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

Towards Real-Time Super-Resolution on Compressed Video	23
Low-Power Hardware Accelerator for Object Detection with Deformable Parts Model	24
Energy-Efficient Reconfigurable Accelerator for Deep Convolutional Neural Networks	25
Low-Power Security-Acceleration Core for the Internet of Things	26
Ultra Low-Power, High-Sensitivity Secure Wake-Up Transceiver for the Internet of Things	27
A Noise-Efficient Chopper Amplifier Using a 0.2V-Supply Input Stage	28
Efficiency Maximization for Wireless Charging	29
Fully Integrated Thermal Energy Harvesting System with 50mV Start-Up	30
0.3V Biopotential Sensor Interface for Stress Monitoring	31
12-bit, 300MS/s CMOS Pipelined Analog-to-Digital Converter	32
A CMOS Flash ADC for GaN/CMOS Hybrid Continuous-Time $\Delta\Sigma$ Modulator	33
High-Performance GaN HEMT Track-and-Hold Sampling Circuits	34
Broadband Inter-Chip Link Using a Terahertz Wave on a Dielectric Waveguide	35
A Fast, Wideband THz CMOS Spectrometer Based on Dual-Frequency Comb Architecture	36
High-Power 1-THz Source Based on a Scalable 2D Radiating Mesh	37
Design, Modeling, and Fabrication of Chemical Vapor Deposition-Grown MoS ₂ Circuits with E-Mode Field-Effect Transistors for Large-Area Electronics	38

Towards Real-Time Super-Resolution on Compressed Video

Z. Zhang, V. Sze Sponsorship: MIT

High-resolution displays are increasingly popular, calling for faster algorithms to upsample the existing low-resolution video content. State-of-the-art super-resolution algorithms mainly address the visual quality of the output instead of real-time throughput. Even the fastest existing super-resolution algorithm, SRCNN, takes 0.4 second to process a single frame of size 256×256. This speed is far behind the requirement of a real-time super-resolution system, which should be capable of processing 30 frames per second on full HD videos (1920×1080).

We propose a framework called Free Adaptive Super-Resolution via Transfer (FAST) to accelerate any image-based super-resolution algorithm running on compressed videos. FAST leverages the similarity between adjacent frames in a video. Given the output of a super-resolution algorithm on one frame, FAST adaptively transfers the super-resolution to the adjacent frame so that we can avoid running the super-resolution algorithm on the adjacent frame. The transferring process has negligible computation cost because the required information including motion vectors, block size, and prediction residual is already embedded in the compressed video for free.

FAST also adapts to video content, which is composed of frames with varying block sizes. It adaptively enables and disables transfer for each block depending on the quality of motion compensation of the video encoder. Note that the blocks are nonoverlapping so that the redundant computations for overlapping blocks in many existing super-resolution algorithms are avoided, which significantly reduce the complexity of the framework. The resulting artifacts are handled by low-complexity deblocking filters.

FAST was evaluated with existing state-of-theart super-resolution algorithms on the common test sequences that were used in the development of the HEVC video compression standard. FAST accelerates all the tested super-resolution algorithms by up to an order of magnitude with acceptable quality loss of up to 0.2 dB. This result proves that FAST can accelerate any super-resolution algorithm, potentially enable running super-resolution algorithms to upsample streamed videos for large screens in real time.



▲ Figure 1: Pipeline of FAST: From the compressed video (1), the video decoder decodes the low-resolution frames (2) and syntax elements (3). The SR algorithm is applied to the first frame to obtain a high-resolution output (4). FAST adaptively transfers it to the second frame (5).



▲ Figure 2: FAST result: Running SRCNN with FAST preserves the rich high-frequency details that SRCNN generates compared to the blurry output of bicubic interpolation.

- C. Dong, C. C. Loy, K. He, and X. Tang, "Learning a Deep Convolutional Network for Image Super-Resolution," European Conference on Computer Vision, New York: Springer, pp. 189-199, 2014.
- C. Liu and D. Sun, "A Bayesian Approach to Adaptive Video Super Resolution," in 2011 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 209-216, 2011.
- G. J. 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*, vol. 22, no. 12, pp. 1649-1668, December 2012.

Low-Power Hardware Accelerator for Object Detection with Deformable Parts Model

A. Suleiman, Z. Zhang, V. Sze Sponsorship: Texas Instruments, DARPA

While fully autonomous cars are still in development, the advanced driver assistant systems (ADAS) is being improved and becoming standard in many cars. Using computer vision to support autonomy is one of the cheapest solutions compared to more complex approaches like Lidar. Increasing object detection accuracy is essential for these types of applications. In our previous project, we developed an object detection accelerator based on histogram of oriented gradients (HOG) features; it supports 12 scales per frame for accuracy and robustness, processes 1080HD videos to detect far objects, runs at 60fps in real time; which gives time to react in fast changing environment, and consumes only 45.3mW.

In this project, we present an object detection hardware accelerator that has higher detection accuracy by using deformable parts models (DPM). The chip processes 1080HD videos in real-time and supports multi-scale detection with a throughput of 30fps consuming only 58.6mW, and it is tested to run up to 60fps. The DPM algorithm is an extension of the HOG-based detection; both use similar HOG features. In DPM, the detection is done on two levels as shown in Figure 1: root and parts. Using the root template, the object is searched for as a whole in the image pyramid with the conventional sliding window approach. To enhance the detection accuracy, 8 different templates are defined for 8 parts of the object and each part is searched for separately at 2x image resolution relative to the root. Eventually, the root score and the 8 parts scores are added together with a deformation penalty. Being able to detect object parts and allowing parts movements by deformation increases the detection accuracy significantly, Figure 2. For example, when DPM and HOG-based detection are compared, the detection accuracy is doubled on INRIA person dataset.



▲ Figure 1: Detection with bicycle DPM templates. 8 part templates are defined at 2x resolution relative to the root. Each part has a deformation penalty. The darker regions indicate lower deformation cost.

▲ Figure 2: Detection examples with different object classes (person, bicycle, horse, and aeroplane). Red boxes highlight the object position, while blue boxes highlight the optimal position for each of the 8 parts.

- A. Suleiman, Z. Zhang, and V. Sze, "A 58.6mW Real-Time Programmable Object Detector with Multi-Scale Multi-Object Support Using Deformable Parts Model on 1920x1080 Video at 30fps," presented at IEEE Symposium on VLSI Circuits, Honolulu, HI, June 2016.
- A. Suleiman and V. Sze, "Energy-Efficient HOG-based Object Detection at 1080HD 60 fps with Multi-Scale Support," IEEE Workshop on Signal Processing Systems, pp. 1-6, October 2014.
- P. F. FelzenszwalR. B. Girshick, D. McAllester, and D. Ramanan, "Object Detection with Discriminatively Trained Part-Based Models," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, pp. 1627-1645, September 2010.

Energy-Efficient Reconfigurable Accelerator for Deep Convolutional Neural Networks

Y.-H. Chen, T. Krishna, J. Emer, V. Sze Sponsorship: DARPA YFA, CICS, Intel

Deep learning using convolutional neural networks (CNNs) give unprecedented accuracy on many computer vision tasks, such as object detection, recognition, and segmentation. These state-of-the-art CNNs are orders of magnitude larger than the CNNs used in the 1990s, when CNNs were first introduced, requiring not only a larger number of layers but also millions of filter weights in varying shapes. For instance, AlexNet uses 2.3 million weights (4.6MB of storage) and requires 666 million MACs per 227x227 image (13 kMACs/pixel). VGG16 uses 14.7 million weights (29.4 MB of storage) and requires 15.3 billion MACs per 224x224 image (306 kMACs/pixel).

The size of these CNNs poses both throughput and energy efficiency challenges to the underlying processing hardware. Specifically, large number of weights, filters, and channels results in substantial data movement, which consumes significant energy. Today, CNN processing is carried out mostly in data centers or by high-end GPUs, where power dissipation often limits the achievable throughput. In the future, CNN processing will also be carried out on embedded devices rather than in the cloud, due to privacy, latency, or communication bandwidth concerns; these battery-powered devices will have even tighter energy and power constraints. Therefore, specialized CNN accelerators, which give higher throughput and improved energy-efficiency over general purpose platforms, will be critical for the implementation of future vision systems.

This work describes a CNN accelerator that can deliver state-of-the-art accuracy with minimum energy consumption in the system (including DRAM) in realtime, by using two key methods: (1) efficient dataflow and supporting hardware (spatial array, memory hierarchy, and on-chip network) that minimize data movement by exploiting data reuse and support different shapes and (2) exploitation of data statistics to minimize energy through zeros skipping/gating to avoid unnecessary reads and computations and data compression to reduce off-chip memory bandwidth, which is the most expensive data movement.



▲ Figure 1: The CNN compute pipeline consists of many layers, each of which performs convolution-like processing to extract abstract features in the image.

CNN	LeNet	AlexNet	VGG16
Year	1989	2012	2014
# of Convolution Layers	2	5	13
# of weights	50k	2.3M	14.7M
Ratio (memory)	1x	46x	294x
# MACs	322k	666M	15.3G
Ratio (Computation)	1x	2067x	47660x

AlexNet

Layer	Filter	# Filters	# of	Strido	
	Size (R)	(M)	Channels (C)	stride	
1	11x11	96	3	4	
2	5x5	256	48	1	
3	3x3	384	256	1	
4	3x3	384	192	1	
5	3x3	256	192	1	

▲ Figure 2: State-of-the-art deep CNNs require not only large number of layers but also millions of filter weights and varying shapes (i.e., filter sizes, # of filters, # of channels).

A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," Advances in Neural Information Processing Systems 25, pp. 1097-1105, 2012.

K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," CoRR, abs/1409.1556, 2014.

Low-Power Security-Acceleration Core for the Internet of Things

U. Banerjee, C. Juvekar, A. P. Chandrakasan Sponsorship: Analog Devices, Inc.

The Internet of Things (IoT) has introduced a vision of an Internet where all computing and sensing devices are interconnected. Digitally connected devices are encroaching on every aspect of our lives, including our homes, cars, offices, and even our bodies. Researchers estimate that there will be over 40 billion wireless connected devices by 2020. On one hand, the IoT enables fundamentally new applications, but on the other, these devices are attractive targets for cybercriminals, thus making IoT security a major concern. According to a report on the state of IoT security in 2015, 90% of IoT devices have collected personal data, but 70% of them used unencrypted network services.

Most commercial IoT transceivers either have no security implementations in hardware or only support symmetric key primitives like Advanced Encryption Standard (AES). To achieve end-to-end security in IoT networks, public key algorithms, like elliptic curve cryptography (ECC) are indispensable. Software implementations of these algorithms involve significant computational costs, and the power consumption presents a bottleneck in resourceconstrained environments. In this work, we propose to design low-power security-acceleration hardware that interfaces with a standard micro-processor and supports ECC for key exchange and digital signatures, along with standard cryptographic components like AES (Figure 1), thus alleviating the security and efficiency trade-off observed in embedded devices.

Our work also focuses on optimizing network security protocols for efficient implementation in embedded devices. Standard implementations of these protocols tend to have a large communication overhead, which becomes an additional concern for battery-powered or energy-harvesting IoT devices. Therefore, our proposed hardware can not only secure private data using low power cryptographic computations, but also reduce energy consumption of the RF transceiver (Figure 2).



▲ Figure 1: Security coprocessor for IoT to accelerate cryptographic primitives like AES, SHA and ECC. The hardware accelerator, which interfaces with a micro-processor, can be used to implement standard transport layer security protocols.



▲ Figure 2: Wireless sensor node using the proposed security coprocessor to send encrypted data all the way to the cloud. Application payloads are fully encrypted so that intermediate nodes can process and modify network addresses without gaining access to the actual data.

D. Miessler, "The State of IoT Security," 2015. [Online]. Available: http://community.hpe.com/t5/Protect-Your-Assets/The-State-of-IoT-Security-2015/ ba-p/6744413.

[•] P. Miranda, M. Siekkinen, and H. Waris, "TLS and Energy Consumption on a Mobile Device: A Measurement Study," *IEEE Symposium on Computers and Communications (ISCC)*, pp. 983-989, 2011.

Ultra Low-Power, High-Sensitivity Secure Wake-Up Transceiver for the Internet of Things

M. R. Abdelhamid, A. Paidimarri, A. P. Chandrakasan

Nanopower "Internet of Things" (IoT) devices deployed in short-range personal health devices, home automation systems and longer-range industrial monitoring systems all consume a large portion of energy on their wireless communication systems. However, long battery lifetimes or energy-harvested operation is desired in order to ease their adoption. In this work, we propose protocol optimizations for sensor-node driven communications. For base-station driven communication, we propose to achieve power reduction through an ultra-low-power wake-up receiver with optimizations in the protocols as well as the circuit design.

Wireless protocols such as Bluetooth low energy (BLE) are optimized for short-length packets with small preambles and reduced header sizes. However, low duty-cycle performance in the default connectedmode of operation is limited by periodic beacons and the requirement that the low-power sensor node absorbs the timing uncertainties and associated guard time intervals. As shown in Figure 1, the analysis of a commercial BLE radio performance shows that the average total power is much higher than the standby power of the radio, which presents opportunities for significant power reduction through protocol optimization. On the wake-up receiver chain, the design can exploit different trade-offs in the protocol to reach a sub- μ W average power consumption while maintaining the specifications by the basestation. Therefore, the sensitivity/power trade-off of the receiver can be mitigated through an optimized protocol and system. Additionally, the tremendous growth of IoT devices allows open communication among all sorts of devices. With such a huge amount of data flowing through the network, security becomes a critical issue. Hence, we propose the wakeup transceiver system shown in the block diagram of Figure 2. The transceiver incorporates a transmitter that provides small amounts of data sporadically upon request and creates a two-way communication channel for secure wake-ups and transmissions.



▲ Figure 1: Commercial BLE performance analysis.



FURTHER READING

C. Salazar, A. Kaiser, A. Cathelin, and J. Rabaey, "A -97dBm-sensitivity Interferer-resilient 2.4GHz Wake-up Receiver using Dual-IF Multi-N-Path
 Architecture in 65nm CMOS," presented at *IEEE International Solid-State Circuits Conference*, San Francisco, CA, February 2015.

O. Seunghyun, N. Roberts, and D. Wentzloff, "A 116nW Multi-band Wake-up Receiver with 31-bit Correlator and Interference Rejection," in Proc. 2013 IEEE Custom Integrated Circuits Conference (CICC), pp. 1-4, 22-25, 2013.

A Noise-Efficient Chopper Amplifier Using a 0.2V-Supply Input Stage

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

Wireless sensor nodes often monitor low-bandwidth signals from sensors that may be small in amplitude. Examples include biopotential signals for medical applications and vibration and strain signals for industrial monitoring applications. Typically, a low-noise instrumentation amplifier (LNIA) is used to provide a high-impedance sensor interface with good common-mode rejection (CMR). Electroencephalogram (EEG) monitoring is one application where LNIA designers have targeted sub-microvolt input-referred noise over a sub-100Hz signal band. Prior work to improve the energy-efficiency of LNIAs includes chopping techniques, current-reuse through amplifier stacking, and low-voltage design reaching 0.45V.

This work explores the limits of supply voltage reduction and focuses on the input stage of the LNIA, which consumes the highest amount of bias current in order to reduce its input-referred noise. Supply voltage



A Figure 1: Supply-squeezed inverter concept. The supply voltage may be reduced to just $2V_{DSAT}$ while keeping both devices saturated.

reduction has typically been limited by topological constraints such as transistor headroom requirements and signal swing, as a conventional differential pair input stage has at least three stacked devices.

We present a chopper amplifier that uses a 0.2V-supply for the input stage followed by a 0.8V-supply amplifier stage. The high input-stage current needed to reduce the input-referred noise is drawn from the 0.2V supply, significantly reducing power consumption. The 0.8V stage provides high gain and signal swing, improving linearity. Biasing and common-mode rejection (CMR) techniques for the 0.2V stage are also presented. The test chip was fabricated in a 0.18-µm CMOS process and achieves sub- μV_{RMS} input noise from 0.5-670Hz while also achieving sub- μW power consumption. The test chip was also used to measure EEG signals, demonstrating high-impedance low-noise sensor measurement capability.



▲ Figure 2: Die micrograph showing major circuit blocks in the system. The signal path occupies 1 mm².

F. M. Yaul and A. P. Chandrakasan, "A Sub-µW 36nV/JHz Chopper Amplifier for Sensors Using a Noise-Efficient Inverter-Based 0.2V-Supply Input Stage," IEEE International Solid State Circuits Conference (ISSCC), San Francisco, CA, 2016.

[•] Y. P. Chen, D. Blaauw, and D. Sylvester, "A 266nW Multi-Chopper Amplifier with 1.38 Noise Efficiency Factor for Neural Signal Recording," *IEEE Symposium on VLSI Circuits*, pp. 1-2, June 2014.

Efficiency Maximization for Wireless Charging

N. Desai, A. P. Chandrakasan Sponsorship: Hon Hai Precision Co., Ltd.

A large number of low-power sensors and wearable devices will operate in indoor environments in the near future. Wireless power transfer is well suited for powering such devices without increasing infrastructure costs. However, inductor-based wireless charging, which is a commonly used solution, suffers from a limited range. Since the receivers are low-power devices, they can also be recharged infrequently by a portable wireless power transmitter (that could be integrated into a smartphone, for example) with short bursts of energy. This is shown in Figure 1, with a wireless charging power amplifier (PA) integrated into a smartphone that powers a transmitter coil that is part of a resonant tank. The same near-field coil that implements wireless charging and/or NFC on most smartphones can potentially be multiplexed to turn the phone into a wireless power transmitter when required.

The receiver in Figure 1 is an Internet-of-Things (IoT) device that has a receiver coil built in, also part of a series tank that is resonant at the same frequency. A rectifier charges the battery on the device. The ac input characteristics of the rectifier can be modeled as an equivalent non-linear resistor whose value depends on both the input ac and output dc voltages since the rectifier is a non-linear device. As the battery on the receiver charges up, the ac input resistance of the rectifier increases.

Figure 2 illustrates the dependence of the end-toend efficiency of a coupled resonator system similar to the one in Figure 1, but with the rectifier replaced by a load resistor for different coupling factors between the transmitter and receiver coils. The efficiency at a given coupling factor has a maximum at a finite value of the load resistance, which changes as the coupling factor changes. Since both the coupling factor and the load resistance seen by the receiver coil can change during charging, the latter due to the increasing battery voltage, the maximum-efficiency point needs to be tracked dynamically by the rectifier. This tracking ensures maximum average efficiency across the entire charging duration of the receiver battery.



▲ Figure 1: System architecture of a portable wireless power transmitter charging a small IoT device.



▲ Figure 2: Theoretical plot of system efficiency of wireless power transmitter using coupled resonators while both the load resistance and the coupling factor are changed.

- 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, January 2008.
- M. Zargham and P. G. Gulak, "The Circuit Theory Behind Coupled-mode Magnetic Resonance-based Wireless Power Transmission," *IEEE Transactions on Circuits and Systems Part I: Regular Papers*, vol. 59, no. 9, pp. 2065-2074, September 2012.

Fully Integrated Thermal Energy Harvesting System with 50mV Start-Up

P. Garcha, M. Araghchini, M. Chen, N. Desai, D. El-Damak, J. Troncoso, D. Buss, J. H. Lang, A. P. Chandrakasan Sponsorship: Texas Instruments

Energy harvesting allows us to use ambient sources of energy for powering small electronic systems. Such self-powered operation can be extremely useful in wearable electronics, remote sensor nodes, and other wireless sensor networks that are widely used for monitoring and sensing applications, as it eliminates the need for battery replacement. Most of the energy harvesters employ boost converters for stepping up voltages, which can operate from as low as 10 mV input voltage. However, they typically need > 200 mV in order to start up initially. Current solutions for achieving a low-voltage start-up require the use of bulky off-chip transformers. Our research goal is to provide a proofof-concept for a fully integrated start-up system, which can cold-start from 50 mV using on-chip magnetics and also be used as a complete energy harvesting system for ultra low power applications.

Our approach involves designing on-chip transformers in Texas Instruments' flux gate technology

(Figure 1) for use in a Meissner oscillator circuit (Figure 2). The much lower Q-factors of these on-chip transformers than their discrete counterparts pose new design and optimization challenges. Hence, we have derived analytical expressions that are well-suited for use with the on-chip magnetics in order to co-optimize the oscillator components.

An optimized depletion-mode MOS device was fabricated and tested with an off-chip transformer and found to start oscillating at much lower voltages than the state-of-the-art oscillators. An on-chip transformer design with the potential of low-voltage start-up has been identified, and will be fabricated in the near future to have an integrated Meissner oscillator circuit. We have also designed and taped-out a switched capacitor DC-DC circuit to be cascaded with the Meissner oscillator block. The switched capacitor chip will rectify and boost the voltage to >1 V to have a complete start-up system for energy harvesting.



▲ Figure 1: Flux gate inductor having a permalloy core (that sits above the top metal layer) with copper windings around it; a) 3D view, b) cross-sectional view, and c) bottom view.



▲ Figure 2: Meissner oscillator circuit, to be built with Texas Instruments' flux gate technology and specially fabricated MOS device, and co-packaged for an integrated proof-of-concept solution.

- A. Shrivastava, N. E. Roberts, O. U. Khan, D. David, B. H. Calhoun, and S. Member, "A 10 mV-Input Boost Converter With Inductor Peak Current Control and Zero Detection for Thermoelectric and Solar Energy Harvesting With Kick-Start," *IEEE Journal of Solid State Circuits*, vol. 50, no. 8, pp. 1820–1832, 2015.
- J. Luo, M. Boutell, and C. Brown, "LTC3108 Ultralow Voltage Step-Up Converter and Power Manager," Data Sheet, vol. 23, no. 2, pp. 1–20, 2010.
- N. V. Desai and A. P. Chandrakasan, "A Bipolar ± 40 mV Self-Starting Boost Converter with Transformer Reuse for Thermoelectric Energy Harvesting," Proc. 2014 Int. Symp. Low Power Electron. Des., pp. 221–226, 2014.

0.3V Biopotential Sensor Interface for Stress Monitoring

S. Orguc, H. S. Khurana, H.-S. Lee, A. P. Chandrakasan Sponsorship: MIT

Miniaturized sensor nodes have a very tight power budget, especially in the case of implantables and health monitoring devices that require long operation lifetimes. Designing these sensor nodes with such a low power budget is a challenging problem, which requires careful design in both analog blocks and back-end digital signal processing blocks. In the present analog-front-end (AFE) solutions, theoretically more power can be saved at lower supply levels, but this comes at the cost of losing dynamic range, speed, and robustness. In order to further reduce the supply without significantly compromising these performance metrics, the analog architectures used in the signal acquisition should be re-designed.

The motivation of this work is to explore the limits of low-voltage design by using simplistic, yet robust circuit topologies. We present a 0.3V biopotential sensor interface (amplifier+ADC) that achieves state-of-the-art power efficiency and ensures enough circuit reliability with reduced dynamic range requirement. The system will provide diagnostic information about stress-related health problems by measuring electromyographic (EMG) signals.

Figure 1 shows the block diagram of the AFE. The AFE has large-signal cancellation ability in order to suppress the effect of unexpected motion-artifact signals coming from the environment and the sensor interface. Figure 2 illustrates the setup that we will use in future experiments.



▲ Figure 1: Block diagram of EMG signal acquisition AFE. The whole system works from 0.3V supply and has single-digit nano-watt power consumption.



▲ Figure 2: The test setup that will be used in the data collection. Once the data is digitized and processed, we will send it wirelessly to a computer for feature extraction and machine learning algorithm development.

[•] P. Harpe, H. Gao, R. van Dommele, E. Cantatore, and A. van Roermund, "A 3 nW Signal-acquisition IC Integrating an Amplifier with 2.1 NEF and a 1.5 fJ/conv-step ADC," IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers, pp. 382–383, February 2015.

[•] Y.-P. Chen, D. Jeon, Y.Lee, Y. Kim, Z. Foo, I. Lee, N.B. Langhals, G. Kruger *et al*, "An Injectable 64 nW ECG Mixed-signal SoC in 65 nm for Arrhythmia Monitoring," *IEEE J. Solid-State Circuits*, vol. 50, no. 1, pp. 375–390, January 2015.

12-bit, 300MS/s CMOS Pipelined Analog-to-Digital Converter

T. Jeong, A. P. Chandrakasan, H.-S. Lee Sponsorship: CICS, Korea Foundation for Advanced Studies

Among many analog-to-digital (ADC) architectures, pipelined ADC covers the widest performance space, and thus various applications adopt it. However, as CMOS technology scaling continues, implementing high-speed, high-accuracy pipelined ADCs has become more difficult. This is mainly due to technology scaling that causes low device intrinsic gain and reduced voltage headroom. Desired characteristics of op-amps for high-performance pipelined ADCs are having highgain, low-noise, and wide-bandwidth. Low gain and bandwidth cause charge transfer error and nonlinearity in the ADC characteristic (Figure 1). In deep submicron technologies, high-gain, high bandwidth opamps are usually implemented at the expense of high power consumption and increased complexity to overcome the low intrinsic gain and reduced headroom.

Numerous techniques have been reported in the literature to address this issue. One approach is to relax the required performance of op-amps. Digital calibration techniques have proposed to remove errors due to low-gain, low-speed op-amps. The drawback of these techniques is that continuous background digital calibration is necessary to track the gain and bandwidth drift due to the power supply or ambient temperature variations. The background calibration consumes a large amount of power. In addition, many background calibration techniques require certain input signal characteristics for the calibration to function properly and therefore are not suitable for general-purpose applications.

Techniques to avoid the usages of op-amps have also been proposed. Zero-crossing-based circuit (ZCBC) is one of the representative examples. In ZCBC-based pipelined ADCs, the ZCBC detects the instant when its input voltage crosses virtual ground rather than requiring it to be virtual ground. By doing so, ZCBC-based pipelined ADCs tend to be more power-efficient than op-ampbased pipelined ADCs. However, ZCBC-based pipelined ADCs entail considerable circuit complexity to deal with signal-dependent voltage drop across switches.

In this project, we seek to develop a digital calibration scheme for op-amp-based pipelined ADCs. The focus of this project is to develop a one-time digital calibration scheme that does not require continuous background calibration. The prototype is currently being designed in 28-nm CMOS technology.



▲ Figure 1: Charge transfer error in switched-capacitor MDAC.

- B. Murmann, "ADC Performance Survey 1997-2015," [Online]. Available: http://web.stanford.edu/~murmann/adcsurvey.html.
- B. Murmann and B. E. Boser, "A 12-bit 75MS/s Pipelined ADC using Open-Loop Residue Amplification," IEEE Journal of Solid-State Circuits, vol. 38, pp. 2040-2050, December 2003.
- L. Brooks and H.-S. Lee, "A Zero-Crossing-Based 8-bit 200MS/s Pipelined ADC," *IEEE Journal of Solid-State Circuits*, vol. 42, pp. 2677-2687, December 2009.

A CMOS Flash ADC for GaN/CMOS Hybrid Continuous-Time ΔΣ Modulator

X. Yang, H.-S. Lee Sponsorship: MIT/MTL Gallium Nitride (GaN) Energy Initiative, ONR

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-thermalnoise 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. 65 comparators are used to achieve the six most significant bits (MSBs).64 interpolators are inserted between the comparators to obtain two extra bits. The input capacitance of this design is ¼ 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 interpolaters.



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

- M. Miyahara, Y. Asada, D. Paik, and A. Matsuzawa, "A Low-Noise Self-Calibrating Dynamic Comparator for High-Speed ADCs," Proc. IEEE Asian Solid-State Circuits Conf. (A-SSCC), pp. 269-272, November 2008.
- Y.-S. Shu, "A 6b 3GS/s 11mW Fully Dynamic Flash ADC in 40nm CMOS with Reduced Number of Comparators," Symp. on VLSI Circuits Dig. Tech. Papers, 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," *IEEE Int. Solid-State Circuit Conf. (ISSCC) Dig. Tech. Papers*, pp. 388-389, February 2014.

High-Performance GaN HEMT Track-and-Hold Sampling Circuits

S. Chung, P. Srivastava, D. Piedra, X. Yang, T. Palacios, H.-S. Lee Sponsorship: MIT/MTL Gallium Nitride (GaN) Energy Initiative, ONR

The performance of emerging applications in ultrafine medical imaging, extremely high-performance cable modem, and data server backbone networks is often limited by analog-to-digital converters (ADCs) whose performance is in turn limited at least partly by a trackand-hold sampling circuits (THSC). The low supply voltage of deeply scaled CMOS transistors determines the THSC input signal range, therefore becoming a fundamental barrier to the signal-to-noise ratio (SNR) of CMOS circuits.

This research ultimately aims to design ultra high-performance THSCs in GaN-on-Si technology, which monolithically integrates GaN HEMTs with Si-CMOS transistors. Operating GaN HEMTs at a high voltage (>30 V) allows a very large input swing (>16 V) and provides performance beyond the limit of CMOS THSCs. As a first step, we designed two GaN HEMT THSCs. The first THSC was fabricated in a commercial GaN foundry technology on SiC substrate, providing 98-dB SNR at 200-MS/s (Figure 1). The second THSC design was fabricated in a GaN technology that was developed at MTL on Si substrate, which operates at 1 GS/s thanks to a higher device transition frequency f_t and external gate-bootstrapping clock (Figure 2).

While these two GaN THSCs achieved high SNR at a given input frequency, they suffered from nonlinearity. We characterized how the static nonlinearity and dynamic memory effects of GaN HEMT THSCs affect the sampled output; we observed that the GaN HEMT dynamic on-resistance does not significantly degrade the THSC linearity because the capacitive load suppresses the impact of on-resistance variation on the sampled voltage. Although dynamic nonlinearity correction techniques are mature with RF power amplifiers (PAs) and improve PA linearity typically by 20-40 dB depending on signal bandwidth and modeling accuracy, these RF PA pre-distortion techniques cannot be directly applied to THSCs. Presently, we are working on a digital post-correction technique, which will demonstrate the accurate cancellation of both static and dynamic nonlinearity in GaN HEMT THSCs.



▲ Figure 1: Pseudo-differential two-stage track-and-hold sampling circuit in 0.25-um GaN HEMT technology on SiC substrate, which demonstrates 200-MS/s 98-dB SNR and 240-MHz track-mode bandwidth with 20-V differential input signal swing.



▲ Figure 2: Track-and-hold sampling circuit with external gate-bootstrapping clock in a GaN technology developed at MTL on Si substrate, which provides over 700-MHz track-mode bandwidth and operates at 1 GS/s.

- S. Chung and H.-S. Lee, "A 200-MS/s 98-dB SNR Track-and-Hold in 0.25-um GaN HEMT," Proc. of IEEE Custom Integrated Circuits Conference, pp. 1-4, 2015.
- A. Zhu, P. J. Draxler, J. J. Yan, T. J. Brazil, D. F. Kimball, and P. M. Asbeck, "Open-loop Digital Predistorter for RF Power Amplifier Using Dynamic Deviation Reduction-based Volterra Series," *IEEE Transactions on Microwave Theory and Techniques*, vol. 56, no. 7, pp. 1524-1534, July 2008.

Broadband Inter-Chip Link Using a Terahertz Wave on a Dielectric Waveguide

J. Holloway, Z. Hu, R. Han Sponsorship: ONR, Lincoln Laboratory

The development of data links between different microchips of an on-board system has 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 solving 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 source) and normally require process modification to the mainstream CMOS technologies. Here, we aim to utilize a modulated THz wave to transmit broadband data.

Similar to the optical link, the wave is confined in dielectric waveguides, with sufficiently low loss (~0.1dB/cm) and bandwidth (>100GHz) for board-level signal transmission (Figure 1). In commercial CMOS/ BiCMOS platforms, we have previously demonstrated high-power THz generation with modulation, frequency conversion, and phase-locking capabilities. In addition, a room-temperature Schottky-barrier diode detector (in 130-nm CMOS) with <10pW/Hz^{1/2} sensitivity (antenna loss excluded) is also reported. The proposed data link will leverage these techniques to achieve a >100Gbps/channel transmission rate with <1pJ/bit energy efficiency. As the first step of this project, we have designed a new broadband chipto-fiber THz wave coupler. In contrast to previous couplers using off-chip antennas, our THz coupler is entirely implemented using the metal backend of a CMOS process and requires no post-processing (e.g., wafer thinning). The structure is also fully shielded, which prevents THz power leakage into the silicon substrate. Conventional on-chip radiators using ground shield work are the resonance type (e.g., patch antenna) and have only <5% bandwidth. In comparison, our design is based on a traveling-wave, tapered structure, which supports broadband transmission. A proof-of-concept is shown in Figure 1: two on-chip couplers are connected with a 2-cm waveguide using Rogers 3006 dielectric material. The entire back-toback setup exhibits only ~11dB insertion loss across over 60-GHz bandwidth (Figure 2).



▲ Figure 1: (Top) High-speed, energy-efficient inter-chip transmission using guided terahertz wave. (Bottom) A test structure including a pair of back-to-back THz integrated couplers separated by a 2-cm dielectric waveguide using Rogers 3006.



▲ Figure 2: The measured back-to-back insertion loss using a two-port network analyzer in the WR-3 band.

- C. Yeh, F. Shimabukuro, and P. H. Siegel, "Low-Loss Terahertz Ribbon Waveguides," Applied Optics, vol. 44, no. 28, pp. 5937-5946, October 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," *IEEE Int. Solid-State Circuit Conf. (ISSCC)*, San Francisco, CA, 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, October 2013.

A Fast, Wideband THz CMOS Spectrometer Based on Dual-Frequency Comb Architecture

C. Wang, R. Han Sponsorship: TSMC, Center for Integrated Circuits and Systems

Terahertz (THz) spectroscopy detects gaseous molecules using the unique characteristic absorption spectrum lines associated with their rotational modes. It is valuable in applications such as industrial noxious leakage monitoring and human breath analyses. A broadband, high-power THz source and sensitive detector are critical for THz spectroscopy. They enable faster spectrum sweeping and better identifications of molecules from samples with complex composition. However, the existing electronic THz spectrometers are still based on single-channel, narrowband transceiver architectures with slow sweeping speed. On the other hand, optical THz spectroscopy provides inadequate frequency resolution for gas spectroscopy and is not compact.

We propose a fast, wideband THz CMOS spectrometer using a THz-comb structure. When used in a chippair configuration (shown in Figure 1), the system simultaneously generates and detects 20 tunable signal tones, which seamlessly cover an ultra-broad band (225–315 GHz). The frequencies of the tones are equally spaced and are precisely controlled by a single input clock reference at lower frequency. Compared to that of the previous single-tone systems, the spectral sweeping time of the proposed scheme is greatly reduced. For each tone, a multi-functional THz circuit is proposed, which enables THz signal generation, radiation, and detection at the same time. Meanwhile, a feedback loop is introduced in this circuit, which greatly improves the THz power generation efficiency without deteriorating the stability. In simulation, 0.5mW output power for each transmitted tone (~5mW for the total comb spectrum) and 15-dB conversion loss for each receiving channel are achieved. The strong output power and low noise heterodyne detection further increase the signal-to-noise ratio of the system (and hence the sweeping speed). The architecture of the spectrometer is highly scalable, and the frequency coverage can be extended by cascading more comb stages. This design uses a 65-nm bulk CMOS process.



▲ Figure 1: System architecture of the proposed wideband THz CMOS spectrometer.

- C. F. Neese, I. R. Medvedev, G. M. Plummer, A. J. Frank, C. D. Ball, and F. C. De Lucia, "Compact Submillimeter/Terahertz Gas Sensor With Efficient
 Gas Collection, Preconcentration, and ppt Sensitivity," Sensors Journal, IEEE, vol. 12, pp. 2565-2574, 2012.
- H. Yi-Da, Y. Iyonaga, Y. Sakaguchi, S. Yokoyama, H. Inaba, K. Minoshima, F. Hindle, Y. Takahashi, et al, "Terahertz Comb Spectroscopy Traceable to Microwave Frequency Standard," IEEE Transactions on Terahertz Science and Technology, vol. 3, pp. 322-330, 2013.
- K. Schmalz, J. Borngräber, W. Debski, P. Neumaier, R. Wang and H. W. Hübers, "Tunable 500 GHz Transmitter Array in SiGe Technology for Gas Spectroscopy," *Electronics Letters*, vol. 51, pp. 257-259, 2015.

High-Power 1-THz Source Based on a Scalable 2D Radiating Mesh

Z. Hu, R. Han Sponsorship: Analog Devices, Inc., IHP Germany

Terahertz waves, which possess unique penetration behaviors through non-polar materials, short wavelength (versus mm-waves), and interaction with the intrinsic motions of molecules, have profound potential in imaging, communication, spectroscopy, etc. Previously, room-temperature, silicon-based sources were able to provide only milliwatt-level radiation power in the low-THz range (0.2~0.5 THz). On the other hand, there is an increasing interest in solid-state microsystems operating in the mid-THz range (~1 THz) due to higher spatial resolution in imaging, more collimated beams (hence smaller path loss) in communications, and enhanced gas spectroscopy sensitivity via high-order rotational modes. Following this trend, this work aims to push the limit of electronics further – building a 1-THz coherent radiation source targeting an output power of ~1mW.

Our design, shown in Figure 1, is based on a scalable slot mesh structure that creates a 250-GHz oscillation for the transistor pair located inside each mesh unit, and then extracts and radiates the 4th harmonic component of the oscillation. Due to the weak highfrequency activity of the transistors, the loss overhead of the above operations needs to be as small as possible. To achieve this, a multi-functional electromagnetic structure is proposed that is based on the synthesis of complex wave patterns inside multiple slot waveguides. By applying certain boundary conditions and topology manipulation (e.g., bending the slots), radiations at fundamental and 3rd harmonic signals are cancelled among adjacent slots in the horizontal direction. Meanwhile, the radiation of the 2ndharmonic signal is also canceled among adjacent slots in the vertical direction. Such configuration minimizes the loss, increases the radiation spectral purity, and creates the optimum conditions of the devices for maximum oscillation and frequency up-conversion.

Lastly, the generated standing waves at the 4th harmonic (~1 THz) are all in phase inside each horizontal slot sections, resulting in efficient, coherent radiation into free space. Due to the compactness of the design, each radiator unit occupies only $\lambda/2 \times \lambda/2$ area, which increases the radiation density and suppresses the side lobe formation. Using a 130-nm SiGe HBT process, we have designed a 1-THz source consisting of 330 coupled radiators. Figure 2 shows the simulated radiation pattern of the array, which predicts ~1.2mW radiated power, 34dBi beam directivity, and 2.8W effective isotropic radiated power.



▲ Figure 1: The architecture of the radiator array: (left) the mutual coupling between cells and (right) a single cell with the locations of 1-THz radiation indicated.



▲ Figure 2: The simulated radiation pattern of our array. The peak directivity is 34dBi (i.e., a beam collimation factor of 2500x).

- D. Mittleman, Sensing with Terahertz Radiation, vol. 85 New York:, Springer, 2013.
- R. Han et al, "A 320Ghz Phase-Locked Transmitter with 3.3 mW Radiated Power and 22.5 dBm EIRP for Heterodyne THz Imaging Systems," International Solid-State Circuits Conference (ISSCC), 2015.
- O. Momeni and E. Afshari, "High Power Terahertz and Millimeter-wave Oscillator Design: A Systematic Approach," *IEEE Journal of Solid-State Circuits*, vol. 46, no. 3, pp. 583-597, 2011.

Design, Modeling, and Fabrication of Chemical Vapor Deposition-Grown MoS₂ Circuits with E-Mode Field-Effect Transistors for Large-Area Electronics

L. Yu, D. El-Damak, S. Ha, X. Ling, U. Radhakrishna, J. Kong, D. A. Antoniadias, A. P. Chandrakasan, T. Palacios Sponsorship: NSF CIQM

The flexibility and the low temperature process for MoS₂ electronics has a great potential for realizing ubiquitous computing systems. Here we present an E-Mode field-effect transistor (FET) based on chemical vapor deposition (CVD) MoS₂ and computer-aided design (CAD) flow to realize this vision. On the device development side, the flow starts with the growth of MoS₂ using CVD, followed by device fabrication and characterization. The CAD flow includes (1) compact models of MoS_2 devices, (2) schematic design based on analytical and simulation results, and (3) layout with parameterized cells using cadence design environment. Then, the full chip layout is exported in GDS format for mask generation, and chip fabrication is performed. This design flow allows for technology-design co-optimization to realize the full potential of such emerging technology. It also allows for capturing the impact of the device parameters on the circuit performance, speeds up the layout process, reduces the number of iterations for system development, and allows for exploring the potential improvements on the system level with the predicted next generation device.

On the device side, the E-Mode device using CVDgrown large area MoS₂ is realized by a gate-first process: passive components are built and optimized before transfer of the atomically thin layer of MoS₂ Statistical distributions of threshold voltage, mobility, and subthreshold swing of E-mode MoS₂ confirm the high uniformity and high yield of this technology. Using our design flow, we built logic circuits such as multistage combinational and sequential circuits (AND, OR, XNOR, latch, NAND, NOR) and power circuits such as switch capacitance regulators. The logic circuits show correct functionality; the regulator generates output voltage regulated by switching frequency. Our device technology, modeling, and design flow bridges gaps between the development stages of MoS₂ to use the full potential of emerging technologies.



▲ Figure 1: The design flow of large-scale MoS₂ integrated circuit with highlight performance of various stages. The red arrows indicate the process procedure, and dashed blue arrow indicating the main feedback and iteration loops.

L. Yu, D. El-Damak, S. Ha, X. Ling, Y. Lin, A. Zubair, Y.-H. Lee, J. Kong, A. Chandrakasan, and T. Palacios, "Enhancement-Mode Single-layer CVD MoS₂, FET Technology for Digital Electronics," *IEDM*, 2015.

L. Yu, D. El-Damak, S. Ha, S. Rakheja, X. Ling, J. Kong, D. Antoniadis, A. Chandrakasan, and T. Palacios, "MoS₂ FET Fabrication and Modeling for Large-scale Flexible Electronics," 2015 Symposium on VLSI Technology Digest of Technical Papers, pp. T144-T145, 2015.