

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

12-bit 330MS/s CMOS Pipelined Analog-to-Digital Converter	3
A 10-bit SAR ADC with Data-Dependent Energy Savings Using LSB-First Successive Approximation	4
Time-interleaved A/D Converters.....	5
High-precision Zero-Crossing/Op-amp Hybrid ADC	6
A Flash ADC with Reduced Number of Comparators.....	7
Continuous-time Delta-sigma Analog-to-digital Converters for Application to Multiple-input Multiple-output Systems	8
A 10-b, 1GS/s Time-interleaved SAR ADC with a Background Timing Skew Calibration	9
High-performance Analog-to-Digital Converter in GaN-on-Silicon Technology.....	10
SCORPIO: A 36-core Research Chip Demonstrating Snoopy Coherence on a Scalable Mesh NoC with In-network Ordering.....	11
A Case for Leveraging 802.11p for Direct Phone-to-phone Communication.....	12
Impact of Digital Clock Tree on LC-VCO Performance in 3D-IC.....	13
A High-throughput CABAC Decoder Architecture for the Latest Video Coding Standard HEVC with Support of High-level Parallel Processing Tools.....	14
Full HD Integrated Video Encoder for H.265/HEVC	15
Energy-efficient Hardware for Object Detection.....	16
Algorithm and Architecture Enhancements for Hardware Speech Recognition	17
Demonstration Platform for Energy-scalable Ultrasound Beamforming	18
Self-powered Long-range Wireless Microsensors for Industrial Applications.....	19
Graphene-CMOS Hybrid Infrared Image Sensor	20
A Multilevel Energy Buffer and Voltage Modulator for Grid-interfaced Micro-inverters	21
Efficient Wireless Charging with Gallium Nitride FETs	22
Efficient Portable-to-Portable Wireless Charging.....	23
A 28-nm FDSOI Integrated, Reconfigurable, Switched-Capacitor Step-up DC-DC Converter with 88% Peak Efficiency	24

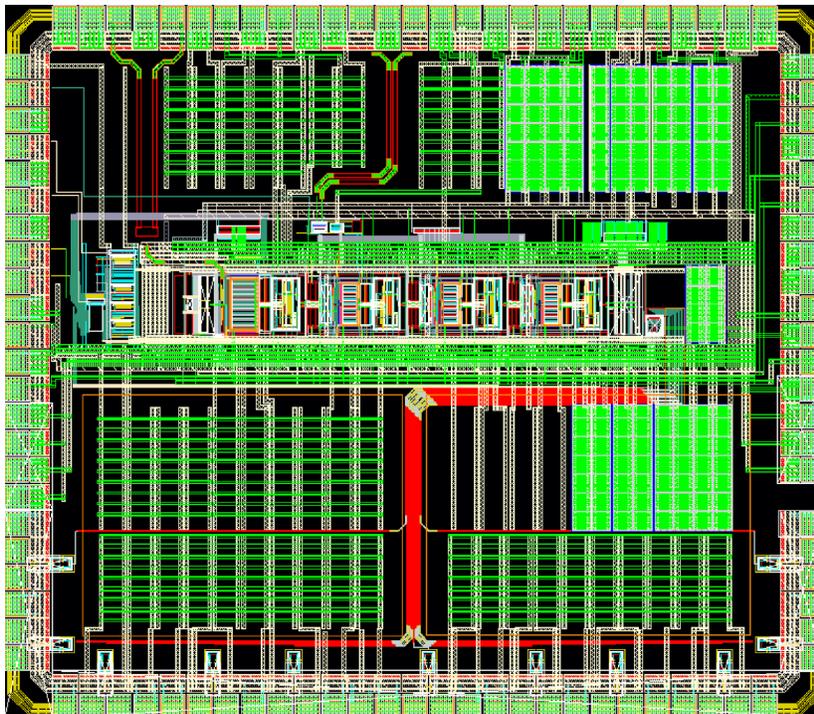
12-bit 330MS/s CMOS Pipelined Analog-to-Digital Converter

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

Sponsorship: Masdar Institute of Science and Technology / MIT Cooperative Program

Pipelined ADC architecture is widely applicable as it offers high-resolution and wide-bandwidth data conversions. The scaling of the CMOS devices, however, accompanies lower supply voltage, reduced intrinsic transistor gain, and degraded signal-to-noise ratio (SNR). The op-amp becomes the performance bottleneck as it suffers from low gain, bandwidth, and noise issues. Approaches have been reported in the literature that address these problems by either employing low-performance op-amps and resolving degradations by digital calibrations or proposing new architectures that replace the op-amps.

We propose innovative circuit techniques to relax the op-amp gain, bandwidth, and noise requirements. The approach enables an energy-efficient high resolution and high sampling speed pipelined ADC. A 12-bit, 330MS/s prototype ADC chip is designed in TSMC 65-nm LP. The differential input signal range is 1.5V peak-peak and the ADC consumes 49mW of power with 1.2V power supply at 330MS/s. The full chip layout is shown in Figure 1, and the core area size is 0.56mm². The fabricated prototype is currently being tested.



◀ Figure 1: Full chip layout. The die area is 2.6mm x 2.26mm; the ADC core area is 1.55mm x 0.36mm.

FURTHER READING

- 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.
- L. Brooks and H.-S. Lee, "A zero-crossing-based 8-bit 200MS/s pipelined ADC," *IEEE J. Solid-State Circuits*, vol. 42, pp. 2677-2687, Dec., 2009.

A 10-bit SAR ADC with Data-Dependent Energy Savings Using LSB-First Successive Approximation

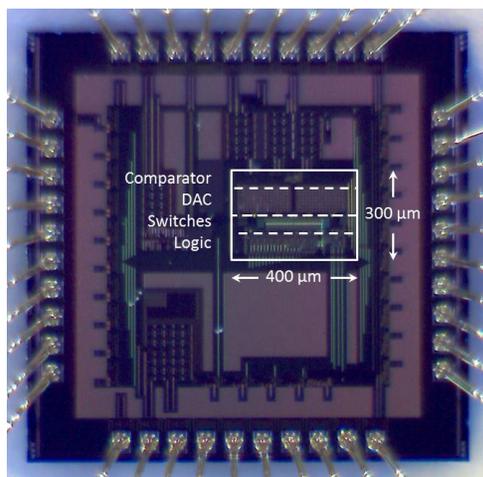
F. M. Yaul, A.P. Chandrakasan
Sponsorship: Shell, Texas Instruments

ADCs used in medical and industrial monitoring often transduce signals with short bursts of high activity followed by long idle periods. Examples include biopotential, sound, and accelerometer waveforms. Current approaches to save energy during periods of low signal activity include variable sample rate and resolution ADCs, asynchronous level-crossing ADCs, and application-specific ADCs with non-uniform quantization boundaries or dead-zones.

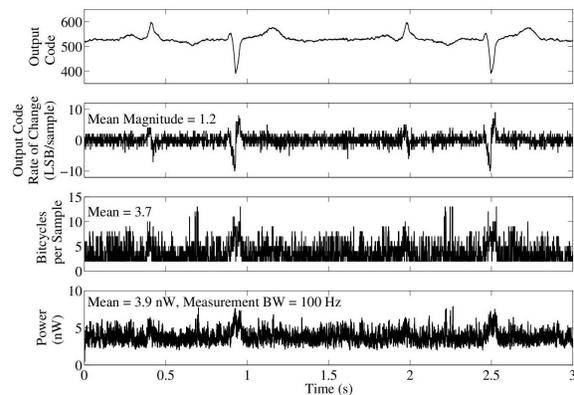
The motivation for this work is to create an ADC which takes advantage of low signal activity to save power, which can help extend the lifetimes of the medical implants or wireless sensor nodes that they are used in. Since low signal activity is common to many sensor signals, the ADC can save power in a broad range of applications. This work leverages the energy-efficient architecture of the highly-digital successive approximation register (SAR) ADC topology, and

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. Because low activity signals are the target application, a good initial guess of the current sample is simply the result of the previous sample taken by the ADC. Since each bitcycle uses an analog comparison, a DAC transition, and many logic transitions, bitcycle reduction saves power in all parts of the SAR ADC.

A more detailed description of the algorithm may be found in the references. Figure 1 shows a photo of the 1 mm² silicon die containing the LSB-First SAR ADC. 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 varies by only 1.2 LSBs/sample on average.



▲ Figure 1: LSB-first SAR ADC chip micrograph showing twin capacitive DACs, comparator, sampling switches, and LSB-First bitcycling logic.



▲ 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 10b 0.6nW SAR ADC with Data-Dependent Energy Savings Using LSB-First Successive Approximation," in *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, Feb. 2014, pp. 198.
- 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.

Time-interleaved A/D Converters

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

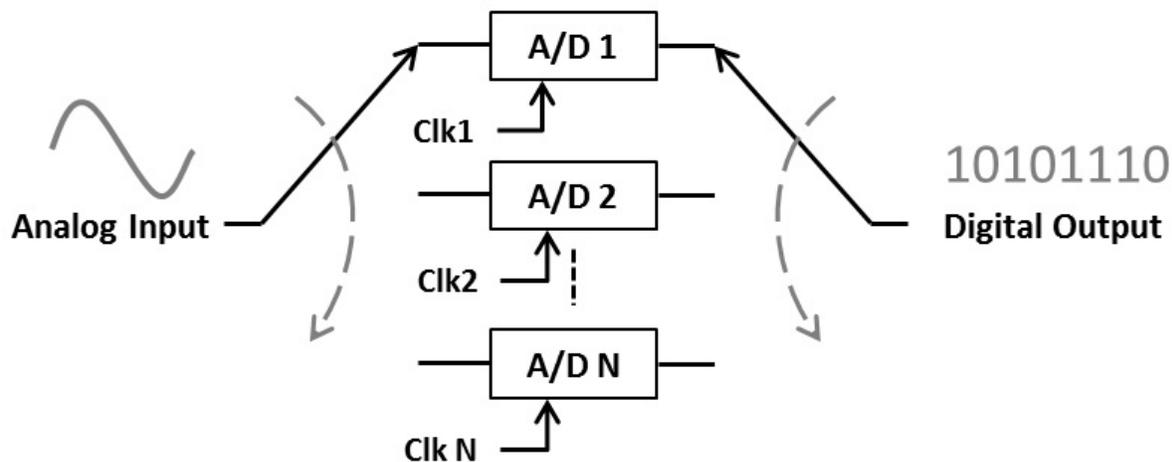
Sponsorship: Masdar Institute of Science and Technology

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 is delayed 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 12-bit, 240MS/s, 4-way time-interleaved A/D is being taped-out to demonstrate the new calibration scheme.



▲ 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

High-precision Zero-Crossing/Op-amp Hybrid ADC

M. Markova, H.-S. Lee

Sponsorship: Center for Integrated Circuits and Systems

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 that replaces the amplifier with a current source and a comparator. The technique has been generalized to zero-crossing based circuits (ZCBCs). It has been demonstrated but not limited to single-ended and differential pipelined ADCs, with an effective number of bits (ENOB) ranging from 8 bits to 11 bits at sampling rates from 10MS/s to 100MS/s.

The purpose of this project was to explore the use of the ZCBC technique for high-precision ADCs. The goal of the project is a 13-bit pipelined ADC operating at up to 100MS/s. A two-phase hybrid ZCBC operation is used to improve the power-linearity tradeoff of the A/D conversion. The first phase approximates the final output value, while the second phase allows the output to settle to its accurate value. Since the output is allowed to settle in the second phase, the currents through capacitors decay, permitting

higher accuracy and power-supply rejection than in standard ZCBCs. Linearization techniques for the ramp waveforms are implemented. Linear ramp waveforms require less correction in the second phase for a given linearity, thus allowing faster operation. We explored techniques for improving linearity beyond using a cascoded current source; these techniques include output pre-sampling and bidirectional output operation. Current steering is used to minimize the overall delay contributing to the first phase error, known as overshoot error. Overshoot error reduction at the end of the first phase improves the linearity requirements of the final phase. We designed a prototype ADC in 1V, 65nm CMOS process to demonstrate the techniques introduced in this work. The prototype ADC achieved 11-bit ENOB at 21MS/s and SFDR of 81dB. The main performance limitations are lack of overshoot reduction in the third pipeline stage in the prototype ADC and mid-range errors, introduced by the bidirectional ramp linearization technique, limiting the attainable output accuracy.

FURTHER READING

- J. K. Fiorenza, T. Sepke, P. Holloway, C. G. Sodini, and H.-S. Lee, "Comparator-Based Switched-Capacitor Circuits for Scaled CMOS Technologies," *IEEE Journal of Solid-State Circuits*, vol. SC-41, pp. 2658-2668, Dec. 2006.
- L. Brooks and H.-S. Lee, "A Zero-Crossing-Based 8-bit 200 MS/s Pipelined ADC," *IEEE Journal of Solid-State Circuits*, vol. SC-42, pp. 2677-2687, Dec. 2007.
- L. Brooks and H.-S. Lee, "A 12b, 50 MS/s, Fully Differential Zero-Crossing Based Pipelined ADC," *IEEE Journal of Solid-State Circuits*, vol. SC-44, pp. 3329-3343, Dec. 2009.
- S. Lee, A. P. Chandrakasan, and H.-S. Lee, "A 12b 5-to-50MS/s 0.5-to-1V Voltage Scalable Zero-Crossing Based Pipelined ADC," *IEEE Journal of Solid-State Circuits*, vol. SC-47, pp. 1603-1614, July 2012.

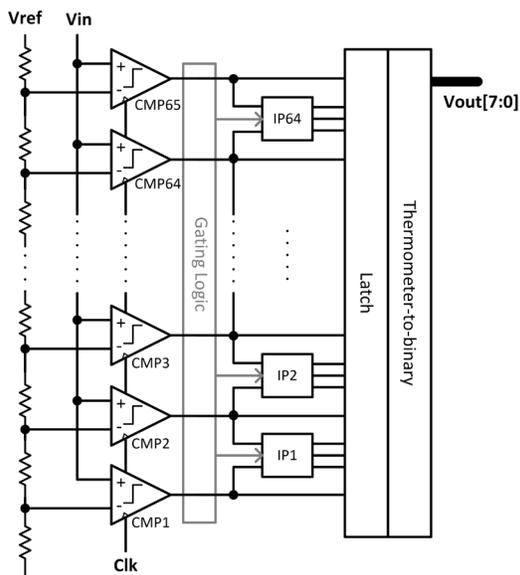
A Flash ADC with Reduced Number of Comparators

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

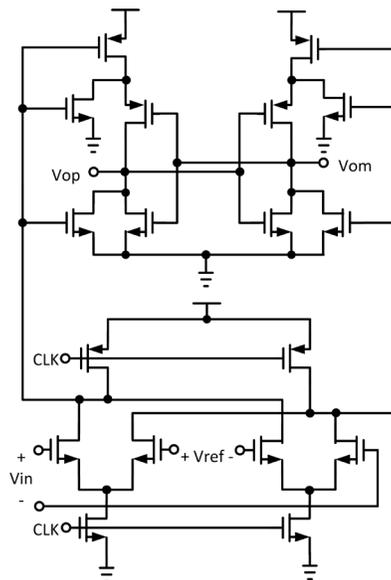
Sponsorship: MIT/MTL GaN Energy Initiative, Office of Naval Research, Samsung

High-speed and low-resolution flash analog-to-digital converters (ADCs) are widely used in applications such as 60GHz 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 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.

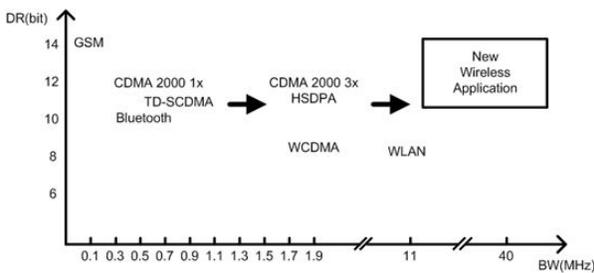
Continuous-time Delta-sigma Analog-to-digital Converters for Application to Multiple-input Multiple-output Systems

D. Yoon, H.-S. Lee
Sponsorship: MediaTek, Inc.

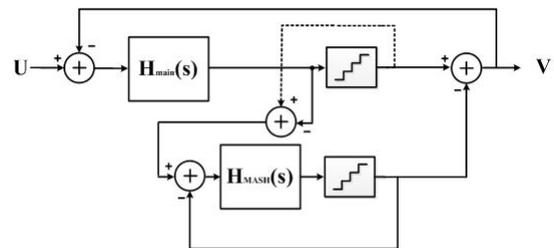
As wireless communication technology is rapidly advancing, new wireless applications are continuously developed. Figure 1 shows each application space and the required dynamic range. The new wireless applications demand wideband (50 MHz) and high resolution data converters (>14 bits). Delta-sigma ($\Delta\Sigma$) analog-to-digital converters (ADCs) are best suited for their ability to achieve high resolution. However, the large bandwidth required poses a significant challenge. A $\Delta\Sigma$ ADC can be implemented in either a discrete-time (DT) or a continuous-time (CT) structure. Since DT $\Delta\Sigma$ ADCs require op-amp settling within each half clock period, the gain-bandwidth requirement for the op-amp is extremely high for the sampling rate required for 50MHz bandwidth. The CT $\Delta\Sigma$ ADCs require much lower gain-bandwidth. Thus, CT $\Delta\Sigma$ ADCs can function at a higher sampling frequency and achieve a wider bandwidth compared to DT $\Delta\Sigma$ ADCs. In addition, since the CT $\Delta\Sigma$ ADCs are more power-efficient and have an inherent anti-aliasing property, they are more suitable

for the demanding new wireless applications.

This project focuses on the design of CT $\Delta\Sigma$ ADCs, specifically for the application in multiple-input multiple-output wireless receivers. For this application, each CT $\Delta\Sigma$ ADC in a channel must provide wide bandwidth and high dynamic range at low power consumption. The state-of-art CT $\Delta\Sigma$ ADCs fail to come close to either wide-enough bandwidth or high-enough dynamic range for such applications. We are investigating a new type of a CT multi-stage noise-shaping (MASH) $\Delta\Sigma$ ADC based on a DT sturdy-MASH $\Delta\Sigma$ ADC. Figure 2 shows the overall structure of a CT MASH $\Delta\Sigma$ ADC. The main advantage of this new type of CT $\Delta\Sigma$ ADCs is that it does not require digital filters that conventional MASH $\Delta\Sigma$ ADCs need to cancel out the quantization error of the first stage. We have developed several new techniques to make a CT MASH $\Delta\Sigma$ ADC faster, more accurate, and robust. The prototype $\Delta\Sigma$ ADC has been designed and taped out in 28nm CMOS technology.



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



▲ Figure 2: Block diagram of a CT MASH $\Delta\Sigma$ ADC.

FURTHER READING

- K. Lee, J. Chae, M. Aniya, K. Hamashita, K. Takasuka, S. Takeuchi, and G.C. Temes, "A noise-coupled time-interleaved delta-sigma ADC with 4.2 MHz bandwidth, 98 dB THD, and 79 dB SNDR," *IEEE J. Solid-State Circuits*, vol. 43, no. 12, pp. 2601-2612, Dec. 2008.
- Y.-S. Shu, J.-Y. Tsai, P. Chen, T.-Y. Lo, and P.-C. Chiu, "A 28fJ/conv-step CT $\Delta\Sigma$ Modulator with 78dB DR and 18MHz BW in 28nm CMOS Using a Highly Digital Multibit Quantizer," *ISSCC Dig. Tech. Papers*, pp. 268-269, Feb. 2013.
- P. Shettigar and S. Pavan, "A 15mW 3.6GS/s CT- $\Delta\Sigma$ ADC with 36MHz Bandwidth and 83dB DR in 90nm CMOS," *ISSCC Dig. Tech. Papers*, pp. 156-157, Feb.
- N. Maghari, S. Kwon, and U. Moon, "74 dB SNDR multi-loop sturdy-MASH delta-sigma modulator using 35 dB open-loop Op-amp gain," *IEEE J. Solid-State Circuits*, vol. 44, no. 8, pp. 2212-2221, Aug. 2009.

A 10-b, 1GS/s Time-interleaved SAR ADC with a Background Timing Skew Calibration

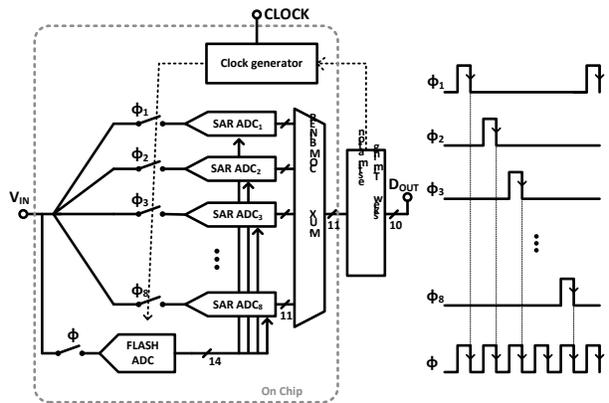
S. Lee, A.P. Chandrakasan, H.-S. Lee

Sponsorship: Center for Integrated Circuits and Systems, Samsung Fellowship

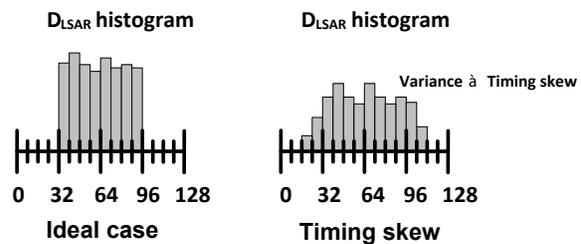
This work presents a time-interleaved (TI) SAR ADC that enables background timing skew calibration without a separate timing reference channel and enhances the conversion speed of each SAR channel. As shown in Figure 1, the proposed ADC incorporates a flash ADC operating at the full sampling rate of the TI ADC. The flash ADC output is multiplexed to resolve MSBs of the SAR channels.

Because the full-speed flash ADC does not suffer from timing skew errors, the flash ADC output is also used as a timing reference to estimate the timing skew of the TI SAR ADCs. However, this work differs from previous works in that no extra channel is required to serve as a timing-skew standard. Figure 2 shows the idea of timing-skew calibration. When the sampling signal of the flash ADC and the sampling signal of the SAR ADC are not aligned due to a timing-skew, the coarse estimation from the flash ADC is inaccurate. To recover the error of the flash ADC, the lower bits of SAR conversion output (D_{LSAR}) go beyond the normal range. Thus, the variance of the DSAR can be used as a measure of the timing-skew. This calibration puts no constraint on the input signal, and the calibration process does not interrupt the normal ADC operation. Thus, this calibration can be run in the background to track variations.

A prototype ADC is designed and fabricated in a 65-nm CMOS process. After background timing skew calibration, 51.4-dB SNDR, 59.1-dB SFDR, and ± 1.0 LSB INL/DNL are achieved at 1GS/s with a Nyquist rate input signal. The power consumption is 18.9mW from a 1.0V supply, which corresponds to 62.3fJ/step FoM.



▲ Figure 1: Block diagram of the proposed TI SAR ADC.



▲ Figure 2: Examples of the DLSAR histogram without timing skew (left) and with timing skew (right).

FURTHER READING

- M. El-Chammas and B. Murmann, "12-GS/s 81-mW 5-bit time-interleaved flash ADC with background timing skew calibration," *IEEE Journal of Solid-State Circuits*, vol. 46, no. 4, pp. 838–847, Apr. 2011.
- S. Lee, A. P. Chandrakasan, and H.-S. Lee, "A 1GS/s 10b 18.9-mW time-interleaved SAR ADC with background timing-skew calibration," in *IEEE ISSCC Dig. Tech. Papers*, pp. 384-385, Feb. 2014.

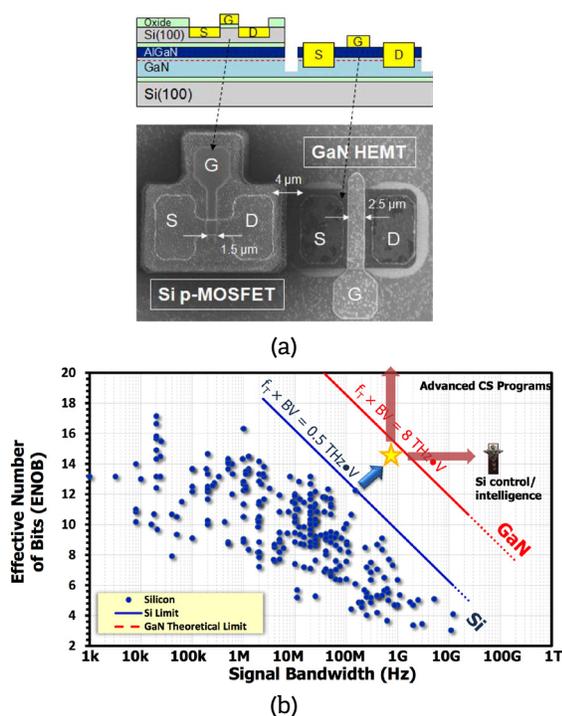
High-performance Analog-to-Digital Converter in GaN-on-Silicon Technology

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

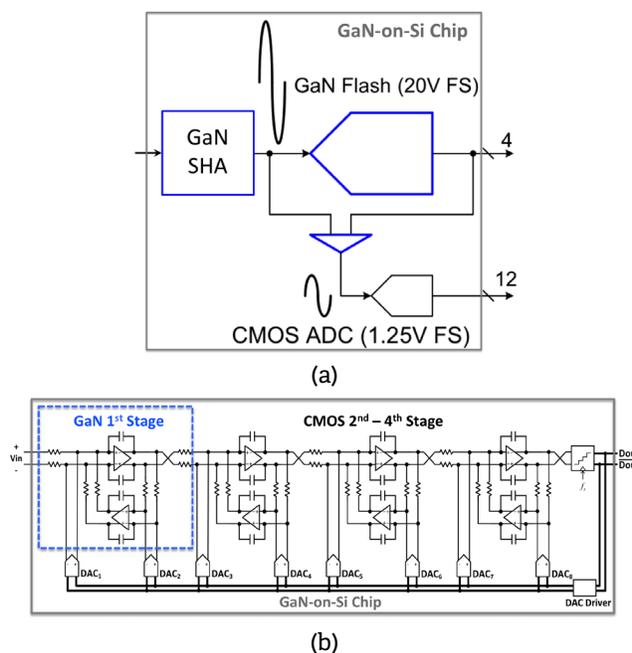
Sponsorship: MIT/MTL GaN Energy Initiative, Office of Naval Research

Mobile multimedia's coming of age with big data applications has spurred the development of extremely high-performance analog-to-digital converters (ADCs) for diverse emerging applications including personal communication, health care, and optical backbone networks. The low supply voltage of deeply scaled CMOS (complementary metal-oxide semiconductor) transistors limits the dynamic range of ADC input signal and consequently becomes a fundamental barrier to the performance of ADCs' built-in silicon technology. Recently, gallium-nitride-based high-electron mobility transistors (GaN HEMT) are reported to provide many advantages over the existing compound semiconductor technologies. The operation of GaN HEMTs at a very high supply voltage over 30V allows us to surpass the fundamental ADC SNR limit originating from ther-

mal noise and the limited signal range of the silicon technology. Due to the high power and relatively large feature size of GaN HEMTs, a hybrid process technology, which monolithically integrates GaN HEMTs with silicon CMOS transistors on a single wafer (GaN-on-Si), will take advantage of both technologies, enabling revolutionary mixed-signal performance, as Figure 1 shows. Our research focuses on the design of unprecedentedly high-performance ADCs in a GaN-on-Si hybrid technology, as in Figure 2. First, we are developing an over-100-dB SNR GaN sampler to culminate the performance of a GaN-on-Si pipeline ADC. Second, we are investigating a wide-swing GaN operational amplifier design for a continuous-time delta-sigma ADC with a very high dynamic range.



▲ Figure 1: GaN-on-Si technology: (a) Fabricated Si MOSFET and GaN HEMT in a monolithically integrated chip. (b) Revolutionary ADC performance expected from the monolithic GaN-on-Si integration.



▲ Figure 2: High-performance ADC architecture in GaN-on-Si hybrid process technology: (a) pipeline architecture for wide bandwidth. (b) 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," in *Proc. of the IEEE*, vol. 96, no. 2, pp. 287-305, Feb. 2008.
- J. W. Chung, J.-K. Lee, E. L. Piner, and 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.

SCORPIO: A 36-core Research Chip Demonstrating Snoopy Coherence on a Scalable Mesh NoC with In-network Ordering

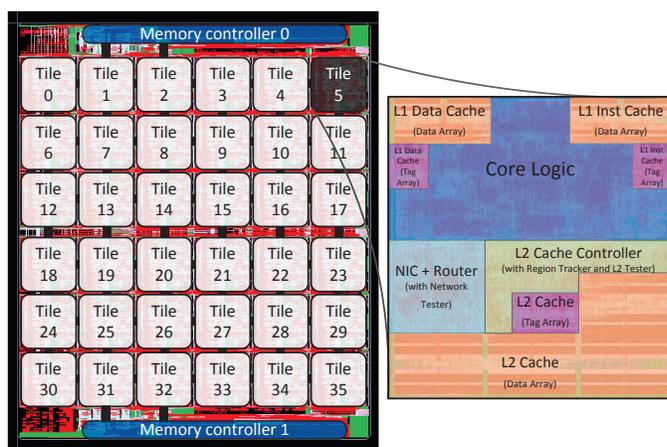
B.K. Daya, C.H.O. Chen, S. Subramanian, W.C. Kwon, S. Park, T. Krishna, J. Holt, A.P. Chandrakasan, L.S. Peh
Sponsorship: Center for Future Architectures (C-FAR), DARPA UHPC Grant (MIT Angstrom)

In the many-core era, scalable coherence and on-chip interconnects are crucial for shared memory processors. While snoopy coherence is common in small multicore systems, directory-based coherence is the de facto choice for scalability to many cores, as snoopy relies on ordered interconnects that do not scale. However, directory-based coherence does not scale beyond tens of cores due to excessive directory area overhead or inaccurate sharer tracking. Prior techniques supporting ordering on arbitrary unordered networks are impractical for full multicore chip designs.

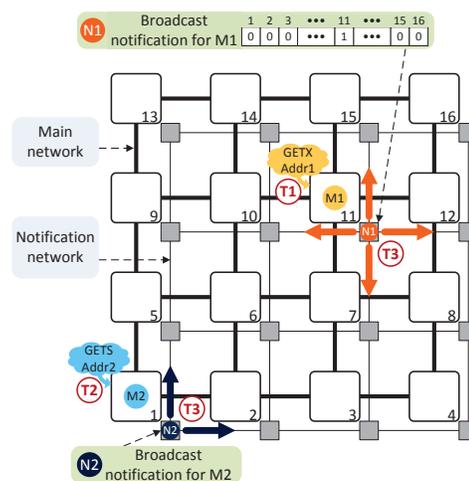
SCORPIO, an ordered mesh Network-on-Chip (NoC) architecture with a separate fixed-latency, bufferless network, achieves distributed global ordering. Decoupling message delivery from ordering allows messages to arrive in any order, at any time, yet be correctly ordered. For each message sent on the main network, a notification is broadcast on the bufferless network. Within a fixed number of cycles, all nodes receive notification. Processing the received

notification messages according to a consistent ordering rule means that all nodes determine locally the global order for messages in the main network. SCORPIO can plug-and-play with existing multicore IPs with priority to practicality, timing, area, and power. Full-system 36- and 64-core simulations on SPLASH-2 and PARSEC benchmarks show average application runtime reductions of 24.1% and 12.9%, vs. distributed directory and AMD HyperTransport coherence protocols, respectively.

Figure 1 shows SCORPIO in an 11 x 13 mm chip prototype, fabricated in IBM 45-nm SOI technology, comprising 36 Freescale e200 Power Architecture™ cores with private L1 and L2 caches interfacing with the NoC via ARM AMBA, plus two Cadence on-chip DDR2 controllers. The prototype achieves 1 GHz post-synthesis operating frequency (833MHz post-layout), estimated power of 28.8W (768mW per tile) with network consuming 10% of tile area and 19 % of its power.



▲ Figure 1: The 36-core fabricated multicore processor layout with SCORPIO NoC. Each tile contains in-order core, split L1 I/D caches, private L2 cache, L2 region tracker for destination filtering, and SCORPIO NoC components. The core assumes a bus is connected to AMBA AHB data and instruction ports, cleanly isolating it from details of the network and coherence support.



▲ Figure 2: At T1 and T2, cache controllers inject cache miss messages M1 and M2 at cores 11 and 1, respectively. Coherence requests are encapsulated into single flit packets and tagged with IDs of sources; IDs are broadcast to all nodes in main network. At T3, notification messages N1 and N2 corresponding to M1 and M2 are generated and sent to notification network.

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 Int'l Symposium on Computer Architecture* June 2014.

A Case for Leveraging 802.11p for Direct Phone-to-phone Communication

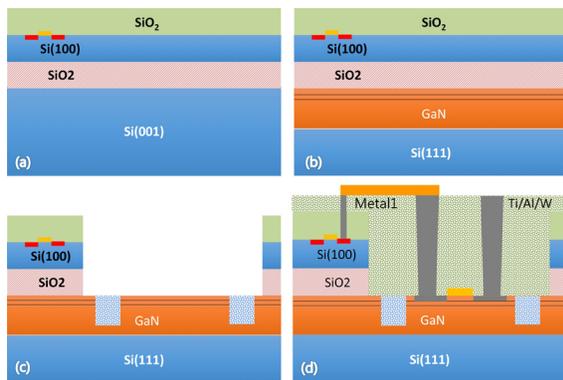
P. Choi, J. Gao, N. Ramanathan, M. Mao, S. Xu, C.C. Boon, S. Fahmy, L.S. Peh
Sponsorship: SMART-LEES program

Direct device-to-device communication between phones is readily supported via standards such as ad-hoc 802.11n or WiFi Direct. However, present WiFi standards cannot effectively handle the demands of these applications, due to insufficient range and poor reliability. We make the case for using 802.11p DSRC instead, which has been adopted as a standard for vehicle-to-vehicle communications, providing lower latency and longer range.

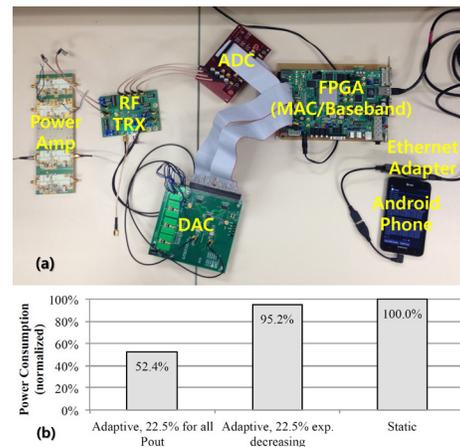
This work is the result of collaboration among materials and device researchers, circuits designers, and mobile systems and software architects. Motivated by a novel fabrication process that deposits both III-V and CMOS devices on the same die, we leveraged the GaN HEMT devices to realize the high-power amplifier tailored for adaptive power control and coupled that with a CMOS transmitter. We designed and fabricated an 802.11p-compliant power amplifier and a transmitter on commercial 0.25- μm GaN and 0.18- μm CMOS

process, respectively. This combination validates our vision for miniaturized and low-power DSRC chipsets. In our system prototype, the fabricated RF front-end is interfaced with an FPGA board implementing 802.11p digital baseband, connected to Android phones through USB. We use RoadRunner, an Android application to control road congestion, as a representative app requiring significant phone-to-phone communication.

The system consumes 0.02 $\mu\text{J}/\text{bit}$ for transmission across 100 m in 64-QAM mode, assuming free space. We demonstrate that application-level power control dramatically reduces power consumption by 47% for our RoadRunner application compared to a case without power control. The 1.98mm² die size demonstrates the feasibility of integrating a RF front-end onto smartphones. Our results show that the GaN-CMOS process can realize a 802.11p front-end within the stringent power and area budget of a mobile phone.



▲ Figure 1: Process integration of GaN and CMOS: (a) Fabricated Si devices; (b) Si CMOS/GaN-on-Si wafer realized by two-step bonding technology; (c) GaN window open and device isolation; (d) Monolithically integrated GaN HEMT devices with final metal interconnection of fabricated HEMTs and Si CMOS devices.



▲ Figure 2: (a) Snapshot of the system prototype; (b) Average power consumption of RoadRunner V2V token exchanges with and without adaptive power control.

FURTHER READING

- US Department of Transportation, "The connected vehicle test bed," 2013, www.its.dot.gov/factsheets/connected_vehicle_testbed_factsheet.htm.
- P. Choi, C. C. Boon, M. Mao, and H. Liu, "28.8 dBm, high efficiency, linear GaN power amplifier with in-phase power combining for IEEE 802.11p applications," *IEEE Microwave Wireless Component Letters*, vol. 23, no. 8, pp. 433–435, Aug. 2013.
- J. Gao and L. Peh, "RoadRunner: Infrastructure-less vehicular congestion control," *CSAIL Technical Report*, Feb. 2014.

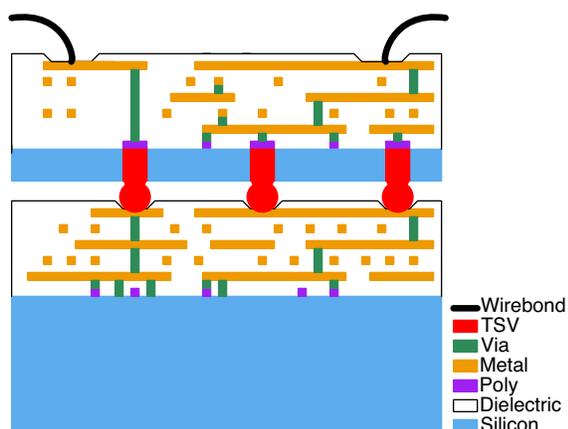
Impact of Digital Clock Tree on LC-VCO Performance in 3D-IC

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

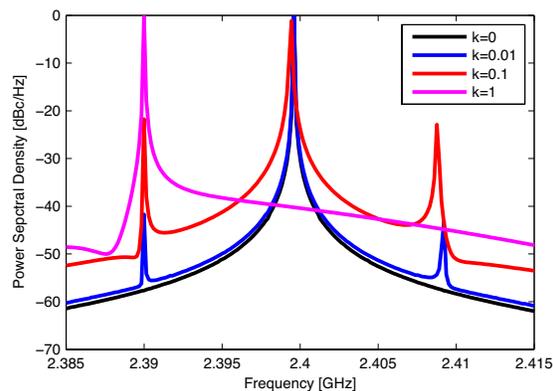
As we approach the physical limits of Moore's law and device scaling, there is increasing interest in other directions to achieve higher performance. Functional diversification, dubbed "More-than-Moore," calls us to bring together domains that have traditionally been separated. One of the suggested directions is three-dimensional integrated circuits (3D-IC), allowing the integration of several dies vertically stacked together and connected via through silicon vias (TSV) as shown in Figure 1. This opens up new possibilities for higher levels of integration to create more versatile and robust Systems-in-Package (SiP). Research is currently being done on integration of such elements as logic, memory, RF, power and sensors, showing the potential to reduce area, power, and cost and to increase data bandwidth. These savings can be achieved by the advent of a shorter interconnect with smaller parasitics and new topologies to utilize the three-dimensional structure. To enable these new technologies, we must overcome challenges caused by mechanical and thermal stress as well as power and signal integrity, requiring solutions at the circuit and system levels.

In this work we explore the impact of closely combining a sensitive analog circuit—an LC voltage

controlled oscillator (VCO) on one tier with a noisy digital clock tree on the other tier. Such a scheme would likely exist in any communication system consisting of the analog front-end leading to a digital processing back-end. We propose a method to obtain insight into the key parameters that affect coupling between clock lines and the VCO inductor structure. The analysis demonstrates the relation between the structure geometries and relative positions using partial inductance matrices. This analysis enables the design of topologies with minimal coupling. We also demonstrate the impact of such noise coupling on the performance of the VCO as manifested in its output spectrum and its phase noise. A simulation of the output power spectrum of a VCO is shown in Figure 2 for several coupling coefficients. As can be seen, the spurs increase by ~20 dB per decade of increased coupling. If the coupling is close and strong enough, we observe a pulling of the VCO frequency as a result of injection-locking. The techniques proposed to mitigate these effects will allow future development of more complex, highly integrated, front-end systems that bring together RF capability with digital signal processing.



▲ Figure 1: Schematic of Back-to-Face 3D-IC stack structure (not to scale).



▲ Figure 2: Simulation of an LC-VCO output power spectral density affected by various coupling levels from independent digital clock tree.

FURTHER READING

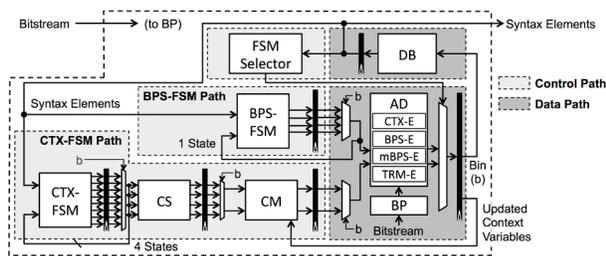
- W. Arden, M. Brillouët, P. Copez, M. Graef, B. Huizing, and R. Mahnkopf, "More-than-Moore white paper," International Technical Roadmap for Semiconductors, vol. 2, p. 14, 2010.
- D. H. Kim, K. Athikulwongse, M. Healy, M. Hossain, M. Jung, I. Khorosh, G. Kumar, Y.-J. Lee, et. al, "3D-MAPS: 3D Massively parallel processor with stacked memory," Int. Solid-State Circuits Conference, 2012, pp. 188–190.
- 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 J. Solid-State Circuits, vol. 46, no. 1, pp. 293–307, 2011.

A High-throughput CABAC Decoder Architecture for the Latest Video Coding Standard HEVC with Support of High-level Parallel Processing Tools

Y.-H. Chen, V. Sze
Sponsorship: MIT

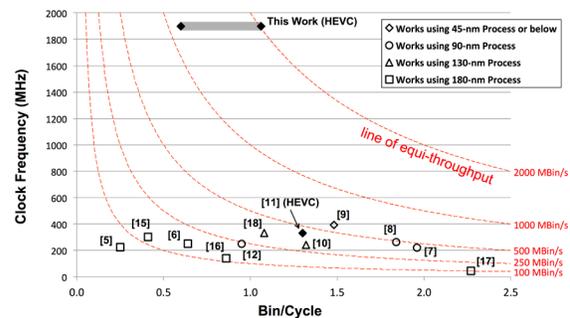
High Efficiency Video Coding (HEVC), developed by the Joint Collaborative Team on Video Coding (JCT-VC) as the latest video compression standard, was approved as an ITU-T/ISO standard in early 2013. Compared to its predecessor, the H.264/AVC standard, HEVC is designed to achieve 2× higher coding efficiency for resolutions up to 4320p (8K Ultra-HD) at 120 fps to support the next decade of video applications. This results in the high-throughput requirements for the context adaptive binary arithmetic coding (CABAC) entropy decoder, which was already a well-known bottleneck in H.264/AVC. To address the throughput challenges, several modifications were made to CABAC during the standardization of HEVC. This work leverages these improvements in the design of a high-throughput HEVC CABAC decoder. It also supports the high-level parallel processing tools introduced by HEVC, including tile and wavefront parallel processing.

The proposed design, as shown in Figure 1, uses a deeply pipelined architecture to achieve a high clock rate. Additional techniques such as the state prefetch



▲ Figure 1: Block diagram shows proposed HEVC CABAC decoder architecture. It consists of two finite state machines (CTX-FSM and BPS-FSM), FSM selector, context selector (CS), context memory (CM), bitstream parser (BP), arithmetic decoder (AD), and de-binarizer (DB). The architecture is deeply pipelined for high-throughput processing.

logic, latched-based context memory, and separate finite state machines are applied to minimize stall cycles, while multi-bypass-bin decoding is used to further increase the throughput. The design is synthesized in an IBM 45-nm SOI process. At the operating frequency of 1.9 GHz, it achieves throughputs up to 2014 and 2748 Mbin/s under common and theoretical worst-case test conditions, respectively. Figure 2 compares the performances of this design and previous works, including designs for both HEVC and H.264/AVC. The throughput advantage of this work comes from both the proposed architectural techniques and the advance in technology. This design is sufficient to decode in real-time high-tier video bitstreams at level 6.2 (8K Ultra-HD at 120 fps) or main-tier bitstreams at level 6.0 (8K Ultra-HD at 30 fps) for applications requiring sub-frame latency, such as video conferencing.



▲ Figure 2: Performance of the CABAC decoder is measured in bin/sec, which is the product of clock frequency and average number of decoded bins per cycle. Plot shows results under common test conditions. Compared to previous works in both HEVC and H.264/AVC standards, proposed design has clear throughput advantage.

FURTHER READING

- G. Sullivan, J. Ohm, T. K. Tan, and T. Wiegand, "Overview of the high efficiency video coding (HEVC) standard," *IEEE Transactions on Circuits and Systems for Video Technology (TCSVT)*, vol. 22, no. 12, pp. 1649–1668, Dec. 2012.
- D. Marpe, H. Schwarz, and T. Wiegand, "Context-based adaptive binary arithmetic coding in the H.264/AVC video compression standard," *IEEE Transactions on Circuits and Systems for Video Technology (TCSVT)*, vol. 13, no. 7, pp. 620–636, July 2003.
- V. Sze and M. Budagavi, "High throughput CABAC entropy coding in HEVC," *IEEE Transactions on Circuits and Systems for Video Technology (TCSVT)*, vol. 22, no. 12, pp. 1778–1791, Dec. 2012.

Full HD Integrated Video Encoder for H.265/HEVC

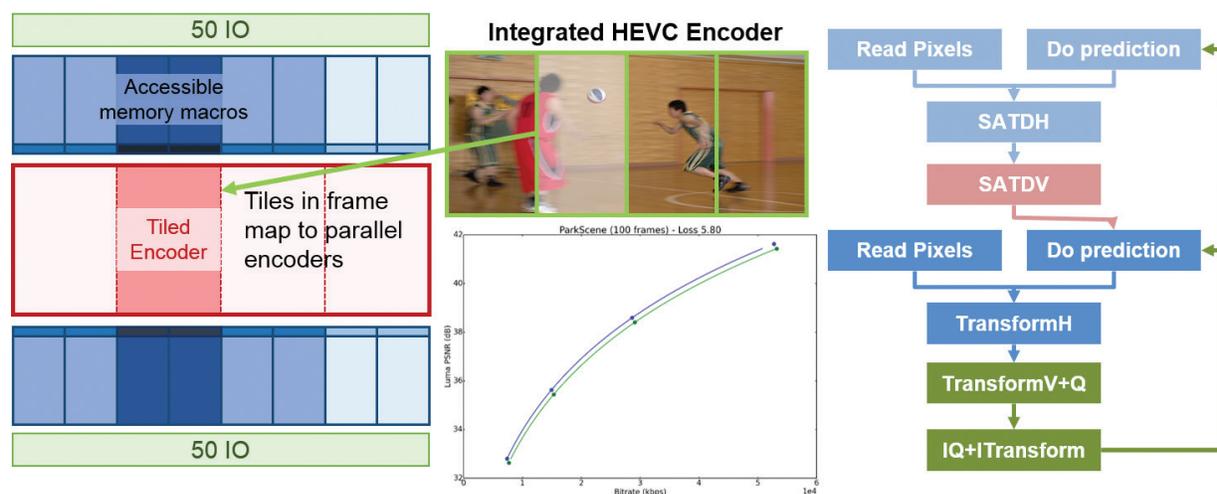
M. Tikekar, C. Juvekar, A.P. Chandrakasan
Sponsorship: Texas Instruments

High-efficiency video coding (HEVC), the latest video standard, uses larger and variable-sized coding units and longer interpolation filters than H.264/AVC to better exploit redundancy in video signals. These algorithmic techniques enable a 50% decrease in bitrate at the cost of increased computational complexity and external memory bandwidth. The added complexity makes building a real-time ASIC encoder a challenging problem. Our work attacks this problem by leveraging new modes of parallelism in HEVC using novel algorithms that are co-designed with the hardware architecture in mind. We target security camera applications where low cost integrated encoders could provide high quality video and still image archival.

Our previous work with HEVC decoders showed that the motion-compensation external memory power is a significant component of the total system's power budget. Since the motion estimation bandwidth for an encoder is typically much greater, we use an on-chip

frame buffer to save power. Use of on-chip memory allows for up to 10x bandwidth with 1/10th the power. The disadvantage is the much lower densities. We use the above features of on-chip storage through a 1-frame tiled motion estimator. Tiling is a frame-level parallelism tool introduced in HEVC that allows us to split the frame into rectangular tiles and then encode them independently. Tiles allow us to perform parallel motion estimation on different parts of the frame, making optimum use of the on-chip frame storage.

Intra estimation in HEVC is an equally challenging problem. Coding gain for intra estimation is typically much harder to realize without performing a very computationally intensive optimization. We tackle this problem by developing a new gradient-based intra algorithm that offers 6% coding loss and a 2.5x runtime improvement. We obtain a 2.5% coding gain, preserving the reconstructed pixel feedback in intra estimation by multi-threading across tiles.



▲ Figure 1: Tile architecture for HEVC encoder with 16 on-chip macros to store decoded picture buffer with support of 4 parallel motion estimation engines. Highlighted macros are accessible from current estimation engine.

FURTHER READING

- G. Sullivan, J. Ohm-Raimier, W.-J. Han, and T. Wiegand, *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 22, no. 12, pp. 1649-1668, 2012.
- C.-T. Huang, M. Tikekar, C. Juvekar, V. Sze, and A. Chandrakasan, "A 249MPixel/s HEVC video-decoder chip for Quad Full HD applications," *Digest of Technical Papers IEEE International Solid-State Circuits Conference (ISSCC)*, pp. 162-163, Feb. 2013.
- M. E. Sinangil, A. P. Chandrakasan, V. Sze, and M. Zhou, "Hardware-aware motion estimation search algorithm development for high-efficiency video coding (HEVC) standard," *Proceedings of the 19th IEEE International Conference on Image Processing (ICIP)*, pp. 1529-1532, 2012.

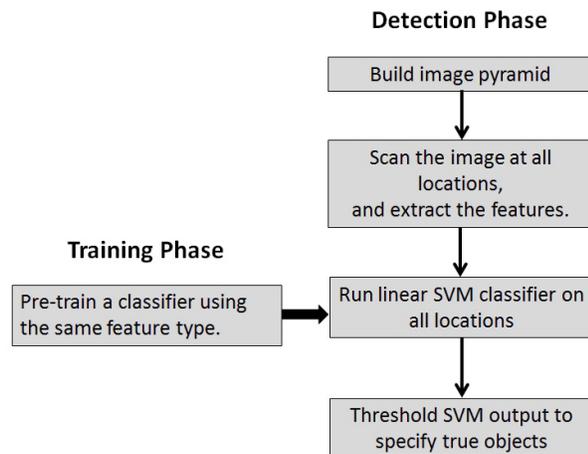
Energy-efficient Hardware for Object Detection

A. Suleiman, V. Sze

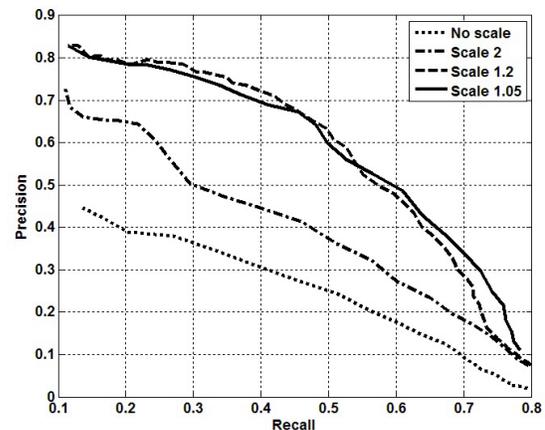
Object detection is needed in many embedded vision applications including surveillance, advanced driver assistance systems (ADAS), portable electronics, and robotics. Requirements for these applications include real-time operation, high-resolution image processing, and energy-efficiency, along with accuracy and robustness. For instance, real-time operation with low latency is necessary for applications such as ADAS and autonomous control in unmanned aircraft vehicles (UAV) to enable faster detection and allow more time for course corrections. High throughput (high frame rate) is also essential for fast reactions to 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 limited energy available in the battery.

This project aims to develop an efficient implementation of the Histogram of Oriented Gradients (HOG) -based object detection that addresses the

throughput and energy requirements, without much degradation in detection accuracy and robustness. This implementation can be achieved through both architectural and algorithmic optimizations. Hardware for object detection faces many challenges, from both complexity and memory points of view. In HOG-based object detection, HOG features are extracted for the entire frame, which is challenging for high-resolution images (e.g., 1080p) under low power constraints. The detection is done using a linear support vector machine (SVM) classifier, which includes a vector dot product requiring large number of multiplications per feature. Another important aspect of object detection is the multi-scale processing, where detection is done on multiple resolutions per frame to detect variably sized objects. This significantly increases the number of pixels being processed per frame, adding complexity to the design. One of the main challenges is to efficiently handle the large amount of computations on pixels and features while adhering to the low area and power requirements of embedded applications.



▲ Figure 1: Flow chart showing the basic steps in object detection algorithms.



▲ Figure 2: Precision-Recall curves for pedestrian detection with HOG using different scaling factors. A scale factor of 1.05 (used in original HOG paper) gives high precision, while not using scales significantly degrades the detection accuracy.

FURTHER READING

- P. Dollar, C. Wojek, B. Schiele and P. Perona, "Pedestrian Detection: An Evaluation of the State of the Art," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 34, no. 4, pp. 743-761, Apr. 2012.
- N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," presented at *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2005.
- S. Bauer, S. Kohler, K. Doll and U. Brunsmann, "FPGA-GPU architecture for kernel SVM pedestrian detection," presented at *IEEE Computer Vision and Pattern Recognition Workshop*, 2010.
- M. Hahnle, F. Saxen, M. Hisung, U. Brunsmann and K. Doll, "FPGA-Based Real-Time Pedestrian Detection on High-Resolution Images," presented at *IEEE Conf. on Computer Vision and Pattern Recognition*, 2013.
- K. Takagi, K. Mizuno, S. Izumi, H. Kawaguchi and M. Yoshimoto, "A sub-100-milliwatt dual-core HOG accelerator VLSI for real-time multiple object detection," presented at *IEEE Inter. Conf. on Acoustics, Speech and Signal Processing* 2013.

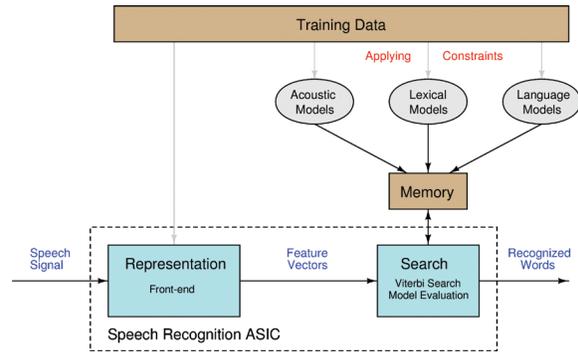
Algorithm and Architecture Enhancements for Hardware Speech Recognition

M. Price, J. Glass, A.P. Chandrakasan
Sponsorship: Quanta Computer, Inc. (Qmulus Project)

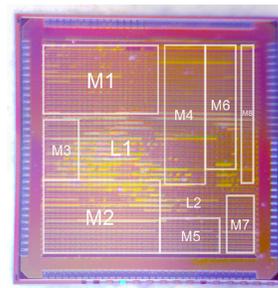
We are developing digital circuits to perform speech recognition within low-power embedded systems. The simplicity, energy efficiency, and scalability of these systems will allow them to substitute for cloud-based speech recognition services in scenarios where Internet connectivity is slow, unreliable, or energetically expensive. The circuits need to recognize speech with a controllable loss of accuracy relative to state-of-the-art software at the desired level of system power consumption.

We have designed an end-to-end (audio in, text out) speech recognition chip that is programmable with industry-standard WFST and GMM models. The chip includes a front-end that transforms audio into the feature representation used by the models and a Viterbi search module that formulates hypotheses based on these features. The weighted finite-state transducer (WFST) is a graph structure providing information about transitions between hidden states in the hidden Markov model (HMM) framework. A Gaussian mixture model (GMM) approximates acoustic observation probabilities.

Evaluating these models takes the bulk of computation and I/O for the chip and is hence the primary target for algorithm and architecture enhancements. GMM memory bandwidth is reduced by a factor of 55 through a combination of caching and parameter compression via Lloyd-Max quantization. A specialized cache for WFST parameters makes memory access more sequential, reducing page access penalties by a factor of 3. A feedback scheme is used to adjust the search pruning threshold dynamically, preventing overflow of limited on-chip memory while accommodating natural variations in search complexity. These techniques were implemented in a 65-nm test chip, which performs real-time speech decoding with a 5,000 word vocabulary. A word error rate of 13.0% is obtained at 50 MHz, with 0.85 V supplies and 6-mW core power consumption.



▲ Figure 1: Architecture of speech recognition system with hardware-accelerated decoder.



Logic regions	Memories
L1: Decoder	M1: Active state list 1
L2: Frontend	M2: Active state list 2
	M3: GMM quantization tables and cache
	M4: WFST cache data table
	M5: Feature vector buffer
	M6: WFST cache hash table
	M7: Feature and audio log
	M8: Frontend and FFT scratch memories

Specification	Value
Process	TSMC 65 nm
Die size	2.5 x 2.5 mm
Package	128-pin LQFP
Logic gates	340k (NAND2 equiv.)
SRAM	2.4 Mb
Supply voltage	0.8 – 1.2 V
Power consumption	5 – 23 mW (core)
Clock frequency	40 – 110 MHz
# WFST states	2.9M
# GMMs	4k x 32 Gaussians

▲ Figure 2: Die photo and summarized specifications of speech recognition chip.

FURTHER READING

- M. Price, J. Glass, and A. Chandrakasan, "A 6mW 5K-Word Real-Time Speech Recognizer using WFST Models," in *2014 International Solid-State Circuits Conference Digest of Technical Papers*, vol. 57, pp. 454-455, Feb. 2014.
- G. He, T. Sugahara, S. Izumi, H. Kawaguchi and M. Yoshimoto, "A 40-nm 54-mW 3x-real-time VLSI processor for 60-kWord continuous speech recognition," in *Proc. 2013 IEEE Workshop on Signal Processing Systems*, pp. 147-152, Oct. 2013.
- I. L. Hetherington, "PocketSUMMIT: Small-Footprint Continuous Speech Recognition," in *Proc. Interspeech 2007*, pp. 1465-1468, 2007.

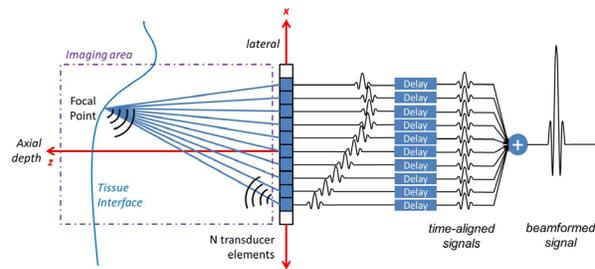
Demonstration Platform for Energy-scalable Ultrasound Beamforming

B. Lam, A.P. Chandrakasan
Sponsorship: SRC/FCRP C2S2, Texas Instruments

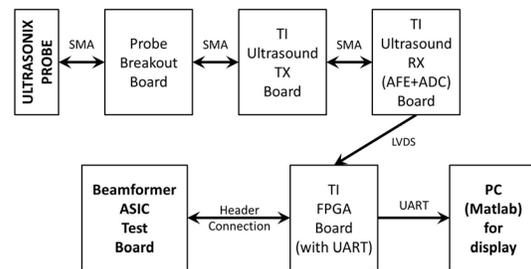
Conventional ultrasound systems employ large arrays of transducer elements to transmit and receive ultrasound waves for image formation. Given the large number of parallel channels, the cables connecting the ultrasound probe to the front-end electronics and back-end processing are necessarily bulky and expensive. With the number of channels supported only to increase (especially with the move to three-dimensional imaging, which requires the use of two-dimensional transducer arrays), miniaturizing the supporting electronics and minimizing the associated power consumption become important next steps in developing next-generation ultrasound imaging systems.

To address this need, previous work in MTL has shown low-power ASIC solutions for analog front-end (AFE) and analog-to-digital converter (ADC) electronics. We designed and fabricated a digital beamforming ASIC using TSMC 65-nm process to demonstrate energy-scalable operation in the digital domain. These three components can potentially comprise a small form factor, low-power, end-to-end solution for two-dimensional beamforming. The beamforming ASIC iteratively processes data from groups of eight transducer channels to form ultrasound images with varying image quality; we analyze the tradeoffs of this scalability.

The system demonstration uses several commercial components in conjunction with the beamforming ASIC to achieve energy scalable two-dimensional imaging. A 128-channel linear transducer probe provided by Ultrasonix Ltd. provides both the transmit and receive paths for analog signals, with the output echo waveforms being iteratively supplied to an eight-channel Texas Instruments AFE and ADC chip. The LVDS (digitized) output of this ASIC is then provided as input to the digital beamforming ASIC, the output of which is sent to a PC via UART for display. Three user-controlled modes of operation allow for a scaling of energy consumption against performance in terms of image quality, with the lower-quality images being used for large feature identification and the high-quality image used for detailed diagnosis.



▲ Figure 1: Graphical representation of beamforming operation.



▲ Figure 2: Block diagram of system demonstration components.

FURTHER READING

- K. Chen, A. Chandrakasan, and C. Sodini, "Ultrasonic imaging front-end design for cmut: A 3-leel 30vpp pulse-shaping pulser with improved efficiency and a noise-optimized receiver," *Asian Solid-State Circuits Conference, IEEE, 2012*.
- S. Lee, A. Chandrakasan, and H. Lee, "A 12 b 5-to-50 ms/s 0.5-1 v voltage scalable zero-crossing based pipeline adc," *IEEE Journal of Solid-State Circuits*, vol. 47, no. 7, pp. 1603-1614, 2012.
- K. Chen, B. Lam, C. G. Sodini, and A. P. Chandrakasan, "System Energy Model for a Digital Ultrasound Beamformer with Image Quality Control," *International Ultrasonics Symposium (IUS), IEEE, 2012*.

Self-powered Long-range Wireless Microsensors for Industrial Applications

N. Ickes, H. Goktas, P. Iannucci, A. Paidimarri, X. Wang, F. Yaul, H. Balakrishnan, K.K. Gleason, A.P. Chandrakasan
Sponsorship: Shell, Texas Instruments, National Science Foundation

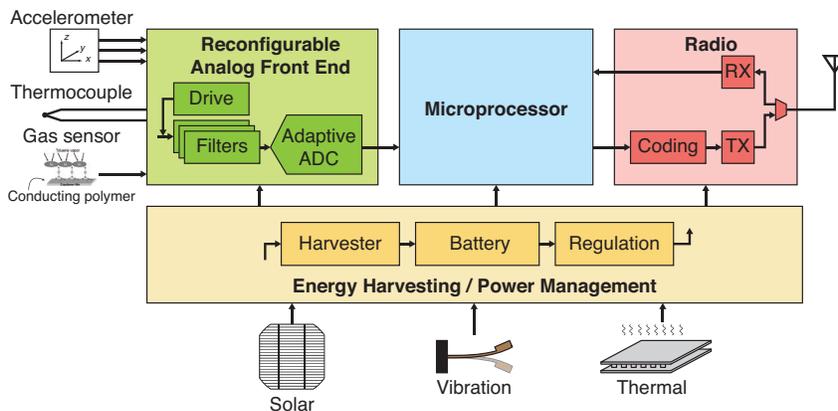
Improved monitoring of industrial equipment through the use of vast networks of small, easily installed, long-lasting, reliable wireless sensor nodes has the potential to increase worker safety, decrease down time, and even minimize unnecessary preventative maintenance. While critical systems are already likely to be highly instrumented and monitored, a large, modern industrial installation (such as an oil refinery, for example) may contain thousands of pumps, fans, motors, and other ancillary equipment that is monitored only through periodic, manual inspections. We are developing wireless microsensors to improve and automate the monitoring of this balance-of-plant equipment.

Figure 1 shows a block diagram of the planned sensor nodes. A reconfigurable analog front end will interface to the sensors themselves. Mechanical failures often develop slowly and can be detected early from abnormal vibrations and temperatures or minor gas leaks. Therefore, while the front end will be adaptable to a variety of sensor types, we will primarily focus on accelerometers and thermocouples, as well as a new

conductive polymer sensor for organic vapors, which is also being developed for this project.

Designing an efficient radio and network protocol is another key effort of this project. The target applications will require long-range (at least 100m) communication between the sensor nodes and the base station and high densities (up to 10,000 nodes per base station). We are investigating ways to improve existing protocols by reducing synchronization requirements and incorporating new coding techniques (such as spinal codes), to shift more of the power consumption to the base station, where energy is less constrained.

Once deployed, the sensors must function for up to twenty years without maintenance, so each node will be entirely powered by energy scavenged from its environment. Solar power is an attractive source for sensors deployed outdoors, but we are also investigating vibration and thermal harvesting, as these energy sources are readily available in many industrial applications and would allow sensors to be deployed in much less accessible locations.



◀ Figure 1: Block diagram of the planned sensor node. Optimized, energy-efficient sensor front end and radio components will allow the entire node to be powered from ambient energy harvested from solar, vibration, or thermal sources.

FURTHER READING

- Yaul, F. M., A. P. Chandrakasan, "A 10b 0.6nW SAR ADC with Data-Dependent Energy Savings Using LSB-First Successive Approximation," in *IEEE International Solid State Circuits Conference (ISSCC)*, 2014, pp. 198–199.
- Angelopoulos, G., A. Paidimarri, A. P. Chandrakasan, M. Medard, "Experimental Study of the Interplay of Channel and Network Coding in Low Power Sensor Applications," in *IEEE International Conference on Communications (ICC)*, 2013, pp. 5126–5130.
- Paidimarri, A., P. Nadeau, P. Mercier, A. Chandrakasan, "A 440pJ/bit 1Mb/s 2.4GHz Multi-Channel FBAR-based TX and an Integrated Pulse-shaping PA," in *IEEE Symposium on VLSI Circuits (VLSIC)*, 2012, pp. 34–35.

Graphene-CMOS Hybrid Infrared Image Sensor

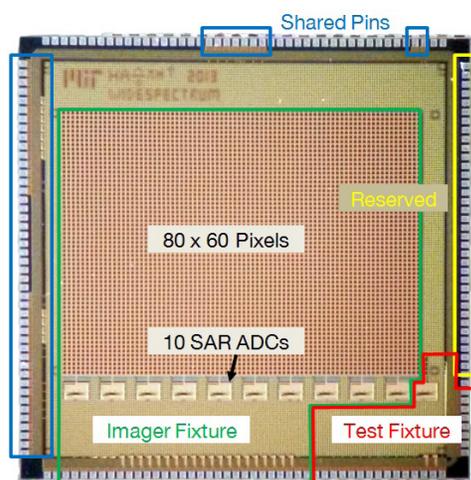
S. Ha, A. Hsu, T. Palacios, A.P. Chandrakasan

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

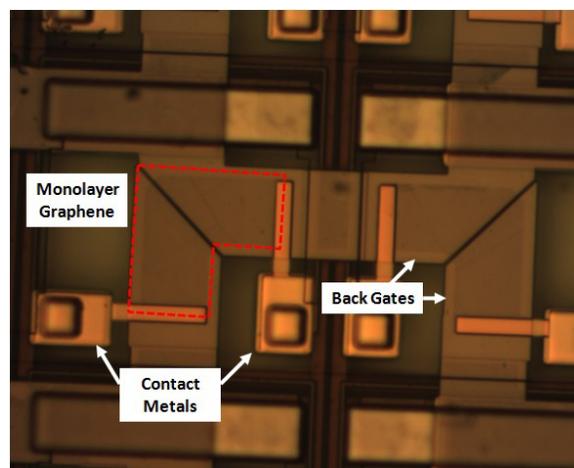
Applications in the mid- and far-infrared spectrum expand from security cameras and medical thermal imaging to spectroscopy for chemistry and astronomy. However, conventional infrared (IR) sensors require high-cost epitaxial growth of InSb or HgCdTe layer for IR absorption, and they also suffer from integration issues and a limited absorption band determined by the epitaxial growth. On the other hand, CMOS image sensors are widely used in digital multimedia applications thanks to their mature production technology, and 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 ($\lambda < 1100\text{nm}$).

We develop a graphene-CMOS hybrid sensing platform that 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. This work presents a prototype chip, with graphene thermocouples fabricated on top of the CMOS IC.

A 5mm x 5mm readout chip is fabricated as shown in Figure 1. The chip is fabricated using commercial 0.18- μm technology, but the design considered post-fabrication steps for graphene IR detectors. Each pixel has two contact metal plugs towards the top surface of the chip for ohmic contact to graphene. The pixel area is $50\mu\text{m} \times 50\mu\text{m}$; a pixel amplifier and signal paths occupy the small portion of each pixel area, leaving over 60% of the area empty and flat. The center area of the chip, 3mm x 4mm, is filled with the 80 x 60 array of the pixels. Sides of the chip are used for the row and column selection logic, current sources, column amplifiers, and ADCs. The layout eases the post-fabrication process, including graphene transfer, by locating the regularly patterned pixel array in the center of the chip and others on the sides. Figure 2 shows the graphene thermocouple fabricated on the chip. Each end of the thermocouple connects to the metal pillar of the pixel amplifier input. An induced p-n junction of graphene generates a thermovoltic signal in response to IR light; the readout IC amplifies and digitizes it to 8-bit code. Ten parallel ADCs can process data from 4,800 pixels at a rate of 1MB/s.



▲ Figure 1: The readout IC for graphene-CMOS hybrid IR imager is fabricated. The chip is shown before the post-fabrication of graphene thermocouples.



▲ 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 underneath.

FURTHER READING

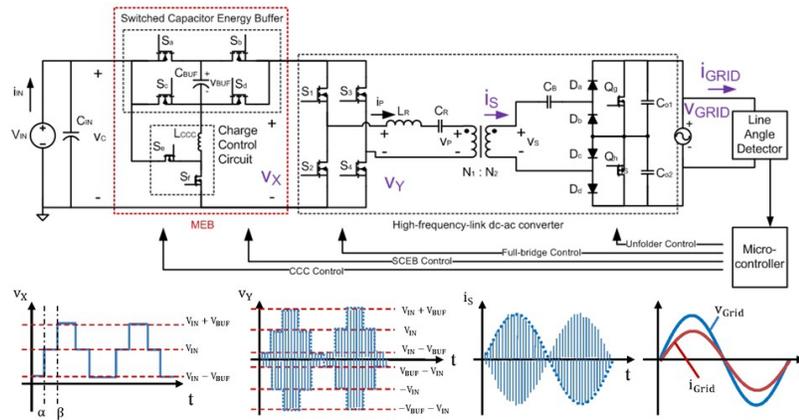
- A. Fish and O. Yadid-Pecht, "Low Power CMOS Imager Circuits," in *Circuits at the Nanoscale: Communications, Imaging, and Sensing*, K. Iniewski Ed. Boca Raton: CRC Press/Taylor & Francis, 2009.
- P. K. Herring, A. Hsu, N. M. Gabor, Y. C. Shin, J. Kong, T. Palacios, and P. Jarillo-Herrero. "Photoresponse of an Electrically Tunable Ambipolar Graphene Infrared Thermocouple." *Nano letters*, vol. 14, no. 2, pp. 901-907, Jan. 2014.
- M. Perenzoni, N. Massari, D. Stoppa, L. Pancheri, M. Malfatti, and L. Gonzo, "A 160x120-Pixels Range Camera with In-Pixel Correlated Double Sampling and Fixed-Pattern Noise Correction," *IEEE Journal of Solid-State Circuits*, vol. 46, pp. 1672-1681, 2011.

A Multilevel Energy Buffer and Voltage Modulator for Grid-interfaced Micro-inverters

M. Chen, K.K. Afridi, D.J. Perreault
Sponsorship: Enphase Energy

Micro-inverters operating in the single-phase grid from solar photovoltaic (PV) panels or other low-voltage sources must buffer the twice-line-frequency variations between the energy sourced by the PV panel and that required for the grid. Moreover, in addition to operating over wide average power ranges, they inherently operate over a wide range of voltage conversion ratios as the line voltage traverses a cycle. These factors make the design of micro-inverters challenging. This paper presents a Multilevel Energy Buffer and Voltage Modulator (MEB) that significantly reduces the range of voltage conversion ratios that the dc-ac converter portion of the microinverter must operate over by stepping its effective input voltage in pace with the line voltage. The MEB partially replaces the original bulk input capacitor and functions as an active energy buf-

fer to reduce the total size of the twice-line-frequency energy buffering capacitance. The small additional loss of the MEB can be compensated for by the improved efficiency of the dc-ac converter stage, leading to a higher overall system efficiency. The MEB architecture can be implemented in a variety of manners, allowing different design tradeoffs to be made. A prototype micro-inverter incorporating an MEB, designed for 27 V to 38 V dc input voltage, 230 V rms ac output voltage, and rated for a line cycle average power of 70 W, has been built and tested in a grid-connected mode. We show that the MEB can successfully enhance the performance of a single-phase grid-interfaced microinverter by increasing its efficiency and reducing the total size of the twice-line-frequency energy buffering capacitance.



◀ Figure 1: One implementation of the MEB micro-inverter and its conceptual operating waveforms.

▶ Figure 2: Pictures of the prototype MEB based micro-inverter. The layout of the MEB stage (switches and capacitors) is optimized, targeting the highest power density.



FURTHER READING

- M. Chen, K. K. Afridi, and D. J. Perreault, "A multilevel energy buffer and voltage modulator for grid-interfaced micro-inverters," IEEE Transactions on Power Electronics, accepted, 2014.
- M. Chen, K. K. Afridi, and D. J. Perreault, "A Multilevel Energy Buffer and Voltage Modulator for Grid-Interfaced Micro-Inverter," in Proc. IEEE Energy Conversion Congress and Exposition (ECCE), Denver, CO, Sept. 2013.
- M. Chen, K. K. Afridi, and D. J. Perreault, "Stacked Switched Capacitor Energy Buffer Architecture," IEEE Transactions on Power Electronics, vol. 28, no. 11, pp. 5183-5195, Nov. 2013.

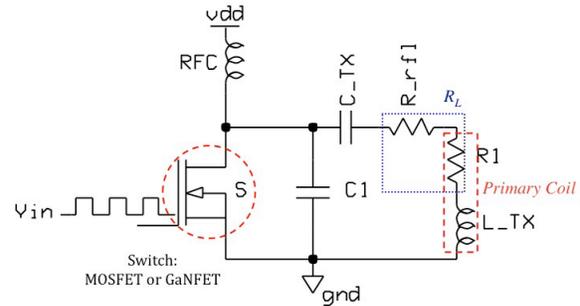
Efficient Wireless Charging with Gallium Nitride FETs

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

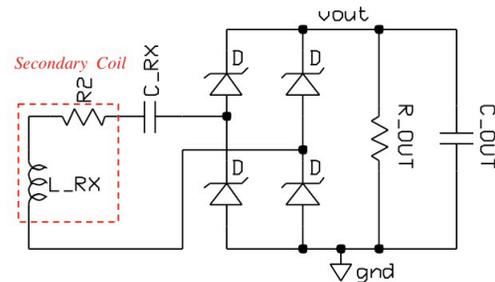
Wireless transfer of power and signals over distances has been rapidly growing. The ability to transfer energy from one device to another without cables improves both the convenience and reliability for consumers. This research focuses on near-field power transfer for ¼ inch up to 1 inch for a device people use in their daily lives: cellphones.

Though wireless charging is more convenient than traditional wired charging methods, it is currently less efficient. This method not only wastes power but also can result in a longer charging time. Reducing the sources of loss in the conversion circuits improves the efficiency of the wireless charging system. In this work, we focus on losses originating from the transistor. We designed and implemented resonant inductive wireless charging systems with different switch implementations to compare efficiencies.

One system utilizes the traditional silicon MOSFET. The other board replaces the MOSFET with a gallium nitride FET (GaNFET), while keeping the circuitry and components as consistent as possible. GaNFETs have many benefits such as lower $R_{ds,on} \times Q_G$, smaller footprints, and potentially higher breakdown voltages. Results show that the GaNFET system has a 5% efficiency gain over the MOSFET system for various distances. The wireless charging systems implemented allow for flexibility in alignment between devices delivering and receiving power and efficiencies in the 30%-50% range.



▲ Figure 1: Transmitter unit. Class E power amplifier used as the DC/RF converter. RF choke provides DC current. S switch is fully either on or off. C1 is shunt capacitor to provide zero voltage turn on. CTX and LTX form the primary tank. LTX is inductance of primary coil; R1 is intrinsic resistance of primary coil.



▲ Figure 2: Secondary unit. A secondary LC tank with full bridge rectifier. LRX is inductance of secondary coil and R2 is intrinsic resistance of secondary coil. CRX adds resonance to circuit for optimal power transfer.

FURTHER READING

- N. Sokal "Class-E RF Power Amplifiers," in *QEX Commun*, 2001.
- T. Lee, *The Design of CMOS Radio-Frequency Integrated Circuits*, Cambridge, UK: Cambridge University Press, 2004.
- R. Sharpeskar, *Ultra Low Power Bioelectronics*, Cambridge, UK: Cambridge University Press, 2010.

Efficient Portable-to-Portable Wireless Charging

R. Jin, A.P. Chandrakasan
Sponsorship: Foxconn, TSMC University Shuttle Program

In today's world of ever-increasing low-power portable electronics, from medical implants to wireless accessories, powering these devices efficiently and conveniently is an growing concern. Currently, these devices are recharged by plugging them into individual chargers. This can be an inconvenience for users and contributes to a large amount of electronic waste. Our alternative solution is to wirelessly recharge these lower-power portable devices through a common magnetic link with a higher-power portable device, such as a smartphone. In a typical case, the user moves the smartphone close to the portable device and charges it in a few minutes for a day's use. Such a method is convenient, environmentally friendly, and cheap to implement.

This portable-to-portable wireless charging application differs from conventional charging pad-based systems in that the transmitter battery life is constrained and valuable, so efficiency is key. Also, since

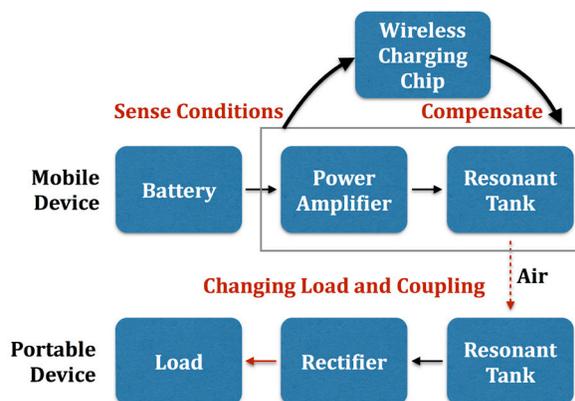
both the transmitter and receiver are portable devices, output load and transmitter-to-receiver coupling are constantly changing, which results in dynamic transmitter loading. These changing conditions affect transfer efficiency.

We develop a resonant inductive wireless charging system operating at 6.78 MHz that transfers energy between portable devices with high efficiency. The system includes a custom integrated circuit that senses changing load and coupling conditions while charging and actively adjusts the transmitter circuit to maintain high system efficiency and consistent power levels.

This portable-to-portable wireless charging system is applicable to many kinds of devices. Applications are demonstrated that use a smartphone to wirelessly recharge fitness trackers, cochlear implants, bicycle lights, MP3 players, wireless keyboards, and calculators, charging most devices in 2 minutes for a typical day's use.



▲ Figure 1: Our wireless charging system uses a smartphone to recharge various lower-power portable devices, such as cochlear implants, wireless keyboards, bicycle lights, fitness trackers, and calculators. The smartphone charges most devices in 2 minutes for a typical day's use.



▲ Figure 2: Changing load and coupling conditions impact the transfer efficiency of the wireless charging system. A wireless charging controller chip detects these conditions and actively adjusts the transmitter to compensate. This maintains high efficiency and consistent power levels as conditions change.

FURTHER READING

- M. Yip, R. Jin, H. H. Nakajima, K. M. Stankovic, and A. P. Chandrakasan, "A fully-implantable cochlear implant SoC with piezoelectric middle ear sensor and energy-efficient stimulation in 0.18 um HVCMOS," in *International Solid-State Circuits Conference*, 2014, pp. 312-313.
- N. Desai, "A low-power, reconfigurable fabric body area network for healthcare applications," Master's thesis, Massachusetts Institute of Technology, Cambridge, 2012.

A 28-nm FDSOI Integrated, Reconfigurable, Switched-Capacitor Step-up DC-DC Converter with 88% Peak Efficiency

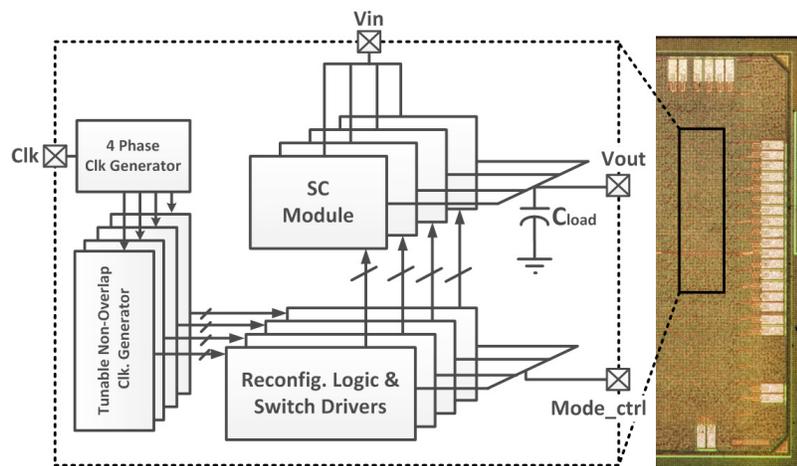
A. Biswas, Y. Sinangil, A.P. Chandrakasan
Sponsorship: DARPA, STMicroelectronics

The increasing integration of analog, digital and RF circuits in modern Systems-on-Chip (SoCs) has created a demand for a wide range of unique power supplies, to cater to different functionalities. Hence, an on-chip power management unit (PMU) is essential to efficiently convert and deliver these diverse power supplies from a single source. With the progress of CMOS scaling, the nominal supply voltage (V_{dd}) of the transistors has substantially decreased. However, certain functionalities, e.g., non-volatile memory, require voltages that are higher than V_{dd} . On the other hand, applications like energy-harvesting need to boost the source voltage to generate a higher output voltage. Thus, step-up DC-DC converters are an important component in the PMU for these kinds of applications. Fully integrated switched-capacitor- (SC) based DC-DC converters can achieve high conversion efficiency and power density, which are key for on-chip implementation.

To benefit from CMOS scaling, SC converters should utilize core transistors as charge-transfer switches. Core transistors offer lower on-resistance (R_{on}) and capacitance compared to I/O transistors.

However, to avoid voltage overstress, core transistors cannot be operated with a gate-to-source/drain voltage of more than V_{dd} . Furthermore, it is desirable to have reconfigurability in the SC converter. It enables the same converter to be efficiently used to generate a wide range of output voltages, rather than using separate converters for each output voltage.

In this work we implement a reconfigurable step-up SC DC-DC converter with 3 conversion ratios of 5/2, 2/1, and 3/2. This converter provides a wide range of output voltage from 1.2V to 2.4V, with a fixed input supply voltage of 1V. The step-up converter has been designed to obviate the need to use high voltage I/O transistors as charge-transfer switches. Additionally, a new topology is proposed for the 5/2 mode, and it improves conversion efficiency by reducing the bottom-plate parasitic loss as compared to a conventional series-parallel topology. The converter (Figure 1) was implemented in a 28-nm FDSOI (fully-depleted SOI) process using only on-chip MOS and MOM (metal fringe) capacitors that require no extra fabrication steps, unlike MIM (metal-insulator-metal) and trench capacitors.



◀ Figure 1: Overall architecture of the integrated, reconfigurable, step-up switched-capacitor DC-DC converter.

FURTHER READING

- D. El-Damak, S. Bandyopadhyay, and A. P. Chandrakasan, "A 93% efficiency reconfigurable switched-capacitor DC-DC converter using on-chip ferroelectric capacitors," *IEEE International Solid-State Circuits Conference (ISSCC)*, pp. 374-375, Feb. 2013.
- Y. K. Ramadass and A. P. Chandrakasan, "Voltage Scalable Switched Capacitor DC-DC Converter for Ultra-Low-Power On-Chip Applications," *IEEE International Power Electronics Specialists Conference (PESC)*, pp. 2353-2359, June 2007.
- W. Jung, S. Oh, S. Bang, Y. Lee, D. Sylvester, and D. Blaauw, "A 3nW fully integrated energy harvester based on self-oscillating switched-capacitor DC-DC converter," *IEEE International Solid-State Circuits Conference (ISSCC)*, pp. 398-399, Feb. 2014.