

Circuits and Systems for Information Processing, Communications, Multimedia, Energy Management, and Sensing

12-bit 250MS/s CMOS Pipelined Analog-to-Digital Converter.....	3
Time-Interleaved A/D Converters	4
A Flash ADC with a Reduced Number of Comparators.....	5
Continuous-Time Delta-Sigma Modulator for Next-Generation Wireless Applications	6
Ultra High-Performance GaN-on-Silicon Analog-to-Digital Converters.....	7
High-Efficiency, Low-Leakage RF Transmitters for Low Duty Cycle Applications.....	8
Broadband Inter-Chip Link Using Terahertz Wave on Dielectric Waveguide	9
Monolithic GaN-MMIC MEMS-Based Oscillators	10
Vertical Noise Coupling Mitigation in 3D-IC Using a Solenoid Inductor.....	11
Energy and Area-Efficient Hardware Implementation of HEVC Inverse Transform and Dequantization.....	12
A Navigation Device with 3-D Computer Vision Processor for Visually Impaired People	13
Energy-Efficient Hardware for Multi-Modal Object Detection.....	14
A Deeply Pipelined CABAC Decoder for High Efficiency Video Coding Supporting Level 6.2 High-Tier Applications ...	15
Towards High-Performance Bufferless NoCs with SCEPTER.....	16
Energy-Efficient SRAM using Data-Dependency	17
Energy-Efficient SRAM Design in 28-nm FDSOI Technology	18
New AC-DC Power Factor Correction Architecture Suitable for High Frequency Operation.....	19
A High-Power-Density Wide-Input-Voltage-Range Isolated DC-DC Converter with a Multi-Track Architecture.....	20
A Systematic Approach to Modeling Impedances and Current Distribution in Planar Magnetics	21
Investigating Magnetic Materials for Power Conversion at High Frequency	22
Energy-Efficient Wireless Power Transmitters and Receivers.....	23
Picowatt Timer for Energy-Constrained, Battery-Less Systems.....	24
Solar Energy Harvesting System with Integrated Battery Management and Startup using Single-Inductor and 3.2 nW Quiescent Power	25
Authentication Tags for Supply Chain Integrity	26
Low-Power Sensor Interfaces for Wireless Sensor Nodes.....	27
A Graphene-CMOS Hybrid Sensor for Thermal Imaging	28

12-bit 250MS/s CMOS Pipelined Analog-to-Digital Converter

H. H. Boo, H.-S. Lee, D. S. Boning

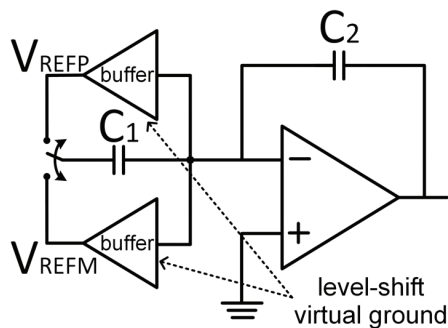
Sponsorship: Masdar Institute / MIT Cooperative Program

The virtual ground reference buffer technique is introduced as a solution to improve the feedback factor of a closed-loop circuit and is demonstrated in a high performance pipelined CMOS analog-to-digital converter (ADC) prototype. The technique enhances the performance of switched-capacitor circuits by improving the feedback factor of the op-amp without affecting the signal gain. The bootstrapping action of level-shifting buffers relaxes op-amp performance requirements in unity-gain bandwidth, noise, open-loop gain, and offset compared with conventional circuits. The improvements substantially reduce the design complexity and the power consumption of the op-amps in applications such as ADCs.

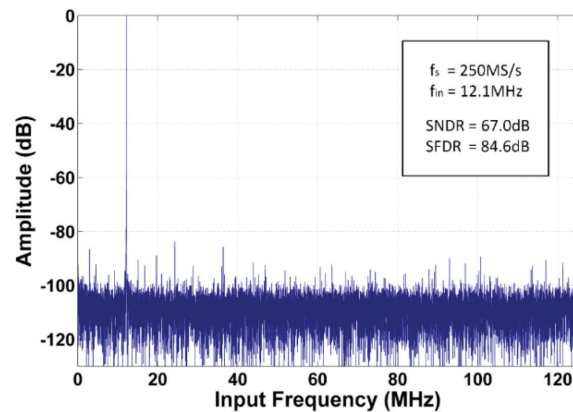
The proposed virtual ground reference buffer technique in its charge-transfer phase is shown in Figure 1. Similarly to conventional pipelined ADC circuits, C_2 flips around the op-amp and C_1 is driven by either the positive or negative reference voltage. Here, however, the reference voltages are referenced to the virtual ground node instead of the system ground, and they are

generated by level-shifting the virtual ground potential. Assuming an ideal buffer, any change in the virtual ground node voltage is reflected at the output of the buffer, effectively bootstrapping C_1 away. Therefore, the C_1 capacitance is removed from the feedback network of the op-amp, resulting in an ideal unity feedback factor independent of the signal gain. The unity feedback factor improves the closed-loop bandwidth and op-amp noise referred to the ADC input by a factor of the signal gain compared with the conventional circuit. Also, op-amp open-loop gain and offset requirements are reduced by a factor of the signal gain.

The proof-of-concept chip was fabricated in 65 nm CMOS technology. Figure 2 shows the output spectra with input frequencies of 12.1MHz at the sampling rate of 250MS/s. SNDR of 67.0dB (10.84-b ENOB) and SFDR of 84.6dB are achieved. The DNL and INL are within $-0.86/+0.52$ LSB and $-0.90/+1.08$ LSB, respectively. The chip operates from a 1.2V power supply and consumes 49.7mW at 250MS/s.



▲ Figure 1: Proposed virtual ground reference buffer technique.



▲ Figure 2: Measured spectrum from $f_{in} = 12.1\text{MHz}$ and $f_s = 250\text{MS/s}$

FURTHER READING

- H. H. Boo, D. S. Boning, and H.-S. Lee, "A 12b 250MS/s Pipelined ADC with Virtual Ground Reference Buffers," in *Proc. International Solid-State Circuits Conference*, 2015, pp. 282-283.
- A. Verma and B. Razavi, "A 10-bit 500MS/s 55-mW CMOS ADC," *IEEE J. Solid-State Circuits*, vol. 44, pp. 3039-3050, Nov., 2009.
- B. Murmann and B. E. Boser, "A 12-bit 75-MS/s pipelined ADC using open-loop residue amplification," *IEEE J. Solid-State Circuits*, vol. 38, pp. 2040-2050, Dec. 2003.

Time-Interleaved A/D Converters

D. P. Kumar, H.-S. Lee

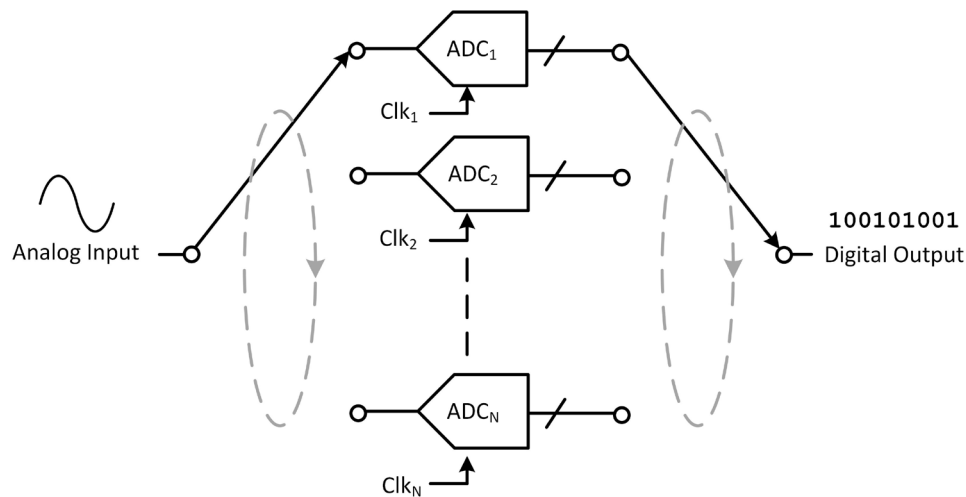
Sponsorship: Masdar Institute of Science and Technology

The demand for high-resolution and high-accuracy A/D converters in communication systems continues to increase. To raise the sampling rates to the GHz range in a power-efficient manner, time-interleaving is an essential technique whereby N A/D channels, each operating at a sampling frequency, f_s , are used to achieve an effective conversion speed of Nf_s , as illustrated in Figure 1.

While time-interleaving enables higher conversion rates in a given technology, mismatch issues such as gain, offset, and sampling clock skew errors between channels degrade the overall A/D performance. Of these issues, sampling clock skew between channels is the biggest problem in high-speed and high-resolution, time-interleaved A/D as errors due to sampling clock skew become more severe for higher input frequencies. A few sources of sampling clock skew between channels exist. Mismatches in the sampling clock path and logic delays are the most obvious. Input signal routing mismatch and RC mismatch of the input sampling circuits also cause sampling clock skew. Previous calibration techniques employ either analog and digital

timing adjustment or digital calibration of output data. The timing adjustment requires an adjustable delay resulting in increased sampling jitter, which cannot be compensated by calibration. The digital calibration of output data requires complex interpolation.

In this research, we are developing a simpler calibration algorithm for sampling clock skew correction whereby the input signal delay is adjusted by controlling the resistance of the input sampling network. The variable time-constant of the input sampling network will result in a linear delay of the input signal if the RC time constant of the input sampling network is much greater than $1/f_{in,max}$, where $f_{in,max}$ is the maximum input signal frequency. This sampling method allows for finely tuned timing-skew corrections, and the impact on noise or power consumption of the system is negligible. A prototype 12-bit, 200MS/s, 4-way time-interleaved A/D implementing the proposed correction technique was taped out in a 65-nm LP CMOS process and is currently being characterized.



▲ Figure 1: Block diagram of a time-interleaved (TI) A/D converter.

FURTHER READING

- N. Kurosawa, H. Kobayashi, K. Maruyama, H. Sugawara, and K. Kobayashi, "Explicit analysis of channel mismatch effects in time-interleaved ADC systems," *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, vol. 48, no. 3, pp. 261-271, Mar. 2001.

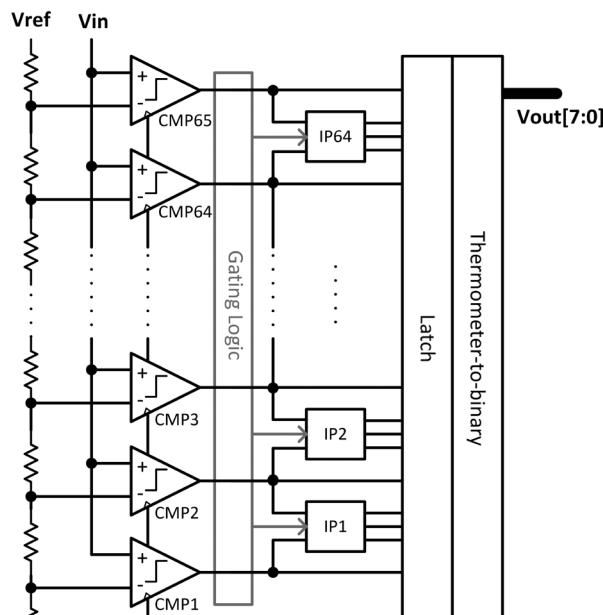
A Flash ADC with a Reduced Number of Comparators

X. Yang, S. Bae, H.-S. Lee

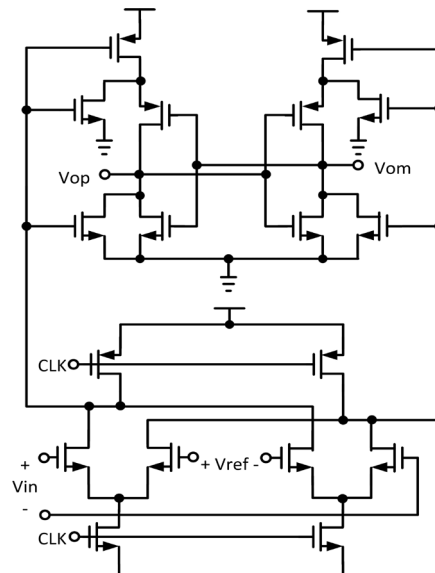
Sponsorship: MIT/MTL GaN Energy Initiative, Office of Naval Research, Samsung

High-speed and low-resolution flash analog-to-digital converters (ADCs) are widely used in applications such as 60-GHz receivers, series links, and high-density disk drive systems, as well as in quantizers in delta-sigma ADCs. In this project, we propose a flash ADC with a reduced number of comparators by means of interpolation. One application for such a flash ADC is a GaN/CMOS hybrid delta-sigma converter. The GaN first stage exploits the high-voltage property of the GaN while the CMOS backend employs high-speed, low-voltage CMOS. This combination may achieve an unprecedented SNR/bandwidth combination by virtue of its high input signal range and high sampling rate. One key component of such an ADC is a flash ADC. To take advantage of the high signal-to-thermal-noise ratio of the proposed system, the quantization noise must be made as small as possible. Therefore, a high-speed, 8-bit flash ADC is proposed for this system. Figure 1 shows the block diagram of the ADC architecture. Sixty-five comparators are used to achieve the 6 most

significant bits (MSBs). Sixty-four interpolators are inserted between the comparators to obtain two extra bits. The input capacitance of this design is only $\frac{1}{4}$ of the conventional 8-bit flash ADC. Therefore a higher operating speed can be achieved. We introduced gating logic so that only one interpolator is enabled during operation, which reduces power consumption significantly. A high-speed, low-power comparator with low noise and low offset requirements is a key building block in the design of a flash ADC. We chose a two-stage dynamic comparator, as in Figure 2, because of its fast operation and low power consumption. With the scaling of CMOS technology, the offset voltage of the comparator keeps increasing due to greater transistor mismatch. A popular offset cancellation technique is to digitally control the output capacitance of the comparator. However, this technique reduces the speed of the comparator because of the extra loading effect. In this project, we also propose a novel offset compensation method that eliminates the speed problem.



▲ Figure 1: Flash ADC architecture, with 65 comparators and 64 2-bit interpolators.



▲ Figure 2: Schematic of the two-stage dynamic comparator.

FURTHER READING

- M. Miyahara, Y. Asada, D. Paik, and A. Matsuzawa, "A Low-Noise Self-Calibrating Dynamic Comparator for High-Speed ADCs," *Solid-State Circuits Conference, 2008. A-SSCC'08. IEEE Asian*, pp. 269-272, Nov. 2008.
- Y.-S. Shu, "A 6b 3GS/s 11mW Fully Dynamic Flash ADC in 40nm CMOS with Reduced Number of Comparators," *VLSI Circuits (VLSIC), 2012 Symposium on*, pp. 26-27, June, 2012.
- M. Miyahara, I. Mano, M. Nakayama, K. Okada, and A. Matsuzawa, "A 2.2GS/s 7b 27.4mW Time-Based Folding-Flash ADC with Resistively Averaged Voltage-to-time Amplifiers," *Solid-State Circuit Conference Digest of Technical Papers (ISSCC), 2014 IEEE International*, pp. 388-389, Feb. 2014.

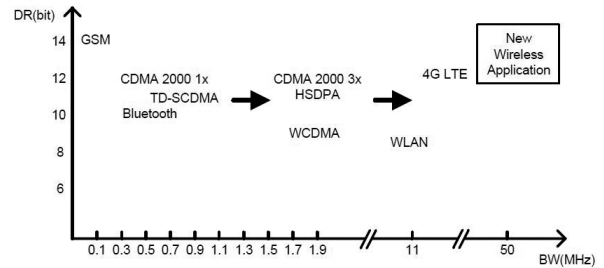
Continuous-Time Delta-Sigma Modulator for Next-Generation Wireless Applications

D.-Y. Yoon, H.-S. Lee

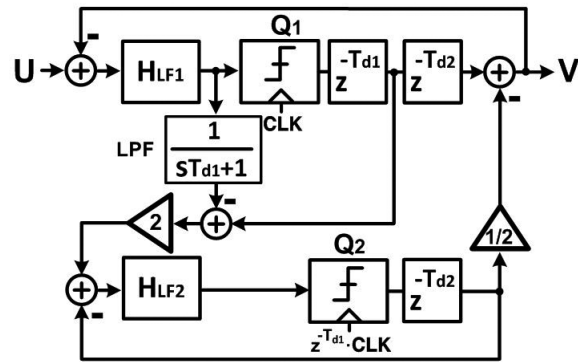
Sponsorship: MediaTek, Inc., Korea Foundation for Advanced Studies

As wireless communication technology is rapidly advancing, new wireless applications are continuously developed. Figure 1 shows each application space and the required dynamic range. The new wireless applications demand wide bandwidth (≥ 50 MHz) and high resolution (≥ 14 bits) data converters. A continuous-time (CT) delta-sigma modulator ($\Delta\Sigma$) is suitable for the demanding new wireless applications due to its high-resolution, wide-bandwidth, and low-power characteristics. In addition, a CT $\Delta\Sigma$ provides an inherent anti-aliasing property. Several state-of-the-art CT $\Delta\Sigma$ s reported recently achieved signal bandwidths greater than the 50 MHz appropriate for the next generation wireless communication. However, the resolution and power consumption still need to be improved.

This project focuses on the design of a CT $\Delta\Sigma$, specifically for modern wireless communication applications. Quantization noise is suppressed aggressively by increasing the effective order of a noise transfer function (NTF) of a $\Delta\Sigma$, instead of increasing the order of a loop filter, to mitigate a stability requirement. To increase the effective order of an NTF, the proposed CT $\Delta\Sigma$ is implemented based on a 2-loop sturdy multi-stage noise-shaping (SMASH) architecture, previously reported in a DT $\Delta\Sigma$. Figure 2 shows the block diagram of the proposed CT SMASH $\Delta\Sigma$. With a wider signal bandwidth, the proposed CT SMASH $\Delta\Sigma$ provides a better quantization noise suppression capability than the original DT SMASH $\Delta\Sigma$ by canceling the quantization noise from the first loop, not simply shaping it. The CT SMASH $\Delta\Sigma$ is implemented with several circuit techniques appropriate for high operation speed. These circuit techniques allow the proposed CT $\Delta\Sigma$ to achieve performance metrics for modern wireless communication applications. As a result, the prototype fabricated in 28 nm CMOS achieves DR of 85 dB, peak SNDR of 74.6 dB, SFDR of 89.3 dBc, and Schreier FOM of 172.9 dB over a 50 MHz bandwidth at a 1.8 GHz sampling frequency.



▲ Figure 1: Dynamic range and bandwidth requirements of ADCs in wireless applications.



▲ Figure 2: Block diagram of a CT SMASH $\Delta\Sigma$ M.

FURTHER READING

- N. Maghari, S. Kwon, and U. Moon, "74 dB SNDR multi-loop sturdy-MASH delta-sigma modulator using 35 dB open-loop Opamp gain," *IEEE J. Solid-State Circuits*, vol. 44, no. 8, pp. 2212-2221, Aug. 2009.
- D.-Y. Yoon, S. Ho, and H.-S. Lee, "15.1 An 85dB-DR 74.6dB-SNDR 50MHz-BW CT MASH $\Delta\Sigma$ modulator in 28nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, 2015, pp. 272-274.

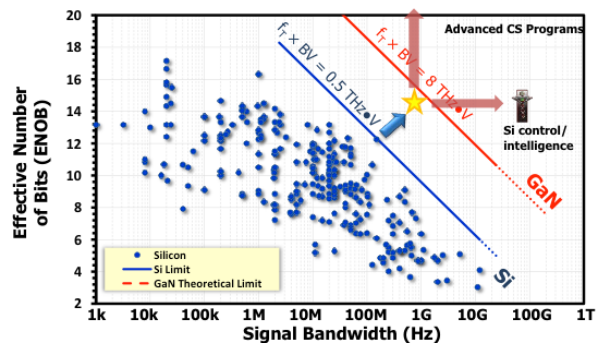
Ultra High-Performance GaN-on-Silicon Analog-to-Digital Converters

S. Chung, X. Yang, H.-S. Lee

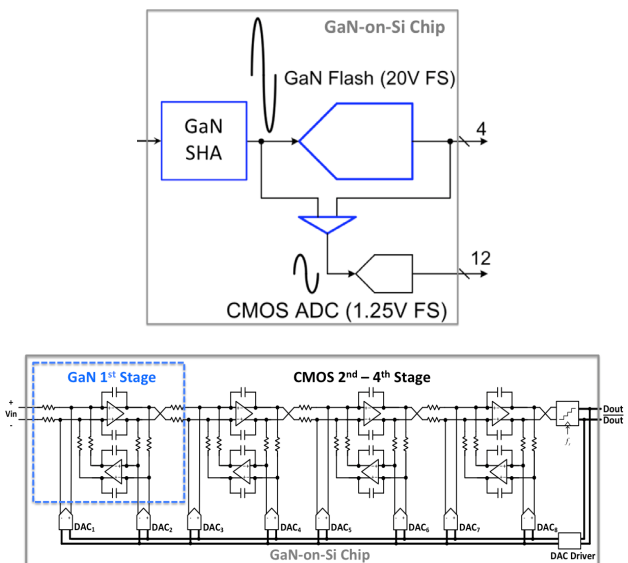
Sponsorship: MIT/MTL GaN Energy Initiative, Office of Naval Research

In this research, we investigate ultra high-performance analog-to-digital converters (ADCs) for diverse emerging applications including personal communication, health care, and an optical backbone network. The low supply voltage of deeply scaled complementary metal oxide semiconductor (CMOS) transistors limits the dynamic range of ADC input signals, thus becoming a fundamental barrier to the performance of silicon ADCs. Recently, high-electron-mobility transistors based on gallium-nitride (GaN HEMTs) are reported with many advantages over the existing compound semiconductor technologies. Operating GaN HEMTs at a very high voltage (30+ V) allows much higher signal-to-noise ratio (SNR) in ADCs than that of CMOS ADCs at given power consumption. A hybrid technology, which monolithically integrates GaN HEMTs with Si-CMOS transistors (GaN-on-Si), will take advantage of both technologies, enabling revolutionary mixed-signal performance (Figure 1). We focus on the design of unprecedentedly high-performance ADCs in a GaN-on-Si hybrid technology (Figure 2). As the first step, we have been developing an over-100-dB SNR GaN sampler for a GaN/Si hybrid pipeline ADC. In addition, we have been investigating the design of a high voltage GaN operational amplifier for a continuous-time delta-sigma ADC with a very high dynamic range.

In the last year, we designed two different GaN track-and-hold (T/H) circuits. The first 250-MS/s T/H was fabricated in a commercial GaN foundry where transistors are optimized for power applications with either low leakage or high gain. These device choices necessitated a two-stage sampler design. The second design operates at a higher speed, 1-GS/s. This circuit was fabricated in a GaN technology that was developed at MTL. This technology offered higher f_t at a lower breakdown voltage. Innovative design techniques allowed a single-stage sampler for a higher sampling rate while avoiding the reverse overvoltage and leakage issues.



▲ Figure 1: Revolutionary ADC performance expected from the monolithic GaN-on-Si integration [Raman 2012].



▲ Figure 2: High-performance ADC architecture in GaN-on-Si hybrid technology: (top) pipeline architecture for wide bandwidth. (bot) continuous-time delta-sigma architecture for high dynamic range.

FURTHER READING

- U. K. Mishra, L. Shen, T. E. Kazior, and Y.-F. Wu, "GaN-Based RF Power Devices and Amplifiers," *Proc. IEEE*, vol. 96, no. 2, pp. 287-305, Feb. 2008.
- J. W. Chung, J.-K. Lee, E. L. Piner, T. Palacios, "Seamless On-Wafer Integration of Si (100) MOSFETs and GaN HEMTs," *IEEE Electron Device Letters*, vol. 30, no. 10, pp. 1015-1017, Oct. 2009.
- S. Raman, "Diverse Accessible Heterogeneous Integration (DAHI) Technology Overview," *Tech. Dig. of DAHI Foundry Tech. Workshop*, p. 15, Aug. 2012.
- S. Chung, H.-S. Lee, "A 200-MS/s 98-dB SNR Track-and-Hold in 0.25-um GaN HEMT," to be presented at *IEEE Custom Integrated Circuits Conference*, San Jose, CA, Sept. 2015.

High-Efficiency, Low-Leakage RF Transmitters for Low Duty Cycle Applications

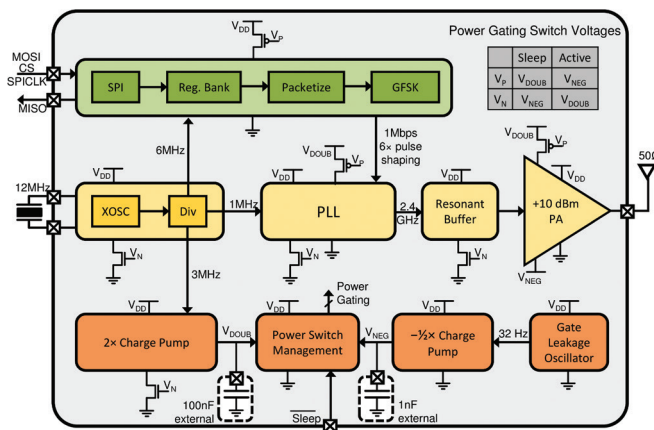
A. Paidimarri, N. Ickes, A. P. Chandrakasan
Sponsorship: Shell, Texas Instruments

With the convergence of ever-improving wireless and energy management technologies, Internet-of-things (IoT) devices for home, industrial, and environmental monitoring have brought significant improvements to lifestyle, safety, and efficiency. Despite these advancements, energy efficiency remains a significant bottleneck, leading to short device lifetimes. This work focuses on RF transmitter optimizations in ultra-low duty cycle applications such as machine vibration monitoring, where time constants of wear and tear are long (hours to days). In this scenario, it is just as important to optimize the leakage power as it is to optimize on-performance.

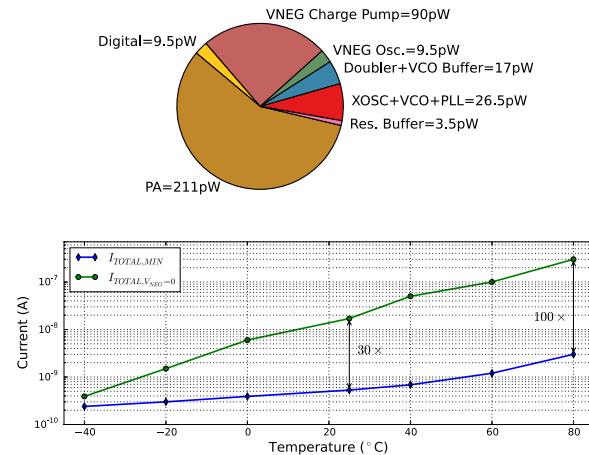
Figure 1 shows the complete block diagram of our Bluetooth low-energy (BLE) transmitter. It operates from a 0.68V supply, where both the switching losses and leakage power are reduced. RF circuits that need a higher voltage for improved efficiency (for example high-Q tunable capacitors) are powered by a voltage doubler that generates 1.2V. In order to have large communication range, the power amplifier is designed

for an output power of +10dBm. A resonant input drive and inductive biasing of the output help achieve high efficiency. The chip also includes a crystal oscillator, PLL, and digital baseband.

Most blocks in the system are power-gated with high- V_T thick-oxide power switches. The on-performance of the switches is improved by turning them on strongly with the 1.2V supply, while the off-leakage is strongly cut off though negative V_{GS} biasing. In order to avoid efficiency penalty due to power switches in the PA, which is the most power-hungry component in the system, the negative biasing is applied to the thin-oxide PA device itself. This achieves significant leakage reduction (up to 100 \times , as shown in Figure 2) without efficiency degradation. The negative voltage of -0.2V is generated by a charge pump and associated oscillator (Figure 1). The design achieves 43.7% system efficiency generating +10.9dBm, and leakage power of 370pW, for an on/off ratio of 7.6×10^7 . This transmitter architecture could enable sub-nW IoT devices.



▲ Figure 1: Block diagram of the BLE transmitter. It shows the digital baseband; RF signal chain including crystal oscillator, PLL, and PA; and leakage management blocks including negative V_{GS} generating charge pump and a 32-Hz oscillator.



▲ Figure 2: Some leakage measurement results. The pie chart shows the distribution of leakage in various blocks. The total leakage is 370pW. The lower graph shows a 30 to 100 \times leakage reduction due to the negative V_{GS} biasing technique.

FURTHER READING

- A. Paidimarri, N. Ickes, and A. Chandrakasan, "A +10dBm 2.4GHz Transmitter with sub-400pW Leakage and 43.7% System Efficiency," *IEEE ISSCC Dig. Tech. Papers*, Feb. 2015, pp. 246-247.
- P. P. Mercier, S. Bandyopadhyay, A. C. Lysaght, K. M. Stankovic, and A. P. Chandrakasan, "A Sub-nW 2.4 GHz Transmitter for Low Data-Rate Sensing Applications," *IEEE J. Solid-State Circuits*, vol. 49, no. 7, pp. 1463-1474, July 2014.
- K. Roy, S. Mukhopadhyay, and H. Mahmoodi-Meimand, "Leakage Current Mechanisms and Leakage Reduction Techniques in Deep-Submicrometer CMOS Circuits," *Proc. of the IEEE*, vol. 91, no. 2, pp. 305-327, Feb. 2003.

Broadband Inter-Chip Link Using Terahertz Wave on Dielectric Waveguide

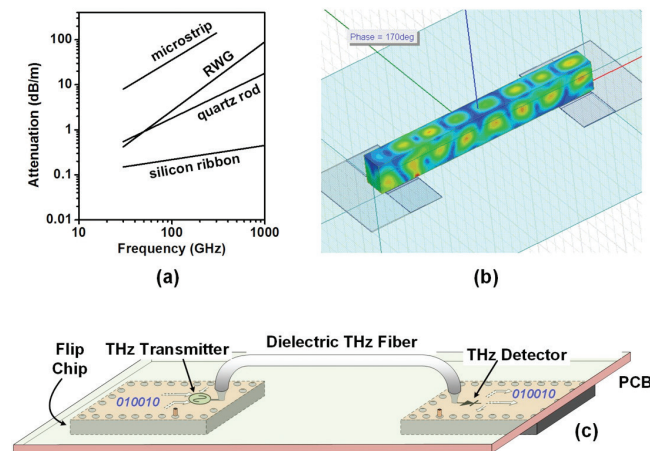
J. Holloway, R. Han

Sponsorship: Office of Naval Research, MIT Lincoln Laboratory

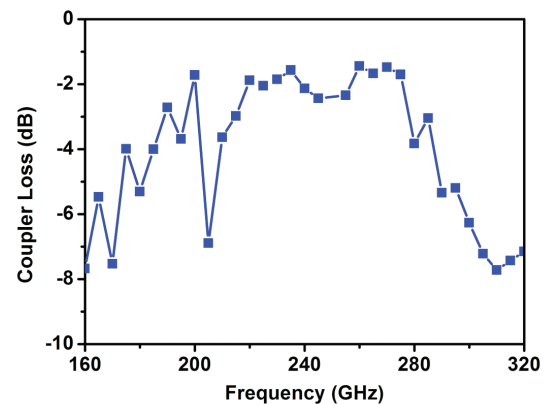
The development of data links between different microchips of an on-board system have encountered a speed bottleneck due to the excessive transmission loss and dispersion of the traditional inter-chip electrical interconnects. Although high-order modulation schemes and sophisticated equalization techniques are normally used to enhance the speed, they also lead to significant power consumption. Silicon photonics provide an alternative path to solve the problem, thanks to the excellent transmission properties of optical fibers. However, the existing solutions are still not fully integrated (e.g., off-chip laser sources are needed) and require process modification to the mainstream CMOS technologies.

In this project, we aim to utilize a modulated Terahertz (THz) wave to transmit broadband data. Similar to the optical link, the wave is confined in dielectric waveguides, with sufficiently low loss (~ 1 dB for 10-cm length) and bandwidth (> 100 GHz) for board level signal transmission (Figure 1). In commercial CMOS/BiCMOS platforms, we have previously demonstrated

high-power THz generation with modulation, frequency up-conversion, and phase-locking capabilities. In addition, a room-temperature Schottky-barrier diode detector (in 130-nm CMOS) with < 10 pW/Hz^{1/2} sensitivity (antenna loss excluded) is also reported. The proposed data link will leverage these blocks in order to achieve > 100 Gbps/waveguide transmission rate with sub pJ/bit energy efficiency. Our current efforts are focused on the design of a chip-to-fiber THz coupler. Different from previous couplers (< 140 GHz) using off-chip antennas, our THz coupler can be implemented entirely with the metal backend of a CMOS process, and does not require any post processing (such as wafer thinning). The structure is also fully shielded, so that there is no THz power leakage into the silicon substrate. Conventional on-chip radiators using ground shield work are resonance type (e.g., patch antenna) and only have $< 5\%$ bandwidth. In comparison, our design is based on a traveling-wave, tapered structure, and exhibits < 3 dB insertion loss across a ~ 60 -GHz bandwidth (Figure 2).



▲ Figure 1: A conceptual illustration of the inter-chip link using THz waveguide (bottom), typical attenuation of various waveguides [1] (top left) and the simulated field of the wave coupled from a chip to a dielectric waveguide ($\epsilon_r = 6$).



▲ Figure 2: The simulated loss of the proposed on-chip coupler design. The bandwidth for < 3 dB loss is ~ 60 GHz.

FURTHER READING

- C. Yeh, F. Shimabukuro, and P. H. Siegel, "Low-Loss Terahertz Ribbon Waveguides," *Applied Optics*, vol. 44, no. 28, pp. 5937-5946, Oct. 2005.
- R. Han, C. Jiang, A. Mostajeran, M. Emadi, H. Aghasi, H. Sherry, A. Cathelin, and E. Afshari, "A 320GHz Phase-Locked Transmitter with 3.3mW Radiated Power and 22.5dBm EIRP for Heterodyne THz Imaging Systems," presented at *IEEE Int. Solid-State Circuit Conf. (ISSCC)*, San Francisco, CA, Feb. 2015.
- R. Han, Y. Zhang, Y. Kim, D. Kim, H. Shichijo, E. Afshari, and K. K. O, "Active Terahertz Imaging Using Schottky Diodes in CMOS: Array and 860-GHz Pixel," *IEEE Journal of Solid-State Circuits (JSSC)*, vol. 48, no. 10, Oct. 2013.

Monolithic GaN-MMIC MEMS-Based Oscillators

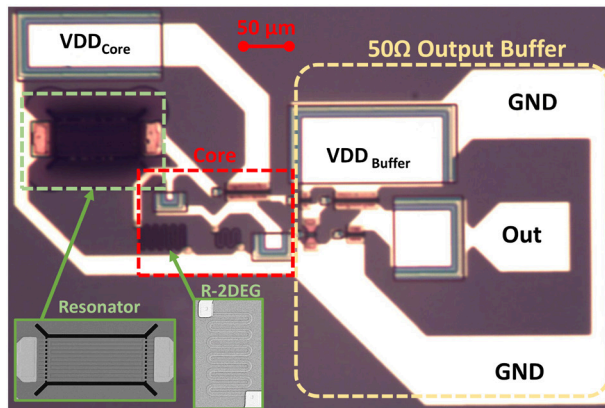
B. Bahr, L.C. Popa, D. Weinstein

Sponsorship: DARPA DAHI Program, DARPA UPSIDE Program

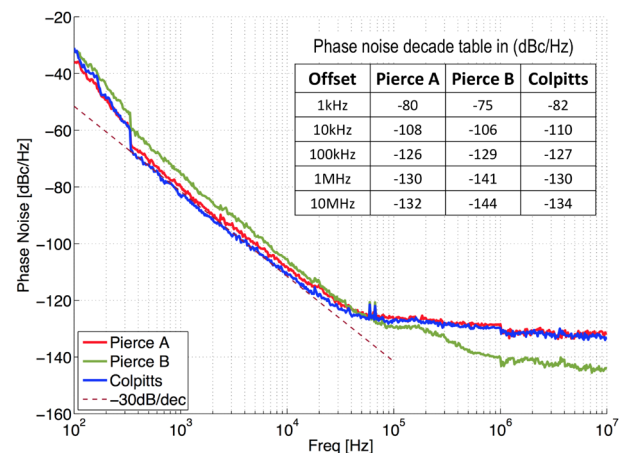
Low phase noise oscillators are critical components in the front end of any communication system. Due to an ever-growing demand for higher data rates and reduced size, weight, and power consumption, efforts to integrate high-Q GHz frequency MEMS resonators with standard circuit technologies have grown dramatically. This work is the first demonstration of a single-chip 1 GHz closed-loop monolithic MEMS-based oscillator circuit implementing both passive and active devices in a standard GaN-MMIC platform, which is a growing technology for various RF front-end applications.

Colpitts and Pierce oscillators were designed and fabricated alongside GaN MEMS Lamb-mode resonators operating at 1 GHz on the same die to realize monolithically integrated GaN-MEMS oscillators. Figure 1 shows an

optical photo of the monolithic Colpitts oscillator, including the MEMS GaN resonator, 2DEG resistors and MIM capacitors incorporated into the standard MMIC process were used for passive components. Fabrication of the oscillator circuit, including the monolithic MEMS resonator and the active HEMT core was carried out in MIT Microsystems Technology Laboratories. Figure 2 shows the measured phase noise of both Colpitts and Pierce oscillators. With power consumption of 1.5 mW, both oscillator topologies out-perform state of the art oscillators of comparable sized devices, a direct result of the significant reduction of parasitics afforded by monolithic integration as well as the high quality factor of the MEMS resonator.



▲ Figure 1: Monolithic Colpitts GaN-MEMS oscillator, showing lamb-mode resonator, 2DEG resistor as well as the core circuit.



▲ Figure 2: Measured phase noise performance of the Pierce and Colpitts oscillators, both operating at 1GHz .

FURTHER READING

- B. W. Bahr, L. C. Popa, and D. Weinstein, "16.8 1GHz GaN-MMIC monolithically integrated MEMS-based oscillators," in *Solid-State Circuits Conference - (ISSCC)*, 2015 IEEE International, 2015, pp. 1-3.
- L. C. Popa and D. Weinstein, "L-Band Lamb mode resonators in Gallium Nitride MMIC technology," in *Proc. IEEE Frequency Control Symposium (FCS 2014)*, 2014, pp. 1-4.
- L. C. Popa and D. Weinstein, "2DEG Electrodes for Piezoelectric Transduction of AlGaIn/GaN MEMS Resonators" in *Proc. European Frequency and Time Forum & International Frequency Control Symposium (EFTF/IFC)*, 2013 Joint, 2013, pp. 922-925.

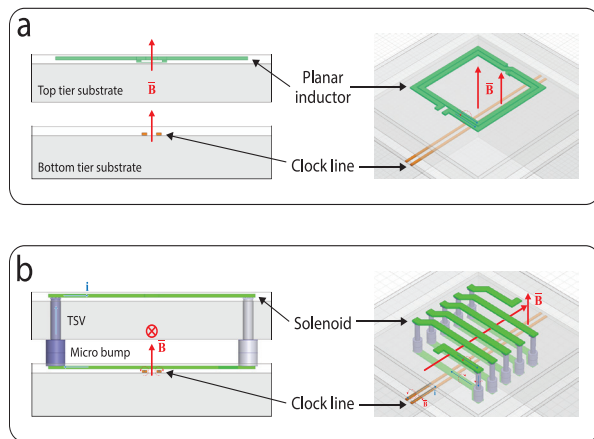
Vertical Noise Coupling Mitigation in 3D-IC Using a Solenoid Inductor

G. Yahalom, A. Wang, A. P. Chandrakasan
Sponsorship: MediaTek

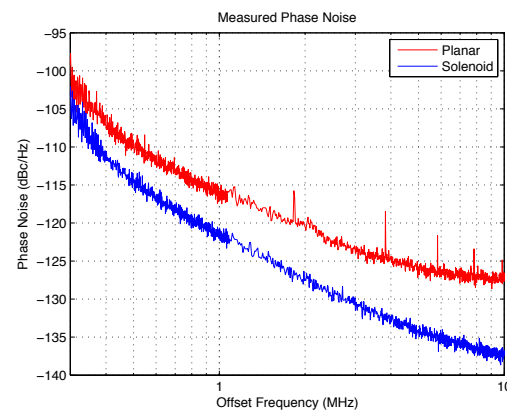
Three-dimensional integrated circuits (3D-IC) have the potential to meet the demand for higher system performance and data rates, while avoiding the increase in cost of scaled CMOS technologies. The ability to stack multiple dies vertically will allow integration of complex systems in a small footprint with short, low-parasitic interconnects. Previous work has shown power and bandwidth benefits of 3D stacking for integrating logic and memory. Designing in three dimensions opens up new possibilities for system- and block-level design; however, care must be taken due to challenges arising from thermal distribution, mechanical stress, power integrity, and signal integrity.

In this work we explore integration of logic devices with RF circuits. Such coexistence may be hindered due to inductive and capacitive coupling between the tiers. Here, a vertical solenoid inductor in 3D-IC is presented to improve the quality factor of the structure and minimize coupling between tiers. The potential coupling between

stacked die tiers is shown in Figure 1 for different inductor structures. The bottom tier die contains signal lines that emulate part of a high-speed digital clock tree. Directly above the clock lines, two different integrated inductor structures were fabricated: the proposed solenoid and a conventional planar structure as a reference design. The planar inductor utilizes the top metal layer and has a patterned ground shield. The solenoid inductor uses the through silicon vias themselves as part of the inductor structure and the redistribution layers on both the top and bottom die tiers. The measured phase noise of two voltage-controlled oscillators using these inductors is plotted in Figure 2. The higher inductance and lower resistance of the solenoid results in a ~70% higher quality factor and lower phase noise. Furthermore, the planar inductor exhibits spurs due to an adjacent low-frequency digital clock, whereas the solenoid does not exhibit spikes due to the clock.



▲ Figure 1: Illustration of 3D-IC stack with clock lines, (a) planar and (b) solenoid inductor structures, along with current directions and magnetic field lines (figure not to scale).



▲ Figure 2: Phase noise measurement of planar (red) and solenoid (blue) inductor-based voltage-controlled oscillators with the presence of an on-chip low speed digital clock.

FURTHER READING

- G. Yahalom, A. Wang, U. Ko and A. Chandrakasan., "A Vertical Solenoid Inductor for Noise Coupling Minimization in 3D-IC," presented at *IEEE Radio Frequency Integrated Circuits Symp.*, Phoenix, AZ, May 2015.
- W. Arden, M. Brillouët, P. Cogez, M. Graef, B. Huizing, and R. Mahnkopf "More-than-Moore white paper," *Version*, vol. 2, p. 14, 2010.
- G. Van der Plas, P. Limaye, I. Loi, A. Mercha, H. Oprins, C. Torregiani, S. Thijs, D. Linten, et al., "Design Issues and Considerations for Low-Cost 3-D TSV IC Technology," *IEEE Journal of Solid-State Circuits*, vol. 46, no. 1, pp. 293–307, 2011.

Energy and Area-Efficient Hardware Implementation of HEVC Inverse Transform and Dequantization

M. Tikekar, V. Sze, A. P. Chandrakasan

Sponsorship: Texas Instruments, National Science Foundation

High Efficiency Video Coding (HEVC) achieves a 50% reduction in bit-rate over Advanced Video Coding (H.264/AVC) at the same visual quality. A key feature of HEVC is the use of large 16×16 and 32×32 inverse discrete cosine transforms (IDCTs), a new 4×4 inverse discrete sine transform, and high-precision 4×4 and 8×8 IDCTs. However, this new feature raises several challenges for hardware implementations which we have addressed in our work.

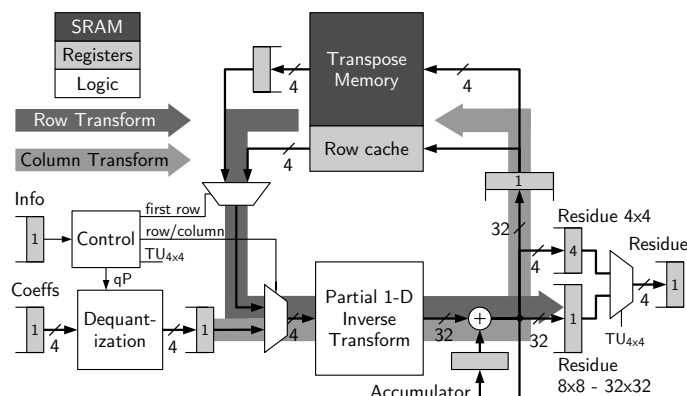
Designing for all the transform sizes (4×4 - 32×32) requires complex control as they need different cycle counts. We developed a pipelining scheme to process all sizes with at least 2 pixels/cycle throughput. Further, we use zero-column skipping to increase throughput by 63%. Zero-column skipping also decreases the number of signal transitions, which reduces energy by 30%.

The large and high-precision transforms have $8 \times$ higher computational complexity, which affects both area and energy. To address the area problem, a Multiple Constant Multiplication-based method had previously

been proposed. We improved the energy efficiency of the design by 17% using data-gating at the cost of 4% area.

The HEVC inverse transform needs a 16 kbit transpose memory as compared to 1 kbit for H.264/AVC. The large transpose memory needs to use static random access memory (SRAM), which is denser but slower and less flexible than registers that are used for H.264/AVC. We used a combination of 4 single-port SRAMs and a small register cache to achieve the desired throughput in a small area.

We implemented the inverse transform and dequantization engine in TSMC 40 nm CMOS. The design, shown in Figure 1, has an area of 126 k gates and energy consumption of 11.9 pJ/pixel. Energy consumption depends on the data being processed. We observed that large transforms typically contain more zeros. Our zero-column skipping method exploits this observation to process the large transforms with better energy efficiency as seen in Table 1.



▲ Figure 1: Architecture of inverse transform and dequantization engine.

Transform size	Fraction of zeros in data	Energy (pJ/pixel)
4×4	85.0%	9.7
8×8	95.5%	11.4
16×16	97.6%	12.6
32×32	99.5%	12.2

▲ Table 1: Exploiting the higher fraction of zeros in large transform to improve energy efficiency.

FURTHER READING

- M. Tikekar, C.-T. Huang, V. Sze, and A. Chandrakasan "Energy and Area-Efficient Hardware Implementation of HEVC Inverse Transform and Dequantization," *Technical Digest of the 2014 IEEE International Conference on Image Processing*, Oct. 2014, p. 2100, 2104.
- M. Budagavi, A. Fuldseth, G. Bjontegaard, V. Sze, and M. Sadafale, "Core Transform Design for the High Efficiency Video Coding (HEVC) Standard," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 6, pp. 1029–1041, 2013.
- M. Tikekar, C.-T. Huang, C. Juvekar, V. Sze, and A. P. Chandrakasan, "A 249-Mpixel/s HEVC Video-Decoder Chip for 4K Ultra-HD Applications," *IEEE Journal of Solid-State Circuits*, vol. 49, no. 1, pp. 61–72, Jan. 2014.

A Navigation Device with 3-D Computer Vision Processor for Visually Impaired People

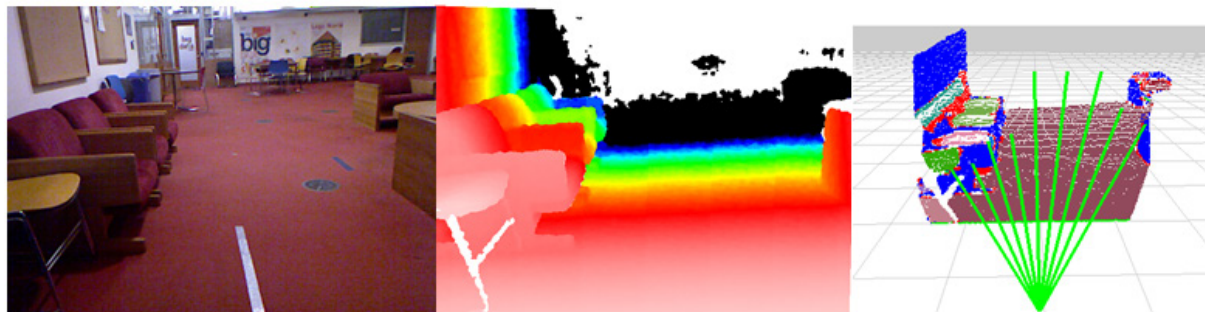
D. Jeon, N. Ickes, P. Reina, I. Ananthabhotla, H. C. Wang, D. L. Rus, A. P. Chandrakasan
Sponsorship: Andrea Bocelli Foundation, Texas Instruments

Computer vision is an emerging area for mobile systems, and there is an increasing need for various applications such as autonomous navigation, gesture recognition, and face identification. In this project, we seek to build a 3-D vision processor targeting multiple applications that achieves good reconfigurability and energy efficiency. We first take a navigation device for the visually impaired as a primary application since it requires various processing algorithms such as surface normal calculation and region growing. The proposed processor will be demonstrated as a core processing part of the entire navigation system.

The portable navigation system processes depth image data captured by a Time-of-Flight camera to detect obstacles in front and inform the user of safe areas to walk (Figure 1). The system must operate for at least few hours relying on a small battery, making the power consumption a key design constraint. We are implementing a low-power computer vision

processor by exploring multiple design spaces ranging from algorithmic optimization and architecture improvement to maximize energy efficiency. We also focus on developing algorithms to save LED illumination power since it can easily outweigh energy savings from the processor. The algorithm will dynamically tune configurations of the ToF camera such as illumination power and refresh rate based on sensory inputs.

Another important direction of our research is miniaturization of the navigation system. The navigation device must be in a small form factor so that it is not easily noticeable for social reasons. We are developing a customized system consisting of an image sensor and peripheral ICs of a ToF camera, FPGA chip and the proposed ASIC vision processor. The system will directly process data captured by on-board ToF camera and detect safe walkable areas. This information will be transferred via Bluetooth to an external haptic array, which gives vibratory feedback to the user.



▲ Figure 1: Original color (left) and depth image (center) captured by RGB-D camera. Processed depth image reveals possible obstacles with different colors and provides walkable distance in different directions (right).

FURTHER READING

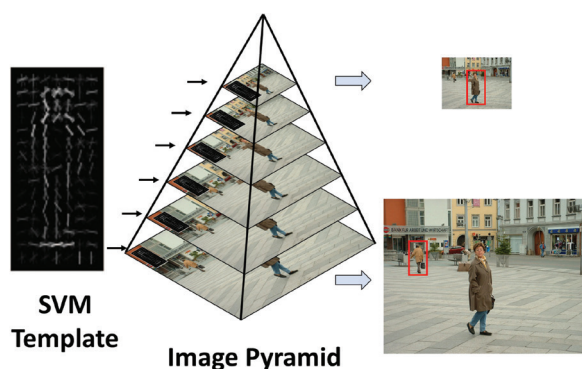
- S. Holzer, R. B. Rusu, M. Dixon, S. Gedikli, and N. Navab, "Adaptive Neighborhood Selection for Real-Time Surface Normal Estimation from Organized Point Cloud Data Using Integral Images," in *Proc. 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012, pp. 2684-2689.
- R. Velasquez, "Wearable assistive device for the blind," *Wearable and Autonomous Biomedical Devices and Systems for Smart Environment*. New York: Springer, 2010, pp. 331-349.
- S. Bhatlawande, A. Sunkari, M. Mahadevappa, J. Mukhopadhyay, M. Biswas, D. Das, and S. Gupta, "Electronic bracelet and vision-enabled waist-belt for mobility of visually impaired people," *Assistive Technology*, vol. 26, no. 4, pp. 186-195, 2014.

Energy-Efficient Hardware for Multi-Modal Object Detection

A. Suleiman, Z. Zhang, V. Sze

Sponsorship: Texas Instruments, DARPA

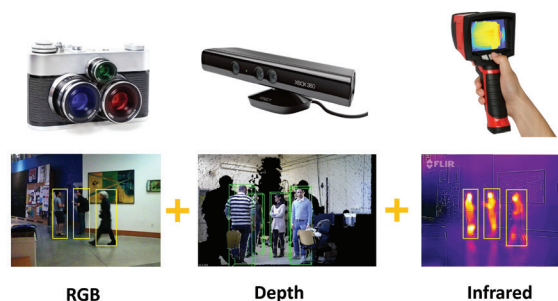
Object detection is needed in many embedded applications, such as surveillance, advanced driver assistance systems (ADAS), consumer electronics, and robotics. Real-time and high throughput are necessary for applications such as ADAS and unmanned aircraft vehicles (UAV) to allow more time for course corrections in case of quick changes in the environment. On the other hand, high-resolution images enable early detection by having enough pixels to identify objects at a distance. Finally, in both navigation and portable devices, energy-efficient object detection is desirable because of the energy-limited battery. In this project, a dedicated application-specific integrated circuit (ASIC) is implemented for histogram of oriented gradients (HOG)-based object detection. Unlike most existing low-power hardware object detectors, this architecture supports multi-scale detection for robustness as shown in Figure 1. The image pyramid is generated on the fly, resulting in 3x more pixels to process. HOG features are then extracted for each scale and classified with a support vector machine (SVM). This detection system



▲ Figure 1: SVM template for pedestrian detection is shown. The template is used to process each level of the image pyramid, hence detecting pedestrians with different sizes/distances in the scene.

can process 1080 HD videos at 60 fps, supporting 12 scales per frame, and consuming only 45.3 mW using 45-nm silicon-on-insulator (SOI) complementary metal-oxide-semiconductor (CMOS) technology.

Our next step is to leverage multi-modality to further boost the detection accuracy. These additional sources, shown in Figure 2, complement the visual red-green-blue (RGB) data and hence enable a vision system to function under challenging but important environments like dark scenes. It has been demonstrated in the literature that these sensors can significantly boost the performance of object detection and recognition systems. However, unlike RGB cameras, the output of these sensors is noisy, low in resolution, and sometimes incomplete. Therefore, it is important to properly process and fuse them with RGB data. We are investigating solutions to these challenges from both algorithmic and hardware perspectives. In the end, we expect to deliver a robust object detection system working on multiple sources in different environments while still delivering energy-efficient, real-time, and high-throughput processing.



▲ Figure 2: RGB, depth, and infrared images captured from different sensors for pedestrian detection. By fusing them correctly, we can enhance the robustness and performance of our detector.

FURTHER READING

- A. Suleiman and V. Sze, "Energy-Efficient HOG-based Object Detection at 1080HD 60 fps with Multi-Scale Support," in *Proc. IEEE Workshop on Signal Processing Systems*, pp. 20-22, Oct. 2014.
- S. Gupta et al., "Learning Rich Features from RGB-D Images for Object Detection and Segmentation," in *Proc. European Conference on Computer Vision*, vol. 8695, pp. 345-360, Sept. 2014.
- Y. Fang et al., "A shape-independent method for pedestrian detection with far-infrared images," *IEEE Transaction on Vehicular Technology*, vol. 53, pp. 1679-1697, Nov. 2004.

Towards High-Performance Bufferless NoCs with SCEPTER

B. K. Daya, L. S. Peh, A. P. Chandrakasan
Sponsorship: Center for Future Architectures

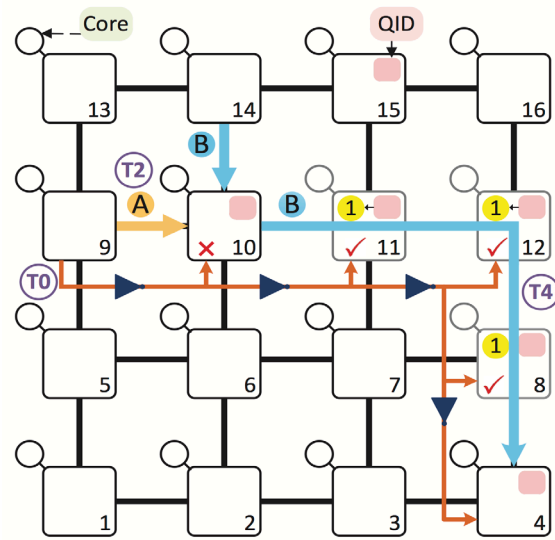
In the many-core era, the network on-chip (NoC) is playing a larger role in meeting performance, area, and power goals. However, the network, especially the buffers, consumes a significant portion of the total power consumption: The MIT SCORPIO NoC connects 36 tiles and expends 18% of the tile power; the Intel TeraFLOPS chip network connects 80 tiles and expends 30% of the total power. Proposals have advocated bufferless NoCs to reduce the NoC power consumption; however, a performance wall has been reached so that high throughput performance has not been extracted.

We present SCEPTER (Single-Cycle Express Path Traversal for Efficient Routing), an NoC architecture that pushes towards high-performance bufferless NoCs. We lower the average network latency of bufferless NoCs by leveraging single-cycle multihop traversals across the network. Thus, even if the flit is sent in a non-preferred direction, a single-cycle path can be potentially traversed, bringing the flit closer to its destination even along non-minimal routes. SCEPTER intelligently prioritizes between flits in the router pipeline, bypassing from faraway, and waiting in the network interface to be injected. It adaptively routes flits in a livelock-free manner while maximizing opportunities to zoom along virtual express paths by opportunistically bypassing. Figure 1 displays an example where a multihop bypass path is preset by Flit A, with the use of switch setup requests (SSRs). Since Flit A is deflected along another direction at node 10, Flit B can either follow the preset bypass path or not. Figure 2 shows the router pipeline and bypass path. When Flit B arrives, a check is performed on whether Flit B's destination quadrant ID (QID) matches that of the preset bypass path. In this example, it matches and Flit B is able to traverse the path in one cycle to reach node 4.

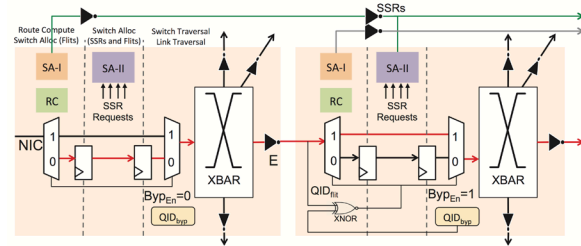
For a 64-node network, we demonstrate an average 62% reduction in latency and an average 1.3x higher throughput over a baseline bufferless NoC for synthetic traffic patterns, with comparable performance to a single-cycle multihop buffered mesh network with 6 flit buffers, per input port, in each router. Early post-synthesis results in IBM 32-nm SOI technology show an average 31% lower area and 33% lower power than optimized buffered router baselines.

FURTHER READING

- B. K. Daya, C. H. O. Chen, S. Subramanian, W. C. Kwon, S. Park, T. Krishna, J. Holt, A. P. Chandrakasan, and L. S. Peh, "SCORPIO: A 36-Core Research Chip Demonstrating Snoopy Coherence on a Scalable Mesh NoC with In-Network Ordering," in *Proc. 41st International Symposium on Computer Architecture* 2014, pp. 25-36.
- T. Krishna, C.-H. Chen, W. C. Kwon, and L.-S. Peh, "Breaking the on-chip latency barrier using SMART," in *IEEE 19th International Symposium on High Performance Computer Architecture*, 2013, pp. 378-389.
- B. K. Daya, L.-S. Peh, and A. P. Chandrakasan, "Towards High-Throughput Bufferless NoCs with SCEPTER," *IEEE Computer Architecture Letters*, 2015, vol. PP, no.99, p. 1, May 1, 2015.



▲ Figure 1: The SCEPTER single-cycle multihop traversals allow multiple hops to be traversed by presetting the crossbar switched. At time T0, Flit A's SSR is sent along a path to the destination node. At time T2, at node 10, Flit A and Flit B contend for output port East. Flit B wins; at time T4 it arrives at node 11 and traverses the preset path to node 4.



▲ Figure 2: The SCEPTER router uses the latches between the pipeline stages to hold the flit's information, as no buffers exist. Thus, even if a flit cannot be routed to its desired output port, it must be deflected along another direction/output port. When a flit arrives, a simple check is performed to see if the bypass path is enabled for this cycle. Subsequently, a check determines whether this flit's destination quadrant ID matches that of the SSR that preset the bypass path. If it matches, the flit skips the first two pipeline stages and traverses the crossbar and link.

Energy-Efficient SRAM using Data-Dependency

C. Duan, A. P. Chandrakasan

Sponsorship: DARPA, National Science Foundation

Embedded static random access memories (SRAMs) are critical components in the design of modern system-on-chips (SoCs). As the capabilities of many digital electronic devices continue to improve, the need for both large and low-power on-chip storage grows in parallel. In IC implementations for various applications, SRAMs occupy a disproportionate amount of total die area and total power consumption. Despite recent progress enabling the low-power operation of digital system blocks at low supply voltages, robust SRAM operation still requires a high operating voltage to guarantee reliability in worst-case scenarios including extreme process, temperature, and/or voltage conditions. Consequently, SRAM has recently become the current bottleneck for further power reduction in many systems and thus necessitates creative energy- and area-efficient solutions. To develop novel approaches to SRAM energy savings,

application-specific data features are to be explored in harmony with state-of-the-art techniques such as voltage scaling. Highly correlated data, introduced intuitively as data with repeated or similar values, has been shown via proof of concept to help memory make predictions, reduce bit-line switching, and ultimately save energy wasted on reading redundant and/or predictable information. Existing data-dependent designs, however, are limited by their narrow applicability as well as significant overhead in area and/or latency. In this work, we propose to address said limitations by first examining inherent data-accessing features of several targeted applications. New bit-cell and architectural-level techniques will be investigated to support low-power SRAM operation. The remaining work focuses on the circuit-level implementation of these developments and the construction of a test chip in 28 nm FD-SOI process.

FURTHER READING

- M. Sinangil and A. P. Chandrakasan, "Application-specific SRAM design using output prediction to reduce bit-line switching activity and statistically gated sense amplifiers for up to 1.9x lower energy/access," *IEEE Journal of Solid-State Circuits*, vol. 49, no. 1, pp. 107–117, Jan. 2014.

Energy-Efficient SRAM Design in 28-nm FDSOI Technology

A. Biswas, A. P. Chandrakasan
Sponsorship: STMicroelectronics, DARPA

As CMOS scaling continues to the sub-32 nanometer regime, the effects of device variations become more prominent. This is very critical in static random access memories (SRAMs), which use very small transistor dimensions to achieve high memory density. The conventional six-transistor (6T)-based SRAM bit-cell, which provides the smallest cell-area, fails to operate at lower supply voltages (V_{dd}). This failure is due to the significant degradation of functional margins as the supply voltage is scaled down. However, V_{dd} scaling is crucial in reducing the energy consumption of SRAMs, which is a significant portion of the overall energy consumption in modern micro-processors. Energy saving in SRAM is particularly important for battery-operated applications, which run from a very constrained power-budget.

This work focuses on energy-efficient 6T SRAM design in a 28-nm Fully Depleted Silicon-On-Insulator (FDSOI) technology. Significant savings in energy per

access of the SRAM can be achieved by V_{dd} scaling. Different read and write assist techniques are evaluated to improve the minimum SRAM operating voltage ($V_{dd,min}$). The different techniques are compared based on various metrics, e.g., energy-overhead, area-overhead, etc. Techniques are proposed to reduce the energy-overhead of different assist methods.

It has been recently shown that correlation in data can be exploited to reduce energy consumption of SRAMs. Applications such as motion estimation in video processing access the same data for multiple read cycles before writing a new data. Sinangil et al. proposed a 10T bit-cell that used data prediction to reduce bit-line switching energy. However, it incurs a significant area overhead due to 10 transistors required to implement the bit-cell. In this work we investigate techniques to incorporate data-prediction in 6T-based SRAM design to benefit from higher density while still saving energy.

FURTHER READING

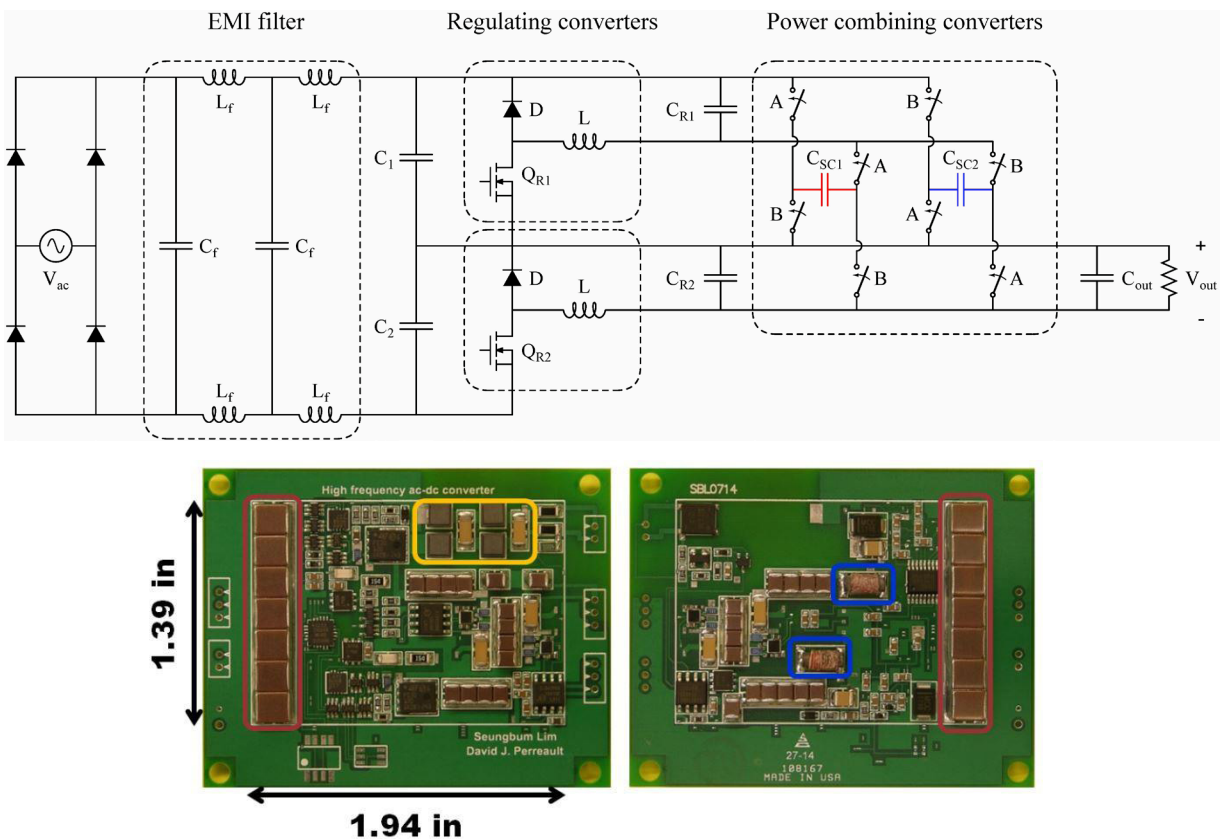
- N. Planes, et. al., "28nm FDSOI technology platform for high-speed low-voltage digital applications," *2012 Symposium on VLSI Technology (VLSIT)*, June 2012, pp. 133-134.
- B. Zimmer, S. O. Toh, H. Vo, Y. Lee, O. Thomas, K. Asanovic, and B. Nikolic, "SRAM Assist Techniques for Operation in a Wide Voltage Range in 28-nm CMOS," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 59, no. 12, pp. 853-857, Dec. 2012.
- M. E. Sinangil and A. P. Chandrakasan, "Application-Specific SRAM Design Using Output Prediction to Reduce Bit-Line Switching Activity and Statistically Gated Sense Amplifiers for Up to 1.9x Lower Energy/Access," *IEEE Journal of Solid-State Circuits*, vol. 49, no. 1, pp. 107-117, Jan. 2014.

New AC-DC Power Factor Correction Architecture Suitable for High Frequency Operation

S. Lim, D.M. Otten, D. J. Perreault
Sponsorship: ARPA-E, Texas Instruments

We present a novel ac-dc power factor correction (PFC) power conversion architecture for a single-phase grid interface. The proposed architecture has significant advantages for achieving high efficiency, good power factor, and converter miniaturization, especially in low-to-medium power applications. The architecture enables twice-line frequency energy to be buffered at high voltage with a large voltage swing, enabling reduction in the energy buffer capacitor size, and elimination of electrolytic capacitors. While this architecture can be beneficial with a variety of converter topologies, it is especially suited for system miniaturization by enabling designs that operate at high frequency (HF, 3–30

MHz). Moreover, we introduce circuit implementations that provide efficient operation in this range. The proposed approach is demonstrated for an LED driver converter operating at a (variable) HF switching frequency (3–10 MHz) from 120Vac, and supplying a 35 Vdc output at up to 30 W. The prototype converter achieves high efficiency (92 %) and power factor (0.89) and maintains good performance over a wide load range. Owing to architecture and HF operation, the prototype achieves a high “box” power density of 50W/ in³ (“displacement” power density of 130W/ in³), with miniaturized inductors, ceramic energy buffer capacitors, and a small-volume EMI filter.



▲ Figure 1: The prototype converter, implemented on a 1.94 in(x), 1.39 in(y), 0.22 in(z) printed circuit board. This figure shows the front and back side of the PCB.

FURTHER READING

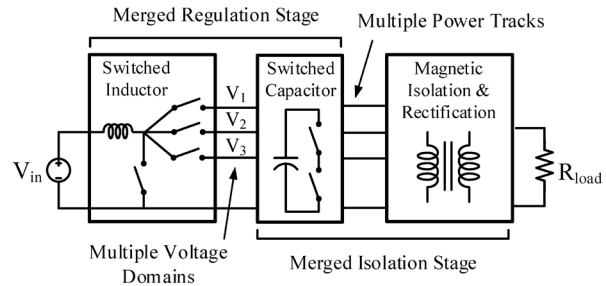
- S. Lim, D. M. Otten, and D. J. Perreault, “Power conversion architecture for grid interface at high switching frequency,” in *Applied Power Electronics Conference and Exposition (APEC), 2014 Twenty-Ninth Annual IEEE*, 2014, pp.1838-1845.

A High-Power-Density Wide-Input-Voltage-Range Isolated DC-DC Converter with a Multi-Track Architecture

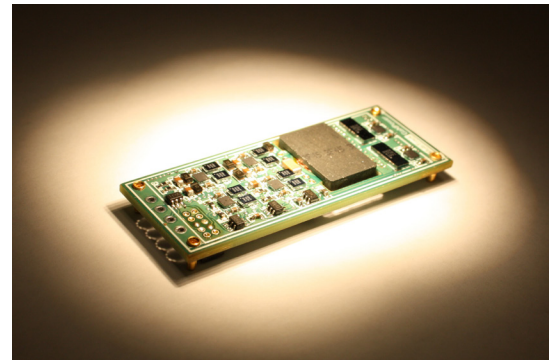
M. Chen, K. K. Afridi, S. Chakraborty, D. J. Perreault
Sponsorship: Texas Instruments, CICS

This project investigates a multi-track power conversion architecture that splits charge into multiple voltage domains and delivers power through multiple tracks, as illustrated in Figure 1. The multi-track architecture reduces the voltage ratings on devices, reduces the voltage regulation stress of the system, improves the component utilization, and reduces the sizes of passive components. The architecture also leverages the complementary strengths of switched-inductor, switched capacitor, and magnetic isolation circuits and gains mutual benefits from the way they are merged. This architecture is suitable to applications that require both isolation and wide-input-voltage range. Compared to a conventional two-stage design, its regulation stage and isolation stage are merged, leading to a hybrid-switched capacitor- magnetics structure that reduces the energy that is “reprocessed” by the two-stages. An 18V-80V_{in}, 5V_{out}, 15A, 800 kHz, isolated dc-dc converter has been built and tested to verify the effectiveness of this architecture. It has a power density of 453 W/inch³ and a peak efficiency of 91.3%. This power density is 3x higher than the power density of the state-of-the-art commercial converters. A picture of the prototype converter is shown in Figure 2.

The proposed multi-track power conversion architecture is one embodiment of a group of generalized distributed power conversion techniques. Comparison of the multi-track power conversion architecture with conventional centralized power conversion architecture can demonstrate and theoretically quantify the advantages of distributed power conversion.



▲ Figure 1: The proposed multi-track power conversion architecture comprising multiple voltage domains and multiple power tracks. Its regulation and isolation stage are merged to reduce the energy that is “re-processed” by the two-stages.



▲ Figure 2: A picture of the high-power-density wide-input-voltage-range isolated multi-track converter.

FURTHER READING

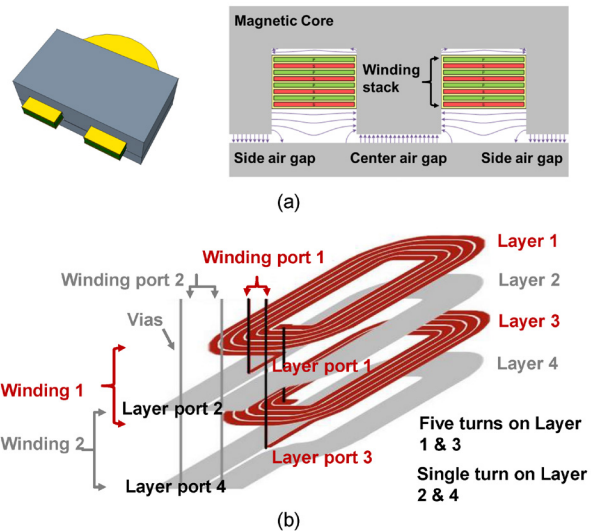
- M. Chen, K. K. Afridi, S. Chakraborty and D. J. Perreault, “A High-Power-Density Wide-Input-Voltage-Range Isolated Dc-Dc Converter having a Multi-Track Architecture,” to be presented at the *IEEE Energy Conversion Congress and Exhibition (ECCE)*, Montreal, 2015.
- M. Chen, M. Araghchini, K. K. Afridi, J. H. Lang, C. R. Sullivan, and D. J. Perreault, “A Systematic Approach to Modeling Impedances and Current Distribution in Planar Magnetics,” *IEEE Transactions on Power Electronics*, vol. PP, no. 99, pp. 1-20, Jan. 2015.

A Systematic Approach to Modeling Impedances and Current Distribution in Planar Magnetics

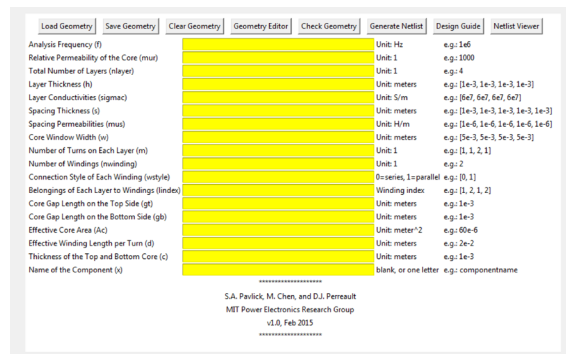
M. Chen, S. A. Pavlick, M. Araghchini, K. K. Afridi, J. H. Lang, C. R. Sullivan, D. J. Perreault
Sponsorship: Texas Instruments, CICS

Planar magnetic components using printed-circuit board (PCB) windings are attractive due to their high repeatability, good thermal performance and usefulness for realizing intricate winding patterns. An example planar magnetic structure is shown in Figure 1. To enable higher system integration at high switching frequency, more sophisticated methods that can rapidly and accurately model planar magnetics are needed. This project develops a systematic approach to modeling impedances and current distribution in planar magnetics based on a lumped circuit model named the *Modular Layer Model* (MLM). Stacked PCB layers are modeled as repeating modular impedance networks, with additional modular impedances representing the magnetic core, air gaps and vias. The model captures skin and proximity effects and enables accurate predictions of impedances, losses, stored reactive energy, and current sharing among windings. The MLM can be used to simulate circuits incorporating planar magnetics, to visualize the electromagnetic fields, and to extract parameters for magnetic models by simulations, among many other applications. The modeling results are checked with results of previous theories and finite-element-modeling approaches, with good matching presented. A group of planar magnetic devices, including transformers and inductors with various winding patterns, are prototyped and measured to validate the proposed approach and clarify the boundaries of its applicability.

A software that can generate SPICE netlists based on planar magnetics geometry information has been developed and is accessible by emailing the authors. The user interface of the current version of the software is shown in Figure 2.



▲ Figure 1: (a) Cross-sectional view of a planar magnetic structure. It comprises a magnetic core, a winding stack, and a set of possible air gaps. (b) Winding stack of an example two-winding, four-layer transformer with 10:1 turns ratio.



▲ Figure 2: A magnetic geometry to SPICE netlists conversion tool – M2Spice.

FURTHER READING

- M. Chen, M. Araghchini, K. K. Afridi, J. H. Lang, C. R. Sullivan, and D. J. Perreault, "A Systematic Approach to Modeling Impedances and Current Distribution in Planar Magnetics," *IEEE Transactions on Power Electronics*, vol. PP, no. 99, pp. 1-20, Jan. 2015.

Investigating Magnetic Materials for Power Conversion at High Frequency

A. J. Hanson, J. Belk, C. R. Sullivan, D. J. Perreault
Sponsorship: Lockheed Martin

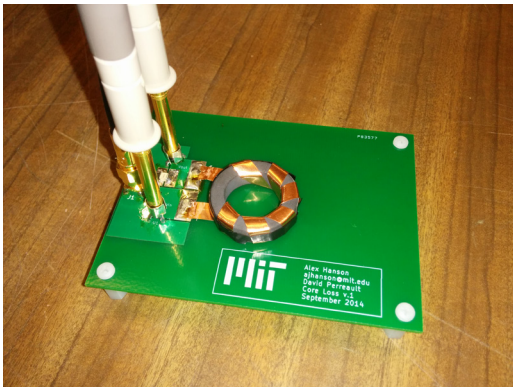
Magnetic components (inductors and transformers) are typically the largest and most lossy components in power converters. While increasing converter switching frequencies can reduce the required size of passive components, size reductions achievable through frequency increases are often limited by magnetic material constraints for components where core loss is a major consideration, such as in transformers and resonant inductors. Nevertheless, recent research has made significant advances in miniaturized power electronics operating in the high (HF) and very high frequency (VHF) ranges (3-300 MHz), well above typical modern designs operating from hundreds of kilohertz to a few megahertz.

While such advances have been substantial, the design of power magnetics is still not fully understood or optimized, especially in the HF range (3-30 MHz) where use of low-permeability RF magnetic materials can play a valuable role. In part, this lack of understanding owes to a lack of data regarding HF magnetic materials. Magnetic materials are typically characterized for power loss; however, such data are simply not available for most magnetic materials above a few megahertz, hindering the design of magnetics at these frequencies.

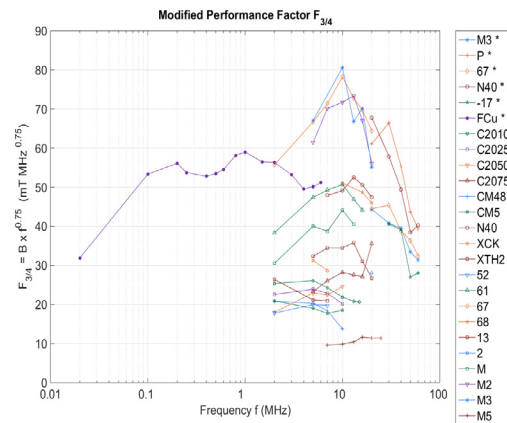
Additionally, the modeling and evaluation methods for magnetic components must be adapted in the HF range.

Great efforts have been made to model magnetic core loss and winding loss, but holistic design and evaluation remain incomplete. For example, the commonly used Performance Factor (a FOM for magnetic materials) assumes that, in a magnetic component to be designed, winding loss is not a function of frequency, which often is not true in the HF range.

Our research is extending both the empirical data and the evaluation methods necessary for magnetics design in power converters operating above a few MHz. Using HF measurement techniques (Figure 1), we have gathered large-signal core loss data for a variety of commercially available materials; these data show room for significant improvements in power conversion by moving to the HF range. We have also developed an extension to the popular Performance Factor to include HF effects (Figure 2), allowing material comparison even at frequencies where the traditional Performance Factor is inapplicable. Our results suggest that significant improvements in performance are possible through operation at HF using commercially available magnetic materials. Performance Factor and Modified Performance Factor facilitate understanding of the implications of these data for selecting an operating frequency and understanding its benefits.



▲ Figure 1: Test setup with a representative magnetic core-under-test (CUT). Core loss measurements at HF require advanced techniques compared to similar characterization in the typical ~100 kHz range.



▲ Figure 2: Plot of Modified Performance Factor for materials characterized in this investigation, compared with materials designed for lower frequencies. The theoretical modification correctly accounts for HF effects on loss.

FURTHER READING

- D. Perreault, J. Hu, J. Rivas, Y. Han, O. Leitermann, R. Pilawa-Podgurski, A. Sagneri, and C. Sullivan, "Opportunities and challenges in very high frequency power conversion," in *Twenty-Fourth Annual IEEE Applied Power Electronics Conference and Exposition (APEC)*, 2009, pp. 1-14.
- Yehui Han; Cheung, G.; An Li; Sullivan, C.R.; Perreault, D.J., "Evaluation of Magnetic Materials for Very High Frequency Power Applications," *Power Electronics, IEEE Transactions on*, vol. 27, no. 1, pp. 425-435, Jan. 2012.
- Di Yao; Sullivan, C.R., "Effect of capacitance on eddy-current loss in multi-layer magnetic films for MHz magnetic components," *Energy Conversion Congress and Exposition (ECCE)*, 2009, pp. 1025-1031.

Energy-Efficient Wireless Power Transmitters and Receivers

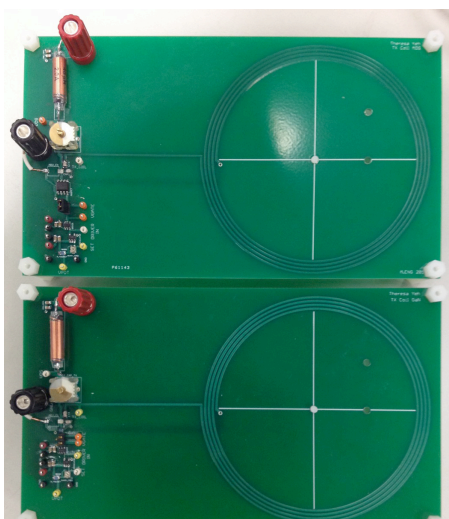
N. Desai, T. Yeh, T. Palacios, A. P. Chandrakasan
Sponsorship: Foxconn, MIT/MTL GaN Energy Initiative

The ability to transfer power wirelessly to a device greatly enhances the convenience of using portable electronics in daily life. Even though far-field radiative transfer of power remains an elusive goal owing to regulatory and/or safety concerns, near-field wireless charging using coupled inductor coils has been receiving a lot of attention over the past few years. However, these systems still need to address issues such as standardization and energy efficiency to be deployed on a wide variety of commercial devices.

Losses in coupled-inductor systems arise primarily due to highly imperfect ($k < 0.1$) coupling between the transmitter and receiver coils and small magnetizing inductance of the air core, which results in a lot of reactive energy sloshing back and forth. Compensating the reactances of the coils using a resonant network can address this problem. However, to keep component sizes reasonably small, this technique implies operation in the several MHz-range, which leads to higher transistor switching losses. To address this issue, we designed a

soft-switched (Class E), resonant transmitter that uses a GaN switch to reduce switching losses on the input of the switch (Figure 1). Results showed a 5% improvement in efficiency just by using a GaN switch over a silicon switch for similar (~ 10 W) transmit power levels.

Another interesting area of research concerns wirelessly charging wearable electronics. As the functionality of these devices grows, they will need multiple bursts of energy over the course of a day from an energy-constrained source such as a cellphone or tablet. This makes maximizing the end-to-end system efficiency very important. The efficiency is a strong function of the load impedance, coupling and other circuit parameters. The load impedance presented by a battery varies across its charging cycle, while the coupling coefficient is a dynamic parameter in a scenario where the user holds the cellphone over the fitness tracker to charge it. We are designing integrated circuits on the receiver side to track these changing conditions and adapt accordingly to maintain maximum efficiency operation.



▲ Figure 1: Design of a zero-voltage switched (ZVS) wireless power transmitter with Si FET (top) and GaN FET (bottom). Due to lower input switching losses with the GaN switch, the transmitter on the bottom reduces losses by 5%.



▲ Figure 2: Using a cellphone to provide a burst of energy to a wearable device such as a fitness tracker can typically charge it within a few minutes.

FURTHER READING

- A. Karalis, J. D. Joannopoulos, and M. Soljačić, "Efficient wireless non-radiative mid-range energy transfer," *Annals of Physics*, vol. 323, no. 1, pp. 34-48, Jan. 2008.
- N. Sokal and A. Sokal, "Class E-A new class of high-efficiency tuned single-ended switching power amplifiers," *IEEE Journal of Solid-State Circuits*, vol. 10, no. 3, pp. 168-176, June 1975.
- T. Yeh, "Efficient wireless charging with Gallium Nitride FETs," Master's thesis, Massachusetts Institute of Technology, Cambridge, 2014.

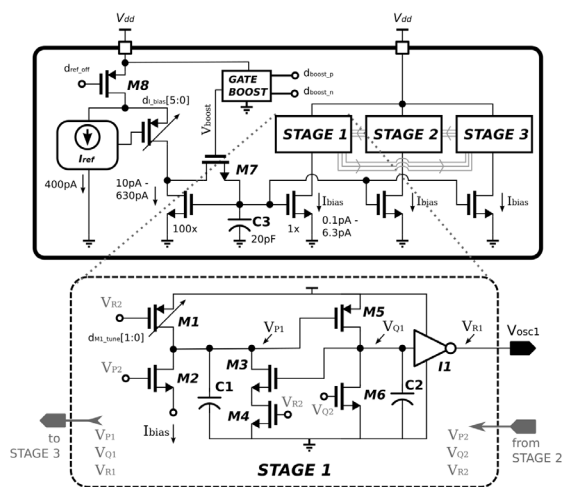
Picowatt Timer for Energy-Constrained, Battery-Less Systems

P. M. Nadeau, A. Paidimarri, A. P. Chandrakasan

Sponsorship: Texas Instruments, TSMC University Shuttle Program, NSERC

Energy harvesting presents an attractive option for powering wireless Internet-of-things devices; however, the most ubiquitous of sources, such as indoor lighting, ambient RF, system vibration, or body heat, all garner an extremely limited amount of average power. With a recent demonstration of harvesting as low as 1 nW of average power, there is a need to design always-on circuits, such as the wake-up timer, to take advantage of these weak sources.

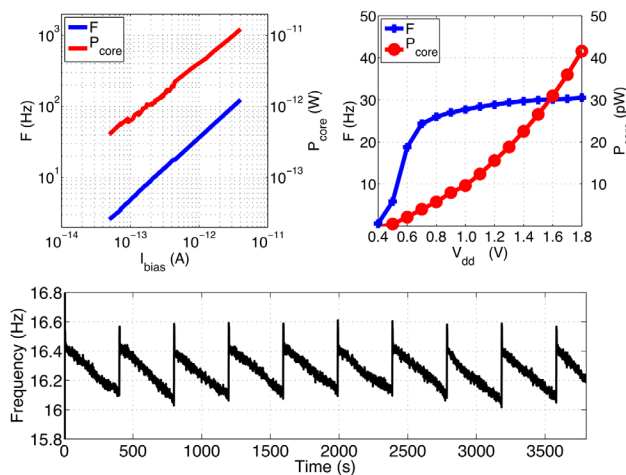
This abstract presents the design of wake-up timer that consumes 4.2 pW of power for 18 Hz of oscillation (0.23 pJ/cycle). The design features a dynamic 3-stage architecture, duty-cycled current-source, and low operating voltage (0.6 V) enabled by a voltage boost circuit. The circuit details are shown in Figure 1. M1 precharges C1 to V_{dd} and M6 precharges C2 to ground. Then, at the beginning of a timing cycle, I_{bias} linearly discharges C1 until M5 begins charging C2 and turns on M3. The positive feedback generated by M5 and M3 quickly flips the cell, generating a sharp edge to subsequent logic, such as inverter I1. A three-stage design provides all of the control signals necessary to precharge and trigger the timing of each stage in sequence.



▲ Figure 1: Circuit diagram of the timer system showing the current reference and interconnection of the stages (top), and detailed circuit for one stage (bottom).

Since I_{bias} should ideally remain stable across voltage and temperature, it is generated by a current source referenced to an on-chip resistor, which is then duty-cycled to save power. At the beginning of operation, M7 is enabled by a boosted voltage for a short period of time so that the current reference can set I_{bias} to the desired level. Then, M7 is disabled, and I_{bias} is fixed by the analog voltage stored on C3. A refresh of the voltage on C3 is performed periodically to compensate for the leakage current through M7, which serves to discharge this voltage over time.

Measurements of the system are shown in Figure 2. The design has scalable performance across 2 orders of magnitude in frequency and power by adjusting the I_{bias} current. Performance versus V_{dd} is also assessed, and the power consumption follows near square-law (e.g., proportional to CV^2f) dependence up to 1.8V, owing to the benefit of the positive regeneration in eliminating short-circuit currents. The figure also shows the performance of the track-and-hold scheme. Refreshing the I_{bias} current every 400 s leads to an average power consumption of 2.2 pW for the reference circuit (overall 4.2 pW for the system) and a $\pm 2\%$ variation in the generated frequency.



▲ Figure 2: Measurements of timer system including performance vs. I_{bias} (top-left), vs. supply voltage (top-right), and performance when integrated with the track-and-hold scheme (bottom).

FURTHER READING

- S. Bandyopadhyay, P. Mercier, A. Lysaght, K. Stankovic, and A. Chandrakasan, "A 1.1 nW Energy-Harvesting System with 544 pW Quiescent Power for Next-Generation Implants," *IEEE J. Solid-State Circuits*, vol. 49, no. 12, pp. 2812-2824, Dec. 2014.
- Y.-S. Lin, D. Sylvester, and D. Blaauw, "A 150pW program-and-hold timer for ultra-low-power sensor platforms," *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, Feb. 2009, pp. 326-327.
- P. Nadeau, A. Paidimarri, and A. Chandrakasan, "4.2 pW Timer for Heavily Duty-Cycled Systems," presented at *IEEE Symposium on VLSI Circuits*, Kyoto, Japan, June 2015.

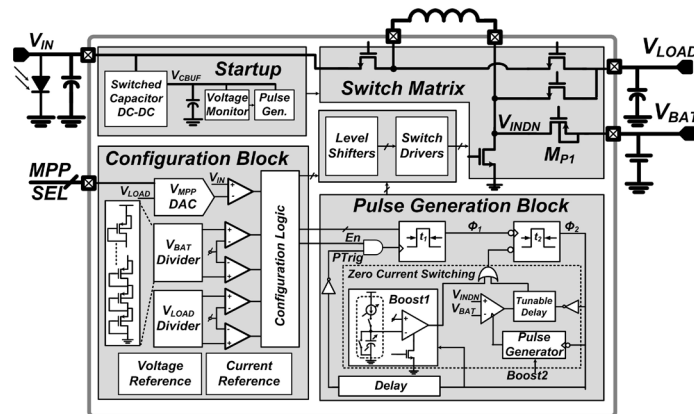
Solar Energy Harvesting System with Integrated Battery Management and Startup using Single-Inductor and 3.2 nW Quiescent Power

D. El-Damak, A. P. Chandrakasan

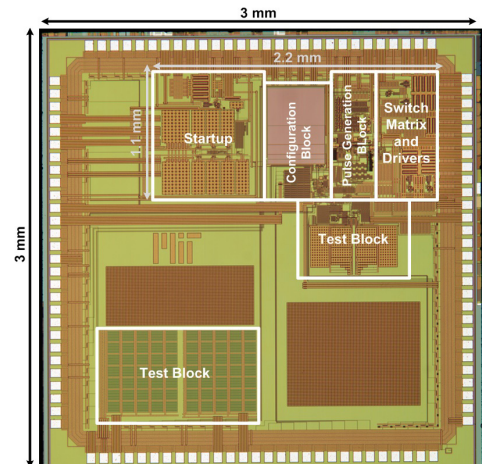
Sponsorship: TSMC University Shuttle Program, Texas Instruments Fellowship for Women in Microelectronics

Energy harvesting systems have allowed the autonomous operation of ultra-low power devices for implantable and wearable applications. Since the power sources available in our surroundings are intermittent, extreme energy-efficiency of the power management circuits is required. Thus, this work presents an energy harvesting system with 3.2 nW of quiescent power for solar-based applications. Figure 1 shows the chip's top-level architecture. The IC integrates the converter switch matrix with the associated configuration logic and drivers, voltage reference, current reference, startup, regulation, and battery management circuits. It can supply a 1V regulated voltage rail, V_{LOAD} , and charge a battery, V_{BAT} using a single 47 μ H inductor. The switch matrix supports three main configurations in addition to the startup: (a) Boost1: from the solar harvester V_{IN} to V_{LOAD} , (b) Boost2: from V_{IN} to V_{BAT} , and (c) Buck: from V_{BAT} to V_{LOAD} , which is used when the power available

from the solar cell is insufficient to keep the 1V output regulated. The control circuit is designed in an asynchronous fashion that scales the effective switching frequency of the converter with the level of the power transferred. The on-time of the converter switches adapts dynamically to the input and output voltages for peak-current control and zero-current switching. The chip operates efficiently with input power that ranges from 10 nW to 1 μ W. Figure 2 shows the micrograph of the fabricated chip with an active area of 2.2x1.1 mm². For input power of 500 nW, the proposed system achieves an efficiency of 82%, including the control circuit overhead, while charging a battery at 3 V from 0.5 V input. In buck mode, it achieves a peak efficiency of 87% and maintains efficiency greater than 80% for output power of 50 nW-1 μ W with input voltage of 3 V and output voltage of 1 V.



▲ Figure 1: Top level architecture of the system.



▲ Figure 2: Chip micrograph.

FURTHER READING

- D. El-Damak and A. Chandrakasan, "Solar energy harvesting system with integrated battery management and startup using single inductor and 3.2nW quiescent power," *IEEE Symposium on VLSI Circuits (VLSIC)*, pp. C280-C281, June 2015.
- S. Bandyopadhyay, and A. P. Chandarkasan, "Platform Architecture for Solar, Thermal, and Vibration Energy Combining with MPPT and Single Inductor," *IEEE Journal of Solid-State Circuits*, vol. 47, no. 9, pp. 2199-2215, Sept. 2012.
- K. Kadirvel, Y. Ramadass, U. Lyles, J. Carpenter, V. Ivanov, V. McNeil, A. Chandrakasan, B. Lum-Shue-Chan "A 330nA Energy-Harvesting Charger with Battery Management for Solar and Thermoelectric Energy Harvesting," *IEEE ISSCC Dig. Tech. Papers*, pp. 106-107, Feb. 2012.

Authentication Tags for Supply Chain Integrity

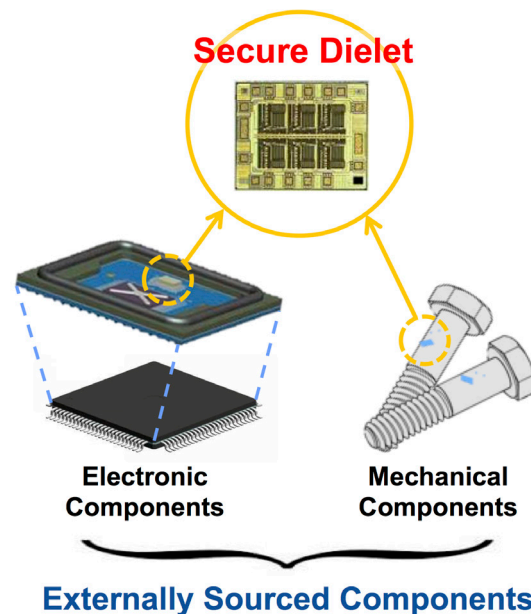
C. Juvekar, H.-M. Lee, A. P. Chandrakasan, J. Kwong (Texas Instruments)
Sponsorship: Texas Instruments, Denso

Counterfeiting is a major problem in commodity markets. A global supply chain exacerbates this problem by making it even harder to maintain the integrity of the components due to the sheer diversity of suppliers. In fact, most counterfeit components make it into the supply chain despite the best effort of the suppliers simply because the suppliers lack the capability to detect them. DARPA has proposed the Supply Chain Hardware Integrity for Electronics Defense (SHIELD) program to address these concerns.

In this project, we implement the concept of electronic fingerprinting to a wide range of components (Figure 1). To achieve this goal we are building integrated tags that act as cryptographically unique identifiers for both electronic and mechanical components. Our tags satisfy the goals of low cost and secure operation through a combination of protocol, circuits, and technology innovation.

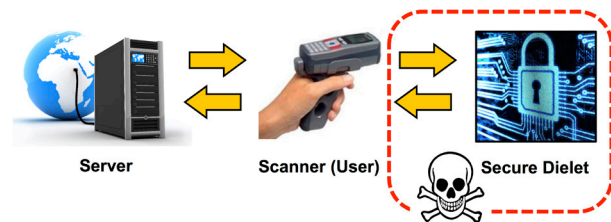
We leverage unique technology features provided by our fabrication partner (Texas Instruments) to build authentication circuits that can operate reliably in the presence of intermittent powering. A dynamic on-chip key memory ensures added protection against conventional side-channel as well as more advanced semi-invasive imaging attacks. To preserve the small form factor and facilitate stand-alone operation, we have integrated wireless power and telemetry circuits with on-chip passives. This enables low-cost system integration on account of zero off-chip components. Finally we have developed a custom cryptographic challenge-response protocol that allows the server to securely authenticate the tag and maintain seamless synchronization with our key storage.

We envision that the tag will be used in conjunction with a server that maintains a database of valid issued tags. To validate the authenticity of purchased components, a user would scan the affixed tag using a handheld scanner and run a challenge response protocol to match cryptographic information stored on the tags with the server database (Figure 2).



Externally Sourced Components

▲ Figure 1: Common stand-alone authentication tags that can be attached to various electronic and mechanical components



▲ Figure 2: Illustration of a complete authentication system that includes the tag, handheld scanner, and server database. Our tags protect against physical and side-channel attacks on the tag.

FURTHER READING

- P. Taylor, "Fake part forces Aston Martin to recall thousands of cars," <http://www.securindustry.com/electronics-and-industrial/fake-part-forces-aston-martin-to-recall-thousands-of-cars-/s105/a1959/>
- M. Sale, "Can EDA Combat the Rise of Electronic Counterfeiting?" presented at *Design Automation Conference (DAC) 2012*, Session 8.1, 2012.

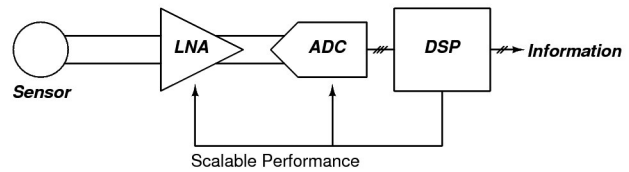
Low-Power Sensor Interfaces for Wireless Sensor Nodes

F. M. Yaul, A. P. Chandrakasan
Sponsorship: Shell, Texas Instruments

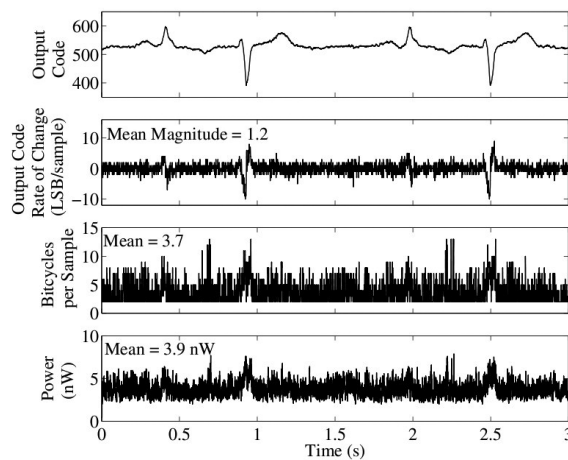
Sensor interface circuits found in wireless sensor nodes (WSNs) typically comprise a low-noise amplifier (LNA), analog signal conditioning circuits, an analog-to-digital converter (ADC), and a digital signal processor (DSP) to extract information from the signals. These blocks are shown in Figure 1. Reducing system power consumption allows for increased sensor node lifetime, which is especially important for nodes that are difficult to physically access. Examples include medical implants and pipeline monitoring nodes.

One component of this project has been an ADC which takes advantage of low signal activity in order to save power. Since low signal activity is common to many sensor signals, the ADC can save power in a broad range of applications. This work introduces an altered successive approximation (SA) algorithm called LSB-First SA, which is designed to reduce the number of bitcycles per conversion, given a good initial guess of the value of the sample. Figure 2 depicts the ADC's response to an ECG input signal and demonstrates the ADC's ability to save power and perform 10-bit conversions in just 3.7 bitcycles/sample on average when the signal is only varying by 1.2 LSBs/sample on average.

Current work on this project involves exploring power reduction techniques in the LNA and DSP blocks, as well as investigating ways to exploit structure in the signal information to save power in the overall system.



▲ Figure 1: Block diagram of signal-to-information chain. Based on the structure of the information, the performance of the analog front-end may be scaled to save power.



▲ Figure 2: ADC response to ECG test input signal with $f_s = 1$ kHz and $V_{DD} = 0.5$ V, demonstrating the ADC's low leakage and data-dependent energy consumption.

FURTHER READING

- F. M. Yaul, and A. P. Chandrakasan, "A 10 bit SAR ADC with Data-Dependent Energy Reduction Using LSB-First Successive Approximation," *IEEE Journal of Solid-State Circuits*, vol. 49, no. 12, Dec. 2014, pp. 2825-2834.
- M. Yip and A. P. Chandrakasan, "A Resolution-Reconfigurable 5-to-10-Bit 0.4-to-1 V Power Scalable SAR ADC for Sensor Applications," *IEEE Journal of Solid-State Circuits*, vol. 48, no. 6, June 2013, pp. 1453-1464.
- N. Verma and A. P. Chandrakasan, "A Micro-Power EEG Acquisition SoC With Integrated Feature Extraction Processor for a Chronic Seizure Detection System," *IEEE Journal of Solid-State Circuits*, vol. 45, no. 4, Apr. 2010, pp. 804-816.

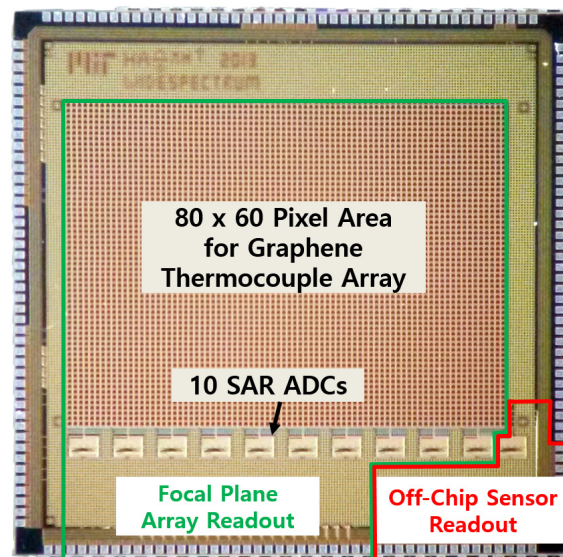
A Graphene-CMOS Hybrid Sensor for Thermal Imaging

S. Ha, A. Hsu, Y. Lin, M. Hempel, C. Mackin, J. Kong, T. Palacios, A. P. Chandrakasan

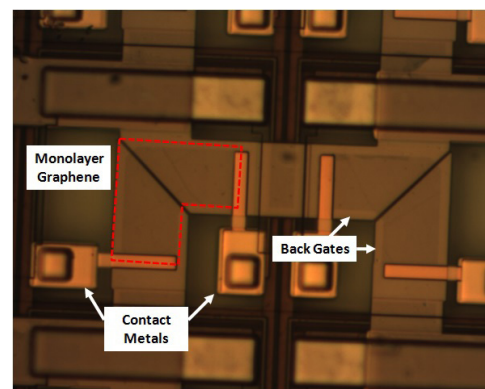
Sponsorship: Center for Integrated Circuits and Systems, Institute for Soldier Nanotechnologies

CMOS image sensors are widely used in digital multi-media applications thanks to their mature production technology; the performance ramps up every year with denser integration, better noise suppression, adjustable dynamic range, and lower power. However, the band gap of silicon fundamentally limits the absorption spectrum to the visible and near-infrared light ($\lambda < 1100\text{nm}$). To overcome the physical limitation of a Si photodetector in longer wavelength detection, we employ a tunable ambipolar graphene-based thermocouple. With an integrated graphene-based thermocouple array, the hybrid IR image sensor can detect the long-wavelength infrared spectrum that is used in security cameras and medical thermal imaging and in spectroscopy for chemistry and astronomy. Furthermore, the present sensing platform solves critical issues in manufacture of conventional IR sensors and enables expanded applications such as a high-speed and high-resolution IR imager and a hyperspectral IR imaging IC.

First, a Si CMOS readout chip using commercially available 0.18- μm technology was designed. The readout circuit consists of pixel transimpedance amplifiers, row-column multiplexers, and parallel column analog-to-digital converters (ADCs). The layout design of the chip required careful considerations on the post-fabrication process to ease the graphene integration processes, or at least make them possible, including standard ebeam lithography, metalization, etching, and graphene transfer steps. For instance, the pixel amplifier and signal paths occupy only the small portion of each 50- μm x 50- μm pixel area, leaving over 60% of the area empty and flat to enable fabrication of back gate electrodes for the graphene thermocouple and high-yield graphene transfer. Secondly, the graphene-integrated readout chip was fabricated as shown in Figure 1. Figure 2 shows a closer look at the fabricated on-chip graphene thermocouple pixel. An electrostatically induced graphene p-n junction of graphene generates a thermoelectric voltage in response to IR light absorption. The terminals of the thermocouple are connected to the M6 pillars of the pixel amplifier input; the signal is amplified and converted into 8-bit codes. In the testing setup, an off-chip processor renders the data in 2D image.



▲ Figure 1: The integrated readout IC for graphene-CMOS hybrid IR imager consists of an array of 80 x 60 graphene thermocouples and pixel amplifiers in the center area and row and column selection logic, current sources, column amplifiers, and ADCs in the periphery to ease the fabrication.



▲ Figure 2: Top view of the pixels with graphene thermocouples connected to the pixel amplifier inputs. The monolayer graphene sheet is biased by two metal gates embedded in the dielectric layer.

FURTHER READING

- P. K. Herring, A. Hsu, N. M. Gabor, Y. C. Shin, J. Kong, T. Palacios, and P. Jarillo-Herrero, "Photoresponse of an Electrically Tunable Ambipolar Graphene Infrared Thermocouple," *Nano Letters*, vol. 14, no. 2, pp. 901-907, Jan. 2014.
- A. Hsu, P. K. Herring, N. M. Gabor, S. Ha, Y. C. Shin, Y. Song, A. P. Chandrakasan, J. Kong, P. Jarillo-Herrero, and T. Palacios, "Graphene-based Thermopile for Thermal Imaging Applications," submitted for publication.