense amplifier," *IEEE Journal of Solid-State Circuits*, vol. 39, pp. 1148-1158, July 2004.

An Adaptive Fractionally Spaced Receive Equalizer

S. Song, V. Stojanović

Sponsorship: National Semiconductor

Based on voltage-to-time conversion technique [1], a pseudo-differential two-way-interleaved adaptive receive equalizer with two 2x-oversample feed-forward taps and one feedback tap has been designed in 90-nm CMOS process. It integrates equalization and phase interpolation functions into one unit to simultaneously address intersymbol-interference cancellation and phase synchronization in a link receiver. According to post-layout simulation, the receiver can operate at 6 Gbps with 6.48 mW of power. The filter is two-way interleaved to overcome technology speed limitations. Figure 1 shows block and circuit diagrams for one direction. The zero-crossing-detector circuit functions as a voltage-to-time converter followed by a dynamic discharger stage that works as a multiplying time and voltage converter. In comparison with [1], a positive feedback PMOS with gate connected to the drain of M1 can reduce t_f by more than 15%. Buffers are placed at the output to increase the driving ability. A data-conditioned sign-sign least-mean squares algorithm has been developed off-chip, to tune the filter coefficients and overcome the convergence issues due to strong signal correlation. Figure 2 shows the power partition from the post-layout simulation, indicating a 1-pJ/bit overall energy-cost for the receiver.

Tuning the external discharge currents I_1, I_2, and bias current for the slicer adjusts the feed-forward/feedback tap weights. On the band-limited channels, 2x oversampling is enough to reconstruct the signal sample completely, regardless of the phase. Tuning the tap weights I1 and I2 makes it possible to reopen the eye, no matter what the phase offset between Tx and Rx. Therefore, it not only equalizes the channel but also alleviates the requirement for an explicit phase synchronization loop. In links that employ decision-feedback equalization at the receiver, these phase synchronization loops are difficult to design due to the excessive amount of intersymbol-interference-induced jitter at symbol edges (i.e., the points at which the circuit derives the timing error information), leading to poor performance of synchronization.

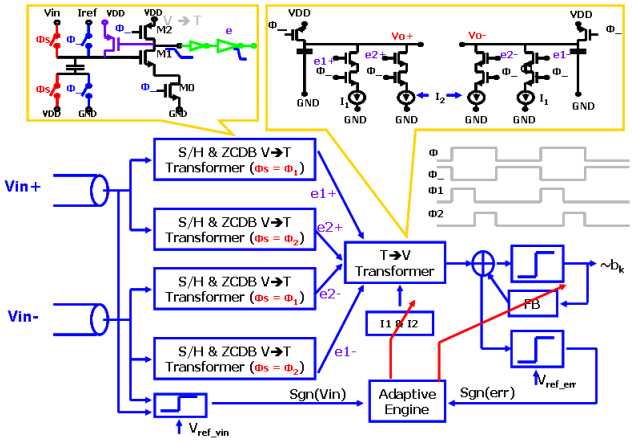

▲ Figure 1: System structure for one way of equalizer. The adaptive engine's inputs are from one direction but its outputs, including I1, I2 and FB tap weight, are shared between the two branches, enabling fully differential operation.

CLK buffer and regeneration	2.88mW
V→T and T→V conversion	1.92mW
Signal and error detectors	0.18mW
DFE	1.5mW
Total	6.48mW

▲ Figure 2: Power partition based on post-layout simulation with 6- Gbps data rate and 1.2-V supply.

Reference

[1] L. Brooks and H.-S. Lee, "A zero-crossing-based 8b 200 MS/s pipelined ADC," *IEEE Journal of Solid-state Circuits*, vol. 42, no. 12, Dec. 2007.

Algorithms and Architectures for Ultra-low-power Video Compression

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

Multimedia applications, such as video playback, are becoming increasingly pervasive. Since the platforms are often energy-constrained devices (cell phones, iPods), the user experience is enhanced by extending the battery life during video decoding. The latest video coding standard is H.264 [1], and it is used in DVB-H and HDTV. While it provides a 50% improvement in compression efficiency over previous standards, this coding efficiency comes at the cost of increased decoder complexity of 4X over MPEG-2 and 2X over MPEG-4 Visual Simple Profile. This increased complexity translates to increased energy consumption, which is a critical concern for mobile and handheld devices.

Our aim is to build an ASIC decoder that exploits techniques such as memory optimization, pipelining, parallelism, ultra-low-voltage operation, and ultra-dynamic voltage scaling [2]. For instance, the Deblocking Filter computation can be parallelized as shown in Figure 2. In video decoders, memory consumes a large portion of overall system power. As a result, the number of redundant memory transfers must be minimized and caching data in on-chip SRAMs/registers is paramount. Using these techniques, the goal is to reduce system power even further than previously published decoders [3]. In addition to optimizing the hardware architecture of the H.264 decoder (Figure 1), we will also focus on the design of next-generation video coding algorithms, e.g., "H.265." We envision that these algorithms will account for the energy and complexity costs of their hardware implementations. By incorporating the energy-awareness into the algorithm, future video coders can provide an explicit energy/PSNR trade-off, along with the existing bitrate/PSNR trade-off curves.

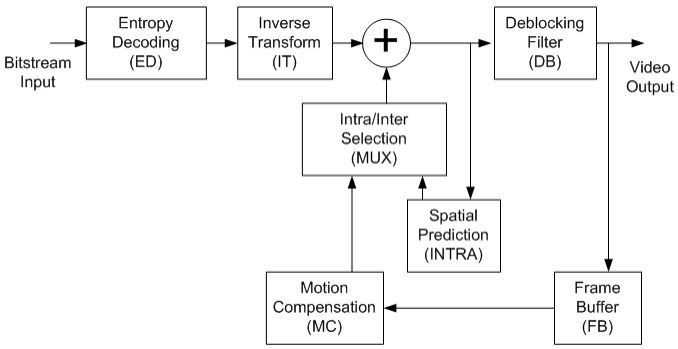

▲ Figure 1: H.264 video decoder architecture.

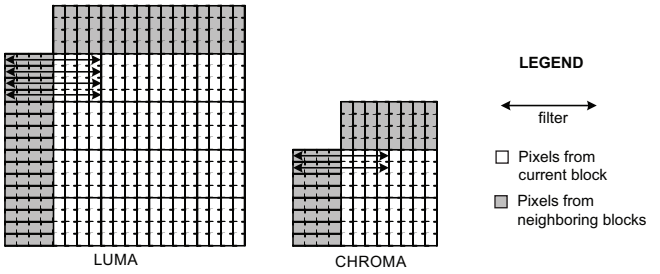

▲ Figure 2: Parallel Deblocking Filters

References

[1] ITU-T Study Group, "Series H: Audiovisual and multimedia systems: Infrastructure of audiovisual services – coding of moving video," in *General Secretariat and Telecom Radiocommunication (ITU-R) Standardization (ITU-T)*, sec. H, no. 264, Sept. 2005.

[2] B.H. Calhoun and A.P. Chandrakasan, "Ultra-dynamic voltage scaling using sub-threshold operation and local voltage dithering in 90nm CMOS," in *Proc. IEEE International Solid-State Circuits Conference*, San Francisco, CA, Feb. 2006, pp. 300-301.

[3] C.D. Chien et. al., "A 252kgate/71mW multi-standard multi-channel video decoder for high-definition video applications," in *Proc. IEEE International Solid-State Circuits Conference*, San Francisco, CA, Feb 2007, pp. 282-283.

A High-density 45-nm SRAM Using Small-signal Non-strobed Regenerative Sensing

N. Verma, A.P. Chandrakasan

Sponsorship: CICS, Intel Foundation Ph.D. Fellowship Program

High-density embedded SRAMs are critical enablers of the tremendous integration trends benefiting integrated circuits every technology node. Their wide use of aggressive minimum-sized features, however, severely aggravates variability, leading to numerous limitations across the entire array. Some of the most critical among these include highly degraded static-noise-margin (SNM) [1] and cell read-current due to the use of small bit-cells; the need to increase sense-amplifier area in order to reduce its variation and input-offset [2] to withstand the lower read-currents; and finally, excessive margining in the sense-amplifier strobe signal, which is required to overcome its severe tracking divergence across operating corners, with respect to the array read-path.

In order to simultaneously address all of these issues, we present the non-strobed regenerative sense-amplifier (NSR-SA), shown in Figure 1 [3]. It performs very simple offset compensation to overcome the sense-amplifier offset-area trade-off without significantly loading the high-speed nodes. The resulting improvement in stability, even in the presence of extreme variation, implies that a small cell read-current can be accommodated for the same array performance, and, accordingly, the bit-cell can be optimized for other parameters, such as read SNM. Lastly, the presence of stable internal voltage references, generated thanks to the offset compensation, is exploited to trigger self-regeneration with respect to the input bit-line voltage itself, rather than to an explicit external control path. The test-chip, fabricated in 45-nm CMOS, employs high-density 0.25-μm² bit-cells and is shown in Figure 2. Measurements from 53 die show an improvement in the sigma of the delay distribution by a factor of four compared to a conventional sense-amplifier, confirming the benefit of offset-compensation and self-regeneration.

IC fabrication is provided by Texas Instruments.

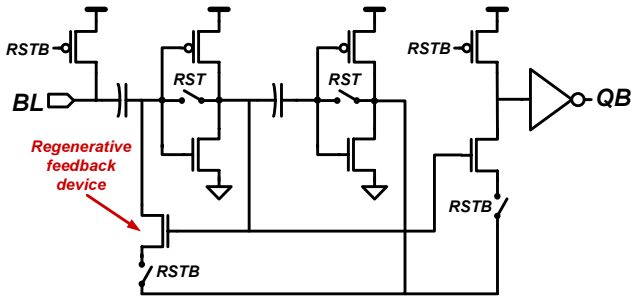

▲ Figure 1: Non-strobed regenerative sense-amplifier circuit schematic.

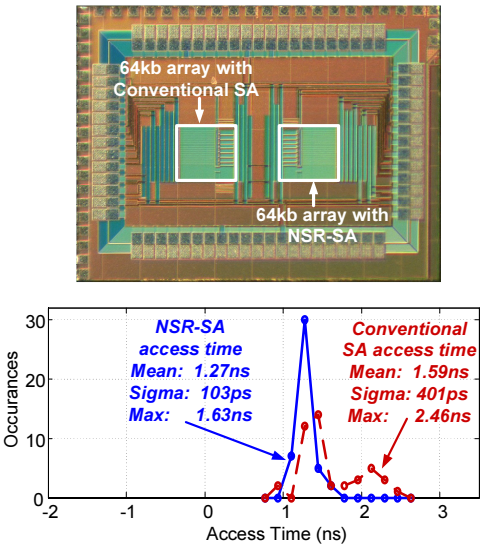

▲ Figure 2: Test-chip die photo and measurement results from 53 chips.

References

[1] A. Bhavnagarwala, X. Tang, and J. Meindl, "The impact of intrinsic device fluctuations on CMOS SRAM cell stability," *IEEE Journal of Solid-State Circuits*, vol. 36, no. 4, pp. 658-665, Apr. 2001.

[2] K. Zhang, K. Hose, V. De, and B. Senyk, "The scaling of data sensing schemes for high speed cache design in sub-0.18μm technologies," in *Proc. IEEE Symposiumon VLSI Circuits*, June 2000, pp. 226-227.

[3] N. Verma and A.P. Chandrakasan, "A high-density 45nm SRAM using small-signal non-strobed regenerative sensing," *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, Feb. 2006, pp. 328-329.

Systematic Information-fusion Methodology for Static and Dynamic Obstacle-detection in ITS

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 perception abilities, we should automatically detect static and dynamic obstacles and obtain their related information, such as locations, speed, possibility of 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- and layered-based information-retrieval methodologies to systematically detect obstacles and obtain their location and timing information for visible and infrared sequences. The proposed obstacle-detection methodologies take advantage of connections among different pieces of information and increase the computational accuracy of obstacle information estimation, thus improving environment-understanding abilities and driving safety. Figures 1 and 2 show two examples.

▲ Figure 1: Segmentation result for urban daytime driving environment.

▲ Figure 2: Segmentation result for night driving environment.

References

[1] A.R. Bruss and B.K.P. Horn, "Passive navigation," *Computer Vision, Graphics, and Image Processing*, vol. 21, no. 1, pp. 3-20, Jan. 1983.

[2] J.E. Tanner, "Integrated optical motion detection," Ph.D. Thesis, California Institute of Technology, Pasadena, CA, Technical Report 5223:TR:86, 1986.

[3] J.E. Tanner and C.A. Mead, "An integrated analog optical motion sensor," *VLSI Signal Processing II, (Proc. of the ASSP Conference on VLSI Signal Processing)*, UCLA, Nov. 5-7, 1986, pp. 59-76..

[4] B.K.P. Horn and S.Negahdaripour, "Direct passive navigation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-9, no. 1, pp. 168-176, Jan. 1987.

[5] B.K.P. Horn, "Motion fields are hardly ever ambiguous," *International Journal of Computer Vision*, vol. 1, no. 3, pp. 259-274, Oct. 1987.

[6] B.K.P. Horn and E.J. Weldon, Jr., "Direct methods for recovering motion," *International Journal of Computer Vision*, vol. 2, no. 1, pp. 51-76, Jun. 1988.

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

[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," special issue on "In-Vehicle Computer Vision Systems" of *IEEE Transactions on Vehicular Technology*, vol. 53, no.6, pp. 1679-1697, Nov. 2004.

[9] B.K.P. Horn, Y. Fang, and I. Masaki, "Time to contact relative to a planar surface," presented at the *IEEE Intelligent Vehicles Symposium*, 2007.