

Circuits & Systems for Communications, IoT, and Machine Learning

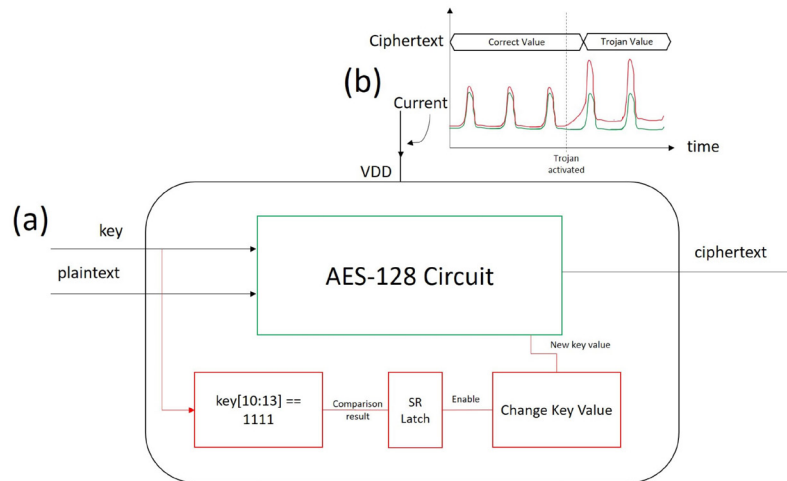
Hardware Trojan Detection using Unsupervised Deep Learning on High Spatial Resolution Magnetic Field Measurements	21
A Low-Power BLS12-381 Elliptic Curve Pairing Crypto-Processor.....	22
Direct Hybrid Encoding for Signed Expressions SAR ADC for Analog Neural Networks.....	23
Efficient Computation of Map-scale Continuous Mutual Information on Chip in Real Time.....	24
Simulation and Analysis of GaN CMOS Logic.....	25
A 0.31 THz CMOS Uniform Circular Antenna Array Enabling Generation/Detection of Waves with Orbital-Angular Momentum.....	26
Stability Improvement of CMOS Molecular Clocks Using an Auxiliary Loop Based on High-Order Detection and Digital Integration	27
A Sampling Jitter Tolerant Continuous-Time Pipelined ADC in 16-nm FinFET.....	28
Bandgap-Less Temperature Sensors for High Untrimmed Accuracy	29
High Angular Resolution THz Beam Steering Antenna Arrays in 22-nm FinFET Technology	30
DC-DC Converter Implementations Based on Piezoelectric Transformers	31
Closed Loop Control for a Piezoelectric-Resonator-Based DC-DC Power Converter	32
Leveraging Multi-Phase and Fractional-Turn Planar Transformers for Power Supply Miniaturization in Data Centers.....	33
Soft-Actuated Micro Aerial Vehicles with High Agility.....	34
Adjusting for Autocorrelated Errors in Neural Networks for Time Series Regression.....	35
Terahertz Wireless Link for Quantum Computing in 22-nm FinFET	36
Energy-Efficient System Design for Video Understanding on the Edge	37
Sparseloop: An Analytical, Energy-Focused Design Space Exploration Methodology for Sparse Tensor Accelerators	38
Multi-Inverter Discrete-Backoff: A High-Efficiency, Very-Wide-Range RF Power Generation Architecture	39
Programming a Quantum Computer with Quantum Instructions.....	40
Silicate-Based Composite as Heterogeneous Integration Packaging Material for Extreme Environments.....	41

Hardware Trojan Detection using Unsupervised Deep Learning on High Spatial Resolution Magnetic Field Measurements

M. Ashok, M. J. Turner, R. L. Walsworth, E. V. Levine, A. P. Chandrakasan
Sponsorship: NSF Graduate Research Fellowship, MITRE Corporation

One major vulnerability of integrated circuits (ICs) is the difficulty of ensuring that an IC fabricated in a third-party foundry is not a maliciously modified version of the original design. Such modifications by attackers, called hardware Trojans (example shown in Figure 1a), can leak private data from an IC, change its functionality, or have other effects. Attackers can design Trojans so that their effects are not visible during simple functional tests, making detection difficult. However, side channel methods (Figure 1b) can measure differences in circuit activity resulting from the modified logic to detect Trojans prior to the presence of functional changes.

In this work, we achieve a method of detecting small footprint hardware Trojans in a field-programmable gate array by performing high spatial resolution and wide field-of-view imaging of the circuit magnetic fields using a quantum diamond microscope. These images are then separated into Trojan-free and Trojan-inserted measurements in an automated framework by using an unsupervised convolutional neural network and clustering. With this framework, we show detection ability comparable to previous literature without requiring any knowledge of the Trojan at test time.



▲ Figure 1: (a) Block diagram of a sample hardware Trojan with malicious effects in a cryptographic circuit. (b) A general method of side channel Trojan detection that measures small differences in IC current prior to the trojan payload activation.

FURTHER READING

- M. J. Turner, N. Langellier, R. Bainbridge, D. Walters, S. Meesala, T. M. Babinec, P. Kehayias, A. Yacoby, et al., "Magnetic Field Fingerprinting of Integrated-Circuit Activity with a Quantum Diamond Microscope," *Physical Rev. Appl.*, vol. 14, no. 1, pp. 014097, Jul. 2020.
- S. Bhunia, M. S. Hsiao, M. Banga, and S. Narasimhan, "Hardware Trojan Attacks: Threat Analysis and Countermeasures," *Proc. IEEE*, vol. 102, no. 8, pp. 1229–1247, Aug. 2014.
- O. Soll, T. Korak, M. Muehlberghuber, and M. Hutter, "EM-based Detection of Hardware Trojans on FPGAs," *2014 IEEE Int. Symposium on Hardware-Oriented Security and Trust (HOST)*, pp. 84–87, May 2014.

A Low-Power BLS12-381 Elliptic Curve Pairing Crypto-Processor

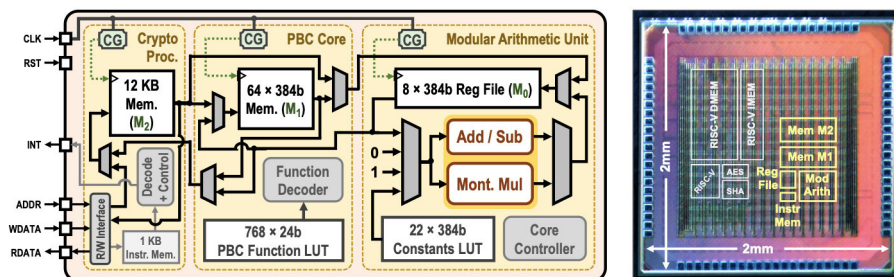
U. Banerjee, A. P. Chandrakasan

Sponsorship: Texas Instruments

Pairing-based cryptography (PBC), a variant of elliptic curve cryptography (ECC), uses bilinear maps between elliptic curves and finite fields to enable novel applications beyond traditional key exchange and signatures, e.g., signature aggregation and functional encryption. These protocols require special pairing-friendly elliptic curves; recent cryptanalysis has compromised the security of commonly used 254b BN curves. Therefore, the new BLS12-381 curve, based on a 381b prime field, is being standardized for PBC applications. However, with strong security, the new curve has higher computational complexity, making implementing low-power embedded devices challenging. To address this challenge, we present the first BLS12-381 elliptic curve pairing crypto-processor, which enables two orders-of-magnitude energy savings through efficient hardware acceleration, implements countermeasures against timing and power side-channel attacks, and provides the flexibility to implement ECC and PBC protocols for securing Internet of Things applications.

Figure 1 shows the architecture of a pairing cryptoprocessor with the chip micrograph. Our test chip was fabricated in TSMC 40-nm low-power

complementary metal-oxide-semiconductor process and supports voltage scaling from 1.1V down to 0.66V. The cryptographic core occupies a 0.2-mm² area consisting of 112k logic gates and 16 KB SRAM. Programming with custom instructions for modular arithmetic, elliptic curve point and line arithmetic, pairing operations, control, and branching is possible. Key building blocks are constant-time and secure against timing and simple power analysis side-channel attacks. For high-security use, our chip can be configured to protect against stronger differential power analysis side-channel attacks at the cost of increased energy consumption. We have evaluated pairing-based public key cryptography protocols on our chip, including signature aggregation, identity-based signatures, identity-based encryption, inner product functional encryption, and multi-party key exchange. Our hardware-accelerated implementations are 130-140× more energy-efficient than software. The programmability of our pairing crypto-processor allows new protocols, algorithm optimizations, and side-channel countermeasures to be easily implemented using one chip.



▲ Figure 1: Architecture of cryptographic core and chip micrograph.

FURTHER READING

- U. Banerjee and A. P. Chandrakasan, "A Low-Power Elliptic Curve Pairing Crypto-Processor for Secure Embedded Blockchain and Functional Encryption," *IEEE Custom Integrated Circuits Conference*, pp. 1-2, Apr. 2021.
- U. Banerjee, T. S. Ukyab, and A. P. Chandrakasan, "Sapphire: A Configurable Crypto-Processor for Post-Quantum Lattice-based Protocols," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2019, no. 4, pp. 17-61, Aug. 2019.
- U. Banerjee, A. Wright, C. Juvekar, M. Waller, A. Arvind, and A. P. Chandrakasan, "An Energy-Efficient Reconfigurable DTLS Cryptographic Engine for Securing Internet-of-Things Applications," *IEEE Journal of Solid-State Circuits*, vol. 54, no. 8, pp. 2339-2352, Aug. 2019.

Direct Hybrid Encoding for Signed Expressions SAR ADC for Analog Neural Networks

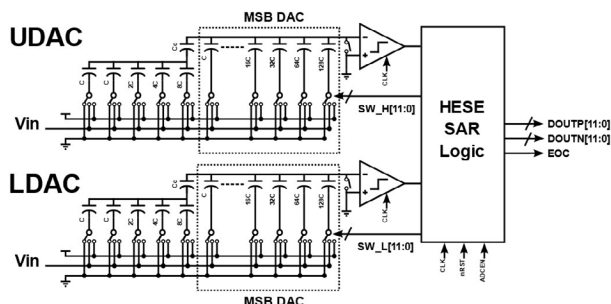
R.-C. Chen, A. P. Chandrakasan, H.-S. Lee
Sponsorship: CICS, DARPA

Artificial intelligence (AI) has proven itself to be one of the most powerful techniques for computer vision, natural language processing, and the automobile industry. Current AI algorithms that are based on deep neural networks (DNNs) are facing a crucial challenge from efficient computing. State-of-the-art DNNs need millions of weights and plenty of computation. The huge energy consumption is neither environmentally friendly nor practical in the battery-constraint edge devices. Conventional DNN hardware is based on fully digital implementation, where data movement is becoming the bottleneck. Data movement typically takes orders of magnitude more energy than the actual computation. Analog neural networks (ANNs) are a promising solution for energy-efficient AI inference. The ANNs perform the in-memory-computing to reduce the energy of data movement. Thus, the analog/digital interface

circuits are a critical part of the ANNs and are often the key bottleneck of the performance, power consumption, and area of the resulting system.

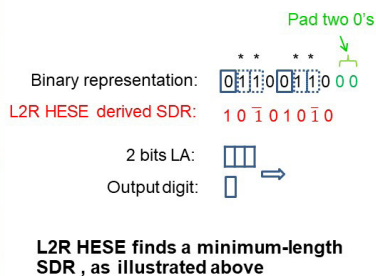
The hybrid encoding for signed expression (HESE) scheme is based on the booth encoding but with additional rules to provide the minimum-length signed-digit-representations (SDR) for efficient encoding for both DNNs including ANNs. This work focuses on a successive approximation register (SAR) analog-to-digital converter (ADC) that produces HESE encoded output on the fly. This ADC has two thresholds for 2-bits look ahead (LA). The proposed SAR ADC can directly encode the analog input to HESE instead of binary encoding. Preliminary results show that in a typical DNN, over 95% of weights can be represented by 5 terms of HESE for a 12-bits resolution.

► Figure 1: Proposed direct hybrid encoding for signed expressions SAR ADC for ANNs. The memory and computation are both inside the PE array. The ADCs are responsible for converting the analog output.



L2R (left-to-right) HESE, starting w/ NOT-IN-A-RUN state

IN-A-RUN (denoted by * in illustration)		
If 2 bit LA is 00 1->	$\bar{1}$	Enter NOT-IN-A-RUN
Else		
0->	$\bar{1}$	Output negative one
1->	0	Output zero
NOT-IN-A-RUN		
If 2 bit LA is 11 0->	1	Enter IN-A-RUN
Else		
0->	0	Output zero
1->	1	Output one



◀ Figure 2: The illustration of direct one-pass HESE. HESE is an efficient one-pass encoding to produce the minimum-length signed-digit-representations (SDR).

FURTHER READING

- Y. Peng, W. Huaqiang, G. Bin, T. Jianshi, Z. Qingtian, Z. Wenqiang, Y. J. Joshua, and Q. He, "Fully Hardware-implemented Memristor Convolutional Neural Network," *Nature*, vol. 577, no. 7792, pp. 641–646, 2020.
- H. T. Kung, B. McDanel, and S. Q. Zhang, "Term Revealing: Furthering Quantization at Run Time on Quantized DNNs," *arXiv preprint arXiv:2007.06389*, 2020.

Efficient Computation of Map-scale Continuous Mutual Information on Chip in Real Time

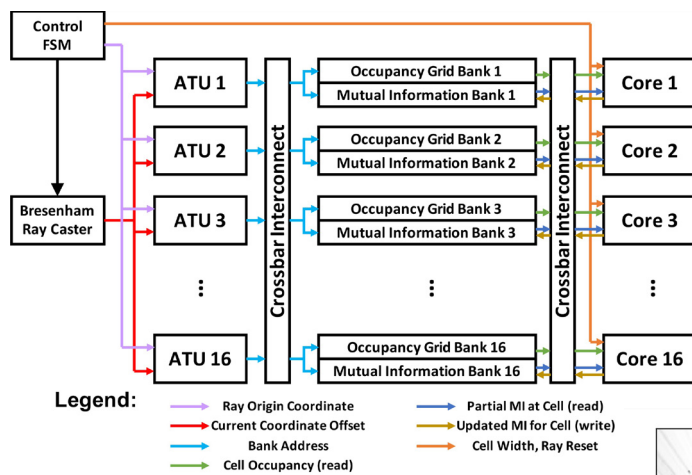
K. Gupta, P. Z. X. Li, S. Karaman, V. Sze
Sponsorship: NSF RTML, NSF CPS

Exploration tasks are essential to many emerging robotics applications, ranging from search and rescue to space exploration. The planning problem for exploration requires determining the best locations for future measurements that will enhance the fidelity of the map, for example, by reducing its total entropy. A widely studied technique involves computing the mutual information (MI) between the current map and future measurements and utilizing this MI metric to decide on the locations for future measurements.

However, computing MI for reasonably sized maps is slow and power-hungry, which has been the bottleneck in fast and efficient robotic exploration. In this paper, we introduce a new hardware accelerator architecture for MI computation that features 16 high-efficiency MI compute cores and an optimized memory subsystem that provides sufficient bandwidth to keep the cores fully utilized. Each core employs interleaving to counter the recursive algorithm and workload

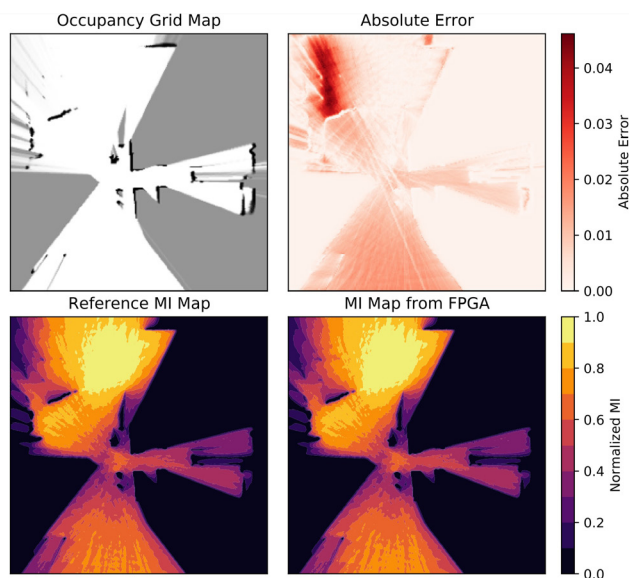
balancing and numerical approximations to reduce latency and energy consumption.

We demonstrate an optimized architecture on a field-programmable gate array (FPGA) implementation, which can compute MI for all cells in an entire 201-by-201 grid (e.g., representing a 20.1-m-by-20.1-m map at 0.1-m resolution) in 1.55 ms while consuming 1.7 mJ of energy, thus finally rendering MI computation for the whole map in real time and at a fraction of the energy cost of traditional compute platforms. For comparison, this particular FPGA implementation running on the Xilinx Zynq-7000 platform is two orders of magnitude faster and consumes three orders of magnitude less energy per MI map compute than a baseline GPU implementation running on an NVIDIA GeForce GTX 980 platform. The improvements are more pronounced when compared to CPU implementations of equivalent algorithms.



◀ Figure 1: Top-level hardware architecture comprised of address translation units (ATU), crossbar interconnects, occupancy grid map, MI map, and 16 FCMI computation cores.

▶ Figure 2: An example input grid and corresponding MI map contours generated from a reference implementation and our FPGA implementation.



Simulation and Analysis of GaN CMOS Logic

J. Jung, N. Chowdhury, Q. Xie, T. Palacios

Sponsorship: MIT Electrical Engineering and Computer Science - Texas Instruments Undergraduate Research and Innovation Scholar

There is an increasing demand for electronics that can operate in high-temperature conditions, such as spacecraft applications and sensors for industrial environments. Electronics based on wide-bandgap materials offer a promising solution, among which gallium nitride (GaN) stands out as a strong candidate due to its excellent material properties and potential for monolithic integration. Most current demonstrations of GaN logic are based on nanometal-oxide-semiconductor (nMOS) technology, which has a high static power consumption. Therefore, we are developing GaN complementary metal-oxide-semiconductor (CMOS) technology, which has lower static power consumption.

This work studies the effect of a p-channel transistor and circuit parameters on the performance of CMOS digital logic circuits. We used the MIT Virtual Source GaN-field-effect transistor (MVSGFET) model to accurately model the behavior of the n-channel and p-channel transistors, which were fabricated on

the developed GaN complementary circuit platform. We simulated and studied several building blocks for digital logic, namely, the logic inverter, multi-stage ring oscillator, and static random-access memory (SRAM) cell, using the developed computer-aided design (CAD) framework. We conducted device-circuit co-design to optimize circuit performance, using a variety of design parameters, including transistor sizing and supply voltage scaling. We projected the high-temperature performance of the circuits through simulations based on experimentally observed device behaviors. The results indicate that GaN CMOS technology based on our monolithically integrated platform has potential for a variety of use cases, including harsh-environment digital computation. We will apply this technique for more complex combinational and sequential logic building blocks, with the eventual goal of realizing a GaN CMOS microprocessor.

FURTHER READING

- N. Chowdhury, J. Jung, Q. Xie, M. Yuan, K. Cheng, and T. Palacios, "Performance Estimation of GaN CMOS Technology," *Proc. Device Research Conference (DRC)*, Jun. 2021.
- N. Chowdhury et al., "Field-induced Acceptor Ionization in Enhancement-mode GaN p-MOSFETs," in *Proc. 2020 IEEE International Electron Devices Meeting (IEDM)*, vol. 4, no. 5, pp. 5.5.1-5.5.4, Dec. 2020. DOI: 10.1109/IEDM13553.2020.9371963.
- N. Chowdhury et al., "Regrowth-Free GaN-Based Complementary Logic on a Si Substrate," *IEEE Electron Device Lett.*, vol. 41, no. 6, pp. 820-823, Jun. 2020. DOI: 10.1109/LED.2020.2987003.


A 0.31 THz CMOS Uniform Circular Antenna Array Enabling Generation/Detection of Waves with Orbital-Angular Momentum

M. I. W. Khan, J. Woo, X. Yi, M. I. Ibrahim, R. T. Yazicigil, A. P. Chandrakasan, R. Han
Sponsorship: NSF EAGER SARE Award


Multiplexing of electromagnetic (EM) waves with different frequencies, polarizations, and coding has been extensively exploited in wireless systems. Recently, another dimension of EM waves—the orbital angular momentum (OAM), is attracting increasing attention. An OAM-based wave possesses a wavefront with a helical phase distribution around the central axis of the beam. Different OAM modes, determined by the handedness and the total phase change ($\Delta\phi$) of the wavefront twist, are orthogonal. Wireless communication uses multi-OAM mode transmission to enhance spectral efficiency and physical-layer security. Conventional OAM-generation approaches incorporate dielectric spiral-phase plates, passive uniform circular antenna arrays, or metasurfaces in conjunction with separate signal drivers. These discrete solutions, however, lead to very bulky and costly systems.

In this project, we demonstrate the first chip-based (at any frequency) CMOS front-end that

generates and receives electromagnetic waves with OAM, shown in Figure 1. The chip, based on a uniform circularly placed patch antenna array at 0.31THz, transmits reconfigurable OAM modes, which are digitally switched among the (plane wave), (left-handed), (right-handed), and superposition states. The chip is also reconfigurable into a receiver mode that identifies different OAM modes with >10dB rejection of unintended modes. The array, driven by only one active path, has a measured EIRP of -4.8dBm and consumes 154mW of DC power in the OAM source mode. In the receiver mode, it has a measured conversion loss of 30dB and consumes 166mW of DC power. The OAM chip output mapped from a repeated Keccak-generated data sequence was verified, and the time-domain outputs of the Rx with different SPP configurations are shown in Figure 2, which shows good correlation with matched modes, partial correlation of multiplexed mode, and rejection of unmatched modes.



▲ Figure 1: The architecture and layout of 0.31 THz chip for OAM generation and reception.



◀ Figure 2: Time-domain output of Rx configured to receive different OAM modes when it is illuminated by same OAM sequence generated by on-chip Keccak.

FURTHER READING

- M. I. W. Khan, J. Woo, X. Yi, M. I. Ibrahim, R. T. Yazicigil, A. P. Chandrakasan, and R. Han, "A 0.31 THz CMOS Uniform Circular Antenna Array Enabling Generation/Detection of Waves with Orbital-Angular Momentum," *2021 IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*, Atlanta, GA, 2021 to be published.
- Y. Yan, G. Xie, M. Lavery, H. Huang, N. Ahmed, C. Bao, Y. Ren, Y. Cao et al. "High-capacity Millimetre-wave Communications with Orbital Angular Momentum Multiplexing." *Nat Commun.* vol. 5, p. 4876, 2014. <https://doi.org/10.1038/ncomms5876>.

Stability Improvement of CMOS Molecular Clocks Using an Auxiliary Loop Based on High-Order Detection and Digital Integration

M. Kim, H. Lee, R. Han

Sponsorship: NASA Jet Propulsion Laboratory, NSF

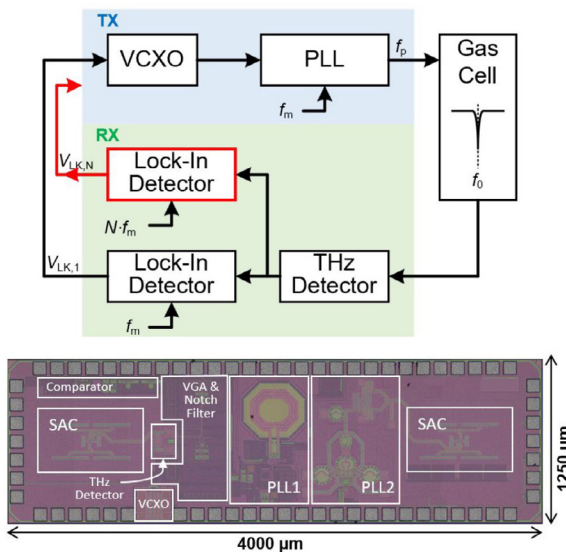
Recently, chip-scaled molecular clocks (CSMC) have achieved high-frequency stability with low power and compact size by using a rotational-mode transition of carbonyl sulfide (OCS) centered around 231.061 GHz as a frequency reference (f_0). In the molecular clock, the probing signal generated from the transmitter is frequency-modulated at f_m around the center frequency (f_c). Since f_c is locked to f_0 in a feedback loop, the output frequency inherits the excellent stability of the OCS transition frequency.

Due to its fully-electronic implementation, CSMC has reduced the cost of high-stability miniaturized clocks. However, the frequency stability is still limited by a finite loop gain of the frequency locked loop and detection non-idealities, which are susceptible to environmental disturbance even though an invariant physical constant is used as the frequency reference. In this work, we propose a new dual-loop CSMC architecture based on both fundamental and high-

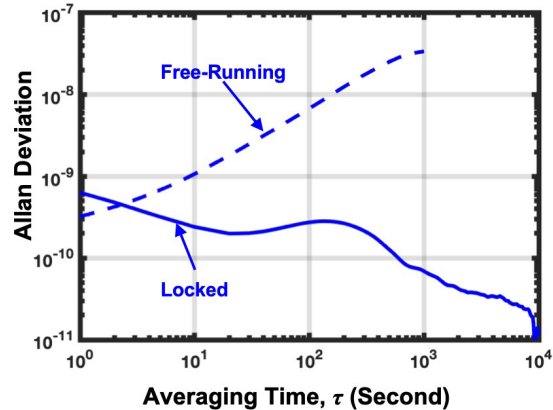
order transition probing as well as digital integration.

In order to achieve a high long-term stability without compromising the signal-to-noise ratio, the fundamental harmonic detection forms the main loop while the higher-order probing is used in an auxiliary loop. The loop fine-tunes the phase-locked loop's frequency multiplication ratio according to the sign of the high-order detection output. With a proper selection of gain and bandwidth in each loop, the main loop enables the fast correction of frequency, and the auxiliary loop responds against long-term frequency variation. Also, the frequency offset between the clock output and the OCS reference can be eliminated when the clock is locked because the auxiliary loop includes a digital integrator to obtain an infinite DC gain.

As a result, the proposed CSMC implemented in 65-nm CMOS process achieved Allan deviation of 5.4×10^{-10} and 2×10^{-11} at 1 s and 10^4 s averaging times, respectively, with 71 mW power consumption.



▲ Figure 1: Simplified block diagram and die micrograph of the proposed CSMC and dispersion curves of the fundamental, third-order, and fifth-order probing.



▲ Figure 2: Measured Allan deviation of the proposed CSMC.


FURTHER READING

- C. Wang, X. Yi, J. Mawdsley, M. Kim, Z. Wang, and R. Han, "An On-chip Fully-electronic Molecular Clock Based on Sub-Terahertz Rotational Spectroscopy," *Nature Electronics*, vol. 1, no. 7, pp. 1-7, Jul. 2018.
- C. Wang, X. Yi, M. Kim, Q. Yang, and R. Han, "A Terahertz Molecular Clock on CMOS using High-harmonic-order Interrogation of Rotational Transition for Medium/long-term Stability Enhancement," *IEEE J. of Solid-State Circuits*, vol. 56, no. 2, pp. 566-580, Feb. 2021.

A Sampling Jitter Tolerant Continuous-Time Pipelined ADC in 16-nm FinFET

R. Mittal, G. Manganaro, A. P. Chandrakasan, H.-S. Lee
Sponsorship: Analog Devices, Inc.

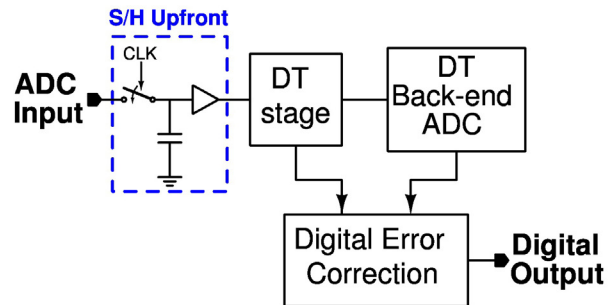
Almost all real-world signals are analog. Yet most of the data is stored and processed digitally due to advances in the integrated circuit technology. Therefore, analog-to-digital converters (ADCs) are an essential part of any electronic system. The advances in modern communication systems including 5G mobile networks and baseband processors require the ADCs to have large dynamic range and bandwidth. Although there have been steady improvements in the performance of ADCs, the improvements in conversion speed have been less significant because the speed-resolution product is limited by the sampling clock jitter (Figure 1). The effect of sampling clock jitter has been considered fundamental. However, continuous-time delta-sigma modulators may reduce the effect of sampling jitter. But since delta-sigma modulators rely on relatively high oversampling, they are unsuitable for high-frequency applications. Therefore, ADCs with low oversampling ratio are desirable for high-speed data conversion.



▲ Figure 1: Performance survey for published ADCs (ISSCC 1997-2019 and VLSI 1997-2019).

In conventional Nyquist-rate ADCs, the input is sampled upfront (Figure 2). Any jitter in the sampling clock directly affects the sampled input and degrades the signal-to-noise ratio (SNR). It is well known that for a given root mean square (rms) sampling jitter σ_t , the maximum achievable SNR is limited to $1/(2\pi f_{in}\sigma_t)$, where f_{in} is the input signal frequency. In an SoC environment, it is difficult to reduce the rms jitter below 100 fs. This limits the maximum SNR to just 44 dB for a 10 GHz input signal. Therefore, unless the effect of sampling jitter is reduced, the performance of an ADC would be greatly limited for high-frequency input signals.

In this project, we propose a continuous-time pipelined ADC having reduced sensitivity to sampling jitter. We are designing this ADC in 16-nm FinFET technology to give a proof-of-concept for improved sensitivity to the sampling clock jitter.



▲ Figure 2: A conventional discrete-time pipelined ADC with a sample-and-hold upfront.

FURTHER READING

- B. Murmann, "ADC Performance Survey 1997-2019," [Online]. Available: <http://web.stanford.edu/~murmann/adcsurvey.html>.
- R. van Veldhoven, "A Tri-mode Continuous-time/spl Sigma/spl Delta/modulator with Switched-capacitor Feedback DAC for a GSM-EDGE/CDMA2000/UMTS Receiver," *Solid-State Circuits Conference, 2003. Digest of Technical Papers. ISSCC. 2003 IEEE International*, pp. 60-477, 2003.

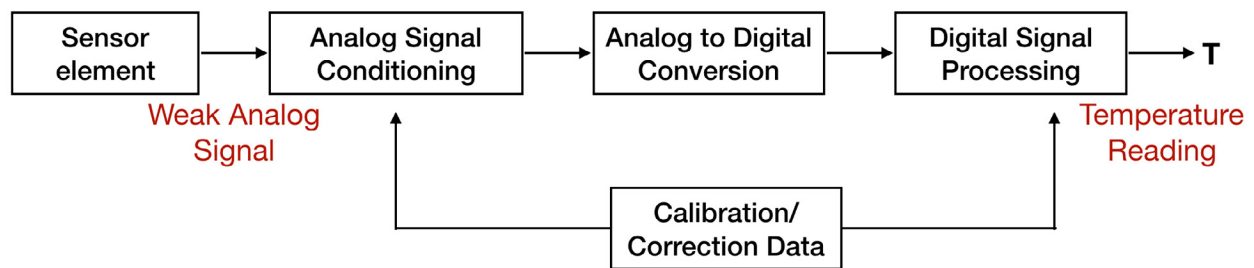
Bandgap-Less Temperature Sensors for High Untrimmed Accuracy

V. Mittal, A. P. Chandrakasan, H.-S. Lee
Sponsorship: Analog Devices Inc.

Temperature sensors are extensively used in measurement, instrumentation, and control systems. A sensor that integrates the sensing element, analog-to-digital converter, and other interface electronics on the same chip is referred to as a smart sensor. Complementary metal-oxide-semiconductor- (CMOS) based smart temperature sensors offer the benefits of low cost and direct digital outputs over conventional sensors. However, they are limited in their absolute accuracy due to the non-ideal behavior of the devices used to design them. Therefore, these sensors require either calibration or gain/offset adjustments in the analog domain to achieve desired accuracies (Figure 1). The latter pro-

cess, also called trimming, needs additional expensive test equipment and valuable production time and is a major contributor to the cost of the sensors. In order to enable high volume production of CMOS-based temperature sensors at low cost, achieving high accuracies without trimming is imperative.

This work proposes the design of a CMOS temperature sensor that uses fundamental physical quantities resilient to process variations, package stress, and manufacturing tolerances to achieve high accuracies without trimming. Simulation results prove that 3σ inaccuracy of less than 10 C can be obtained with the proposed method.



▲ Figure 1: System level diagram of a smart temperature sensor.

FURTHER READING

- G. Meijer, M. Pertijs, and K. Makinwa, "Smart Sensor Systems: Emerging Technologies and Applications," Hoboken, New Jersey: John Wiley & Sons, 2014.
- Y. Li, H. Lakdawala, A. Raychowdhury, G. Taylor, and K. Soumyanath, "A 1.05V 1.6mW 0.450C 3σ -resolution $\Delta\Sigma$ -based Temperature Sensor with Parasitic-resistance Compensation in 32nm CMOS," in *Solid-State Circuits Conference*, 2009. Digest of Technical Papers. 2009 IEEE International, pp. 340-341, Feb. 2009.

High Angular Resolution THz Beam Steering Antenna Arrays in 22-nm FinFET Technology

N. Monroe
Sponsorship: Intel Corp.

THz phased arrays are a promising emerging technology for many applications, including THz imaging, radar, communications, and other sensing applications. This is largely a result of the smaller wavelength at THz frequencies and accordingly smaller array size and weight. However, challenges exist in their design, particularly the design of THz phase shifters, which are often lossy, power hungry, and physically large, precluding their use in dense arrays. These losses often arise from the high-resolution nature of the phase shifters. In addition, lossy on-chip transmission lines significantly degrade system performance. In this work, we apply phased array principles to yield dense THz antenna arrays with only one bit of phase resolution, yielding performance benefits in terms of DC power, THz loss, size, bandwidth, and simplicity. In addition, by distributing radio frequency (RF) power spatially, we mitigate many of the losses with RF signal distribution. This approach is termed reflectarray (reflector array). We demonstrate

our approach on complementary metal-oxide semiconductor silicon in the form of a 4x4 mm² chip containing 7x7 antenna elements, operating at 260 GHz. The chip is designed in Intel 22-nm FinFET process so that multiple chips can be tiled to create large arrays that can be scaled in size based on performance requirements. The use of one-bit phase shifters comes at a cost in system-level performance by introducing sidelobes in the radiation pattern. Our work introduces a number of approaches to mitigate this, allowing the one-bit phased array design to approach the performance of a phased array with a continuous, analog phase shifter. While still in progress, this work pushes towards practical large-scale THz phased arrays.

DC-DC Converter Implementations Based on Piezoelectric Transformers

E. Ng, J. D. Boles, J. H. Lang, D. J. Perreault

Sponsorship: Texas Instruments, UROP, NSF Graduate Research Fellowship


Power converters play major roles in many applications ranging from power generation and distribution in electric grids to everyday devices such as mobile phones and computers. As many applications require small form factors, there has been a significant demand to miniaturize power converters while maintaining high performance. Typical converters rely on magnetic energy storage components, but the achievable power densities of magnetics fundamentally decrease at low volumes and therefore limit converter miniaturization.

Piezoelectrics, which have more favorable power density scaling properties than magnetics, are a promising energy storage alternative to meet the demands for low-volume power electronics. Furthermore, multi-port piezoelectric transformers (PTs) offer the additional benefits of galvanic isolation and inherent voltage conversion ratios. Despite their potential, PTs have seen little use in converters without magnetics, and such design attempts have unreported or limited efficiencies.

In this work, we systematically enumerate isolated and non-isolated converter switching sequences and

topologies that best utilize PTs as their only energy storage components. We constrain this search for (1) high-efficiency behaviors such as zero voltage switching (ZVS) and all-positive instantaneous power transfer and (2) practical characteristics such as voltage regulation capability and control simplicity. To evaluate the selected switching sequences, we also develop a model for estimating the PT's efficiency.

Initial experimental results of these converter designs demonstrate promising high-efficiency behaviors and peak whole-converter efficiencies higher than reported for most magnetic-less PT-based converters in the literature. The prototype displayed in Figure 1 is based on a commercially available PT and achieves a peak efficiency of 89.3%, which is close to our estimation model's predictions. These results suggest that PT-based converters can offer high efficiencies in addition to the low-volume scaling benefits of piezoelectrics. Such characteristics can be advantageous to high-voltage, low-power applications such as portable electronics and biomedical devices, particularly those requiring galvanic isolation.



▲ Figure 1: Photograph of the PT-based converter prototype.

FURTHER READING

- J. D. Boles, J. J. Piel, and D. J. Perreault, "Enumeration and Analysis of DC-DC Converter Implementations Based on Piezoelectric Resonators," *IEEE Transactions on Power Electronics*, vol. 36, no. 1, pp. 129-145, 2021.
- E. Ng, J. D. Boles, J. H. Lang, and D. J. Perreault, "Non-isolated DC-DC Converter Implementations Based on Piezoelectric Transformers," *Proc. IEEE Energy Conversion Congress and Exposition (ECCE)*, Vancouver, Canada, Oct. 2021.
- J. D. Boles, E. Ng, J. H. Lang, and D. J. Perreault, "High-efficiency operating Modes for Isolated Piezoelectric-transformer-based DC-DC Converters," *Proc. IEEE Workshop on Control and Modeling for Power Electronics (COMPEL)*, Aalborg, Denmark, Nov. 2020.

Closed Loop Control for a Piezoelectric-Resonator-Based DC-DC Power Converter

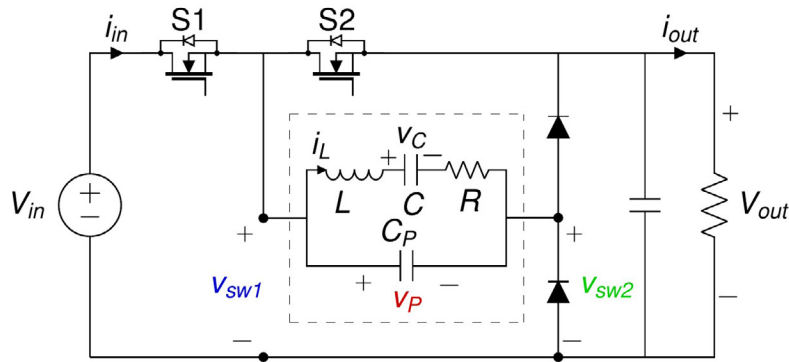
J. Piel, J. Boles, D. Perreault

Sponsorship: Texas Instruments, NSF GRFP, MIT UROP

Electronics such as computers, mobile phones, household appliances, and even electric vehicles can vary greatly in terms of supply requirements; power electronics are necessary to power these devices from standard sources. Reducing the sizes of power converters allows them to be more cost-effective and useful to a wider range of applications. Traditional DC-DC power converters make use of magnetics for energy storage, but these are less efficient and power-dense when scaled down to small sizes. Our prior work has explored the use of piezoelectric resonators (PRs) as alternative energy storage mechanisms for DC-DC converters, and we successfully demonstrated a magnetics-less PR-based converter with >99% efficiency. However, our initial prototypes depended on open-loop switching times that were manually tuned, meaning the converter could not dynamically handle transients or adjust operation when the load or temperature changed.


This work presents a closed-loop control scheme

for the PR-based DC-DC converter. For high efficiency, the converter is designed to cycle through a specific 6-stage “switching sequence” during each PR resonant cycle. In this sequence, the PR is switched between fixed-voltage energy transfer stages and resonant transition stages (shown in Figure 1), which is challenging to implement in a simple manner. The converter is controlled by two active switches, as shown in Figure 2. Both switches are triggered to turn on purely by voltage measurements of the PR node voltages. Switch 1’s on-time is modulated to control power output, and switch 2’s on-time is modulated to reach the specific high-efficiency point. Simulation results have shown that this control scheme is effective, and we are currently validating it on hardware. The successful implementation of this closed-loop control scheme will allow the PR-based converter to operate on its own, paving the way for use of these small and efficient DC-DC converters in commercial applications.



► Figure 2: The specific converter topology implemented for validation. S1 and S2 are the controlled active switches. The PR equivalent circuit is shown inside the box. This topology steps down to output voltages less than ½ the input voltage.

◄ Figure 1: Plot of PR voltage waveforms. A indicates voltage-sensed turn-ons. B indicates power-regulating turnoff. C indicates zero voltage switching controlled turnoff. The value of V_p at D is used to control turnoff at C.



FURTHER READING

- J. D. Boles, J. J. Piel, and D. J. Perreault, “Enumeration and Analysis of DC-DC Converter Implementations Based on Piezoelectric Resonators,” *IEEE Transactions on Power Electronics*, vol. 36, no. 1, pp. 129-145, 2021.
- J. D. Boles, J. E. Bonavia, P. Acosta, Y. K. Ramadass, J. H. Lang, and D. J. Perreault, “Evaluating Piezoelectric Materials and Vibration Modes for Power Conversion,” (submitted).
- J. D. Boles, J. J. Piel, and D. J. Perreault, “Analysis of high-efficiency operating modes for piezoelectric resonator-based dc-dc converters,” *Proc. IEEE Applied Power Electronics Conference and Exposition (APEC)*, Mar. 2020.

Leveraging Multi-Phase and Fractional-Turn Planar Transformers for Power Supply Miniaturization in Data Centers

M. K. Ranjram, D. J. Perreault

Sponsorship: Cooperative Agreement between the Masdar Institute of Science and Technology and MIT

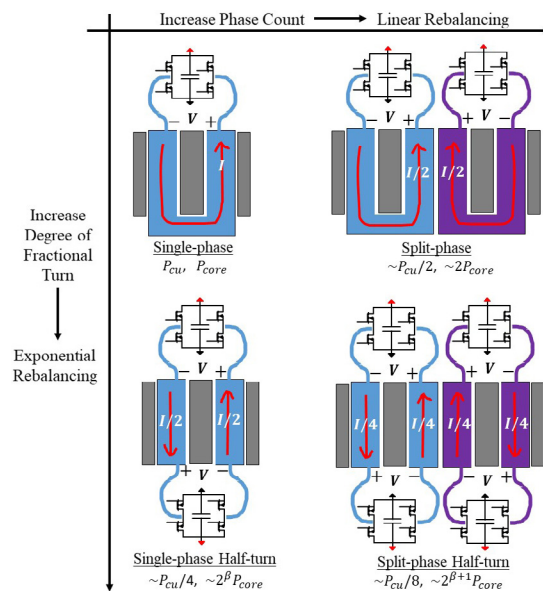
Data centers are the backbone of the Internet. Their servers represent an important and growing electrical load, and there is strong interest in miniaturizing the supplies that power them. Miniaturization is challenging as it requires both a reduction in volume and an increase in efficiency and is bottlenecked in this application by the need for a high-current transformer.

A common approach toward improving the current carrying capability of the transformer is to increase its phase count by employing multiple identical transformers in parallel. Every phase that is added proportionally decreases the “copper loss” (ohmic loss) of the transformer while proportionally increasing its core loss (i.e., loss in the magnetic material). We call this “linear rebalancing.”

In this work, we fundamentally re-think the nature of the transformer to maximally leverage the connecting electronics. In particular, by careful placement of the active switching devices required in a converter around and the passive copper and

magnetic material comprising the transformer, we can create a “fractional turn” transformer. Employing a half-turn fractional transformer reduces copper loss by a factor of four while increasing core loss by a factor of 2^β , where β is between 2 or 3 depending on the core material. Thus, fractional turn transformers yield an “exponential rebalancing” of core and copper loss.

We show that the fractional turn concept can also be combined with the common approach of adding transformer phases, enabling multi-phase fractional-turn transformers. For example, a split-phase half-turn transformer (SPHTT) combines the linear and exponential rebalancing of each of those transformers and allows a designer to get closer to the true optimum loss trade-off for a given application. We show that a SPHTT is optimal for a data center application, yielding 3.1x lower loss than a single-phase transformer in the same volume and demonstrating its clear miniaturization benefit.



▲ Figure 1: Multi-phase and fractional-turn techniques define a “map” to find the optimal loss trade-off.

Soft-Actuated Micro Aerial Vehicles with High Agility


Z. Ren, K. Chen

Sponsorship: Research Laboratory of Electronics, MIT

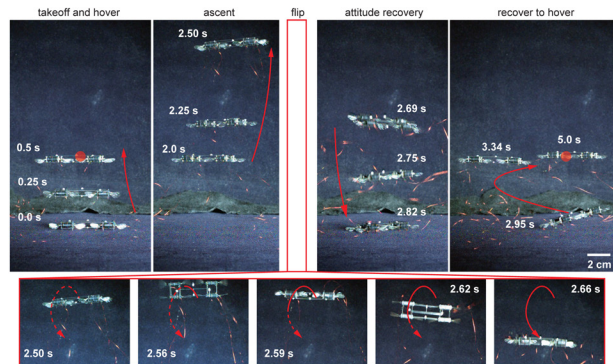
Developing agile and robust micro-aerial-vehicles (MAVs) that can demonstrate insect-like flight capabilities poses significant scientific and engineering challenges. Previously, we chose dielectric elastomer actuators (DEAs) to substitute for rigid actuators and achieved the first take-off and controlled flight of a soft-actuated MAV. In this work, we substantially improve the robot's flight capability through redesigning the actuator, robot wings, and transmission. The new MAV weighs approximately 665 mg and can complete a somersault within 0.16 s. Furthermore, its vertical ascending speed exceeds 70 cm/s, which makes it among the fastest soft mobile robots, and it outperforms rigid-powered subgram MAVs. A major contribution to this excellent performance is that we switch to a less viscoelastic elastomer, Elastosil P7670. Compared to our previously used elastomer (5:4 mixture of Ecoflex 0030 and Sylgard 184), this new elastomer has a higher resonance peak, which implies a larger displacement at

the resonant frequency. In addition, it has higher a dielectric strength and a shorter pot time.

Based on our measurement, the new MAV achieves a high lift-to-weight ratio of 2.2:1, which is 83% better than our previous work. The large lift force enables us to demonstrate hovering flight, ascending flight, in-flight collision recovery, and—more impressively—a somersault. As shown in Figure 2, the MAV takes off and hovers, accelerates upward, flips along its body pitch axis, recovers attitude, and finally returns to hover. The somersault is completed in 0.16 s; during the body flip, the motion capture system loses tracking for approximately 0.1 s. This loss results in the MAV's hitting the ground before recovering its attitude. Despite experiencing disturbance caused by the collision, the MAV quickly stabilizes its attitude and returns to hover. This is the first time that a soft-driven MAV performs agile tasks that rigid-driven MAVs have not yet demonstrated.



▲ Figure 1: Perspective view of a 665-mg MAV that is powered by DEAs. It consists of four identical modules; each module consists of an airframe, a DEA, a pair of wing hinges, wings, and transmissions.



▲ Figure 2: Sequence of composite images that show a 5-s somersault flight. The images illustrate the five flight phases: takeoff, ascent, flip, recovery, and hover.

FURTHER READING

- Y. Chen, S. Xu, Z. Ren, and P. Chirattananon, "Collision Resilient Insect-scale soft-actuated Aerial Robots with High Agility," *IEEE Transactions on Robotics*, to be published, 2021, DOI: 10.1109/TRO.2021.3053647.
- Y. Chen, H. Zhao, J. Mao, P. Chirattananon, E. F. Helbling, N. P. Hyun, D. R. Clarke, R. J. Wood, "Controlled Flight of a Microrobot Powered by Soft Artificial Muscles," *Nature*, vol. 575, no. 7792, pp. 324-329, 2019, DOI: 10.1038/s41586-019-1737-7.
- H. Zhao, A. M. Hussain, M. Duduta, D. M. Vogt, R. J. Wood, and D. R. Clarke, "Compact Dielectric Elastomer Linear Actuators," *Adv. Funct. Mater.*, vol. 28, no. 42, 2018, Art. no. 1804328, DOI: 10.1002/adfm.201804328.

Adjusting for Autocorrelated Errors in Neural Networks for Time Series Regression

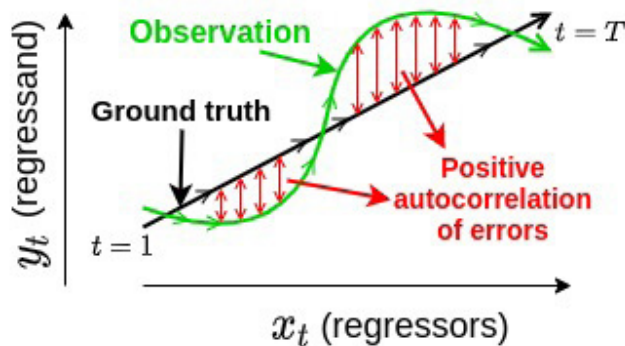
F.-K. Sun, D. S. Boning
Sponsorship: Lam Research

Time series data are ubiquitous. Researchers in many fields, including the social sciences, operations research, and engineering often collect time series data to create models for systems without prior or precise knowledge of the model structure and, in turn, provide insight for such systems. During this process of collection and creation, errors inevitably occur. Usually, the assumption is that the errors are uncorrelated at different time steps. However, in practice, errors can be autocorrelated when (1) the function space of the model and the true underlying system do not intersect, (2) some key explanatory variables are not collected, or (3) a measurement error at a current time step carries over to future time steps.

To solve this issue, previous literature, such as the Cochrane–Orcutt estimation, focuses only on cases where the model is linear or contains only predefined nonlinearity. This focus greatly limits usage, as many systems today (such as in semiconductor manufacturing) are almost certainly nonlinear while

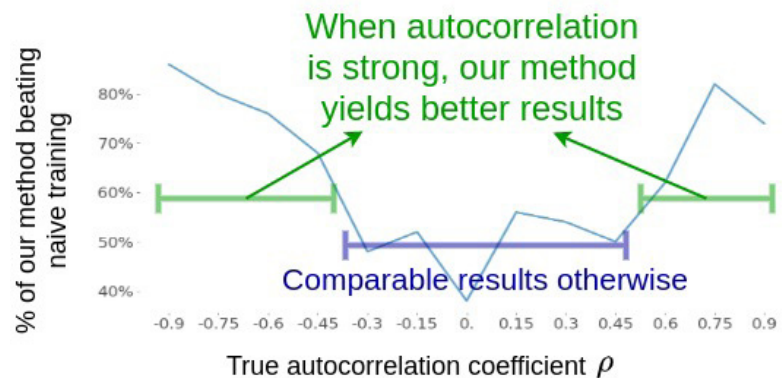
the underlying nonlinearity is unknown.

Here, we propose to use neural networks (NNs) to approximate the unknown nonlinearity and treat the autocorrelation coefficient ρ as a trainable parameter. The input to our model is a vector of features (i.e., regressors) at time t , and the output is the target scalar (i.e., regressand) also at time t . During training, we jointly optimize model parameters with the autocorrelation coefficient to adjust for the autocorrelated errors. This optimization enables us to train a NN that can fit the nonlinearity and adjust correspondingly to its autocorrelated noise. Compared to previous methods, this one has the advantages of (1) fitting unknown nonlinearity with autocorrelated noise and (2) better optimization via joint training of model parameters and autocorrelation coefficient. Our experimental results show that we obtain a better estimate of the autocorrelation coefficient and improve the model performance especially when the autocorrelated errors are substantial.



◀ Figure 1: Schematic illustration of autocorrelated errors.

▶ Figure 2: Applying our method improves the model performance.



FURTHER READING

- F.-K. Sun, C. I. Lang, and D. S. Boning, “Adjusting for Autocorrelated Errors in Neural Networks for Time Series Regression and Forecasting,” *arXiv preprint arXiv:2101.12578*, 2021.


Terahertz Wireless Link for Quantum Computing in 22-nm FinFET

J. Wang, M. I. Ibrahim, R. Han
Sponsorship: Intel University Shuttle


Quantum computing can provide exponential speed-up in solving many of today's intractable problems such as quantum chemistry, RSA encryption, DNA analysis, etc. In order to implement an error-protected quantum computer (QC), we will require approximately a million or thousands of qubits. State-of-the-art QCs have only around 100 qubits but still demand large-form-factor room-temperature electronics with many radio-frequency (RF) cables to realize the control and readout of quantum processors. These RF cables routed from room temperature to cryogenic temperature consume a non-negligible power due to the heat load, limiting the scalability and practical implementations of QCs.

We propose a terahertz (THz) wireless link to efficiently deliver the control signals to the cryogenic environment, reducing the heat loss due to the physical

conductive links (Figure 1). We implement a cryogenic THz receiver to send multi-Gb/s control signals modulated on a THz carrier (e.g., 260-GHz). The THz operation allows for a small antenna aperture size, high data rate, and minimal interference with the operation of the qubits, working around a few GHz. For the demodulation of the sub-THz downlink control signal, a THz square-law detector, operating with zero drain bias, is used first to rectify the input to baseband, and then a low-power transimpedance amplifier followed by a post-amplifier are used to boost the baseband signal so that the subsequent digital circuits can operate reliably. Figure 2 shows the chip photo of this prototype. This system opens the door for scalable and practical realization of cryogenic quantum systems.



▲ Figure 1: Block diagram of the proposed THz cryo-CMOS link system.



◀ Figure 2: Die photo of the RX part of the proposed THz cryo-CMOS link.

FURTHER READING

- J. C. Bardin, E. Jeffrey, E. Lucero, T. Huang, S. Das, D. T. Sank, O. Naaman, A. E. Megrant *et al.*, "Design and Characterization of a 28-nm Bulk-CMOS Cryogenic Quantum Controller Dissipating Less Than 2 mW at 3 K," *IEEE J. Solid-State Circuits*, vol. 54, no. 11, pp. 3043–3060, 2019.
- R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus *et al.*, "Superconducting Quantum Circuits at the Surface Code Threshold for Fault Tolerance," *Nature*, vol. 508, no. 7497, pp. 500–503, 2014.

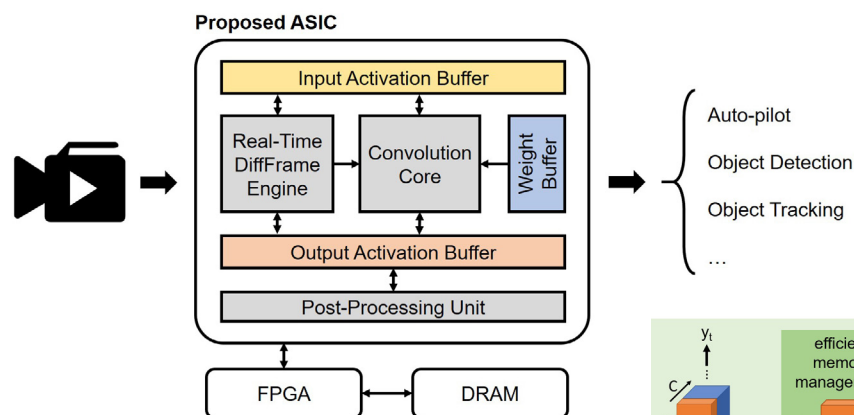
Energy-Efficient System Design for Video Understanding on the Edge

M. Wang, Z. Zhang, J. Lin, Y. Lin, S. Han, A. P. Chandrakasan
Sponsorship: Qualcomm Technologies, Inc.

With the rise of various applications including autonomous driving, object tracking for unmanned aerial vehicles, etc., the need increases for accurate and energy-efficient video understanding on the edge. Although plenty of deep learning chips designed for images exist, little work has been done for videos. Video understanding on the edge has three major challenges. First, video understanding requires temporal modeling. For example, it identifies the difference between opening and closing a box, which is distinguishable only with temporal information. Second, many applications are delay-critical, such as self-driving cars. Third, high energy efficiency matters for edge devices with a tight power budget. Due to temporal continuity, consecutive frames might share much information, providing a potential to improve processing efficiency. However, an image-based processing system, which processes frames individually, cannot utilize that.

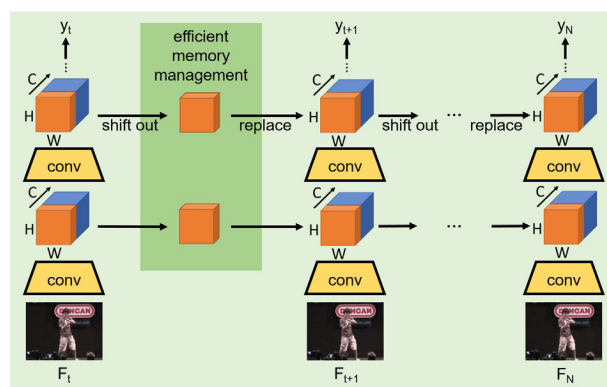
In this project, we co-designed algorithms and hardware for energy-efficient video processing on

delay-critical applications (Figure 1). We applied temporal shift module (TSM) on the backbone built on 2D convolutional neural network (Figure 2). To the best of our knowledge, our work is the first chip with temporal modeling support. Moreover, we propose a Real-Time DiffFrame method to reduce on-chip energy and DRAM traffic. It is based on the linearity of convolution, which has $\text{Conv}(f_t) = \text{Conv}(f_t - f_{t-1}) + \text{Conv}(f_{t-1})$, where f_t and f_{t-1} are the successive frames. Due to temporal continuity, $f_t - f_{t-1}$ is usually sparse. Instead of the ordinary sparsity-aware convolution in previous work, our method utilizes SparseConv, which does not dilate the input pattern and further improves energy efficiency. The load and store of $\text{Conv}(f_{t-1})$ are the overhead of the DiffFrame method. We propose a scheme to reduce memory traffic for real-time processing. The preliminary results show that our method achieves 1.6x reduction in DRAM traffic over previous work and 1.8x estimated reduction in computation and memory access over the baseline.



► Figure 2: TSM delivers state-of-the-art temporal modeling without any increase in computation, and a proposed TSM-aware mapping scheme efficiently handles the data movement.

◀ Figure 1: The proposed system supports temporal modeling and utilizes temporal redundancy in videos with a Real-Time DiffFrame Engine to improve processing efficiency for real-time applications.



FURTHER READING

- J. Lin, C. Gan, and S. Han, "TSM: Temporal Shift Module for Efficient Video Understanding," *Proc. IEEE International Conference on Computer Vision (ICCV)*, pp. 7083–7093, 2019.
- M. Riera, J.-M. Arnau, and A. González, "Computation Reuse in DNNs by Exploiting Input Similarity," *ACM/IEEE 45th Annual International Symposium on Computer Architecture (ISCA)*, pp. 57–68, 2018.
- C. Choy, J. Y. Gwak, and S. Savarese, "4D Spatio-temporal Convnets: Minkowski Convolutional Neural Networks," *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3075–3084, 2019.

Sparseloop: An Analytical, Energy-Focused Design Space Exploration Methodology for Sparse Tensor Accelerators

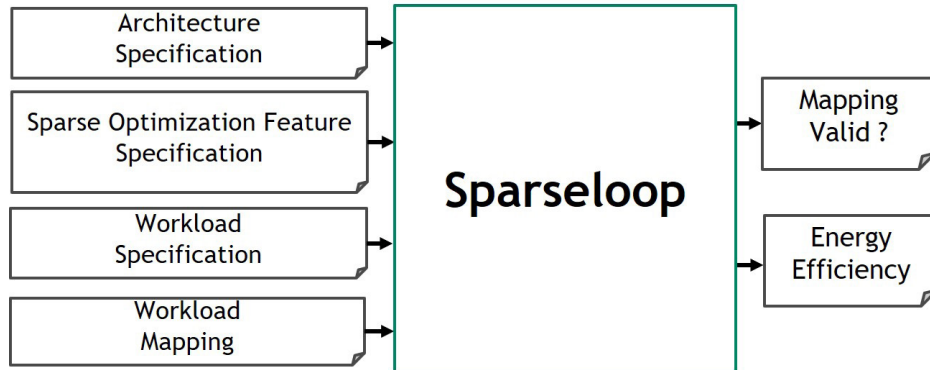
Y. N. Wu, P.-A. Tsai, A. Parashar, V. Sze, J. S. Emer
Sponsorship: DARPA

Many popular applications (e.g., deep neural networks) involve tensor computations (e.g., cross products) whose operand and result tensors can have high sparsity. Due to the nature of multiplication, zero multipliers always result in zero products. Such computations (which are called ineffectual) can be exploited by hardware sparse optimization features to improve energy efficiency and throughput. We classify these sparse optimization features into three categories: zero-gating, zero-skipping, and zero-compression. Zero-gating improves energy efficiency by keeping the associated hardware components idle for ineffectual computations. Zero-skipping further improves throughput by skipping cycles where ineffectual computations would have taken place. Zero-compression reduces required storage by storing only nonzero values.

In recent years, a variety of sparse tensor accelerators have been proposed. Based on the designer's intuitions, each design applies variations of the aforementioned sparse optimization features differently to the storage and compute levels of the

architecture. However, these specific designs are just points in a large and diverse space of sparse tensor accelerators. A fast, flexible, and accurate modeling framework would enable architects to perform early design space exploration in the complete space instead of picking specific points based on intuition.

Existing tensor accelerator models are either very detailed and design-specific, leading to slow and limited design space exploration, or fast and flexible but unable to systematically evaluate the impact of sparse optimization features, resulting in inaccurate modeling. In this work, we propose Sparseloop, an analytical modeling infrastructure for performing fast design space exploration of sparse accelerators that vary in both (1) properties associated with sparsity (e.g., compression formats, ineffectual operations' gating/skipping, and workload attributes) and (2) architecture properties (e.g., organization of the storage hierarchy). To the authors' knowledge, Sparseloop is the first analytical model that allows systematic evaluation of sparse tensor accelerators.



▲ Figure 1: Sparseloop high-level framework. Workload mapping describes the data movement and compute scheduling in space and time of the workload running on the specified architecture.

FURTHER READING

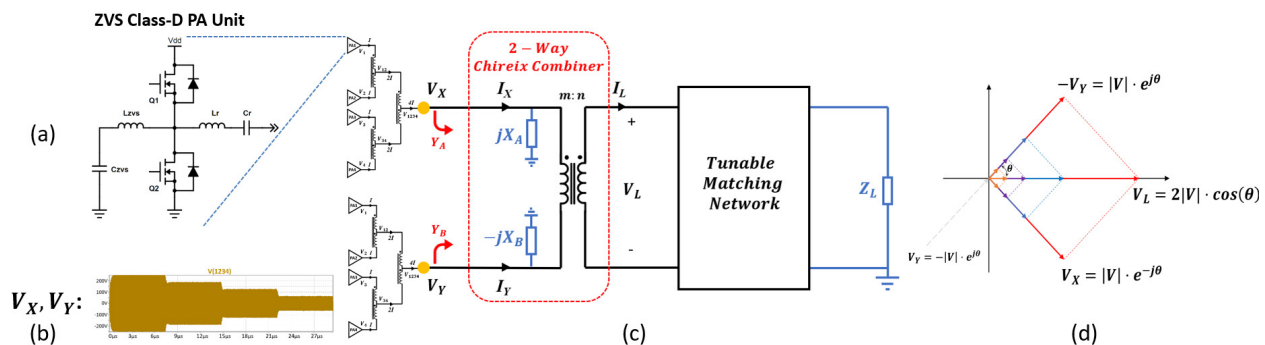
- Y. N. Wu, P.-A. Tsai, A. Parashar, V. Sze, and J. S. Emer, "Sparseloop: An Analytical, Energy-Focused Design Space Exploration Methodology for Sparse Tensor Accelerators," presented at *2020 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, [city, country], 2021.
- Y. N. Wu, J. S. Emer, and V. Sze, "Accelergy: An Architecture-Level Energy Estimation Methodology for Accelerator Designs," presented at *2019 International Conference on Computer Aided Design (ICCAD)*, 2019.
- A. Parashar et al., "Timeloop: A Systematic Approach to DNN Accelerator Evaluation," presented at *2019 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, Mar. 2019.

Multi-Inverter Discrete-Backoff: A High-Efficiency, Very-Wide-Range RF Power Generation Architecture

H. Zhang, A. Al Bastami, A. S. Jurkov, A. Radomski, D. J. Perreault
Sponsorship: MKS Instruments, Inc.

Radio-frequency (RF) power amplifiers (PAs) for industrial applications, e.g., plasma generation for semiconductor processing equipment, operate into variable load impedances at high frequency (e.g., tens of MHz) and power levels (e.g., peak power in kW), and often with wide overall power ranges and high peak-to-average-power ratios. To meet the evolving needs for semiconductor processing, goals for RF PAs in these applications include (1) operation over a wide load impedance (as determined by the plasma load); (2) operation across a very wide range of output power (e.g., 100:1 or 20 dB); (3) very fast dynamic response to output commands (e.g., at μs scale); and (4) high peak and average efficiency (to reduce cooling requirements and electricity costs). Unfortunately, meeting all these goals has not been possible to date, and efficiency is often sacrificed in order to meet the other performance metrics.

This work introduces a scalable power amplifier architecture and control approach suitable for such applications. The architecture consists of modular PAs organized in groups and employs (1) a technique which we call Multi-Inverter Discrete Backoff (MIDB), which losslessly combines the outputs of parallel-grouped switched-mode PAs and modulates the number of active PAs within the same group to provide discrete steps in RF output voltage, and (2) outphasing among the voltage outputs of PA-groups, for fast-response and continuous output power control over a wide range. To further expand the high-efficiency output power range of the system, discrete drain modulation may be optionally employed. In doing so, the MIDB-based architecture can maintain high efficiency and fast RF power control across a very wide range of power backoff.



▲ Figure 1: Example MIDB system implementation: (a) PA unit, (b) Achievable voltage magnitudes ($|V|$) with 4 PA units per PA-group, (c) PA system with differential Chireix combiner enabling outphasing between PA group voltage outputs V_X & V_Y and shunt compensation, and (d) Output power control vector diagram.

FUTHER READING

- H. Zhang, A. Al Bastami, A. S. Jurkov, A. Radomski, and D. J. Perreault, "Multi-Inverter Discrete Backoff: A High-Efficiency, Wide-Range RF Power Generation Architecture," 2020 IEEE 21st Workshop on Control and Modeling for Power Electronics (COMPEL), Aalborg, Denmark, 2020. doi: 10.1109/COMPEL49091.2020.9265702.
- A. Al Bastami, H. Zhang, A. Jurkov, A. Radomski, and D. Perreault, "Comparison of Radio-Frequency Power Architectures for Plasma Generation," 2020 IEEE 21st Workshop on Control and Modeling for Power Electronics (COMPEL), Aalborg, Denmark, 2020. doi: 10.1109/COMPEL49091.2020.9265700.

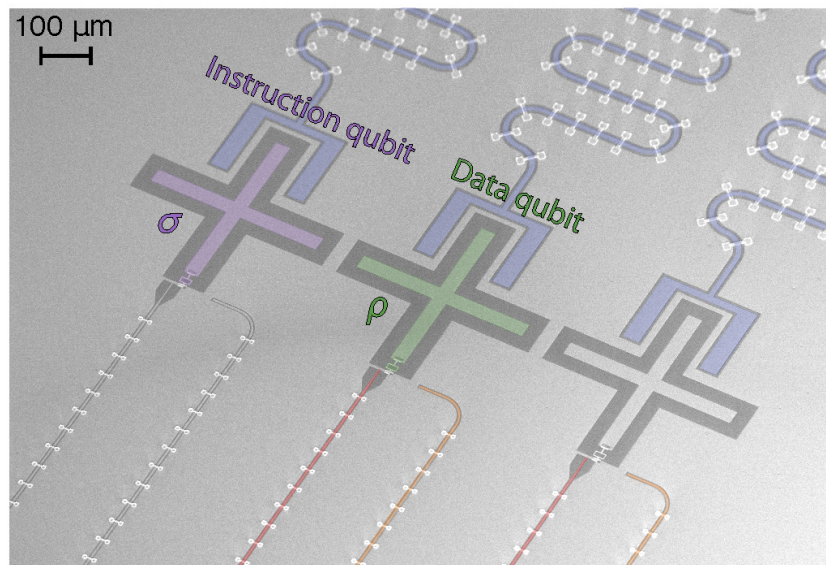
Programming a Quantum Computer with Quantum Instructions

M. Kjaergaard, M. E. Schwartz, A. Greene, G. O. Samach, A. Bengtsson, M. O’Keeffe, C. M. McNally, J. Braumüller, D. K. Kim, P. Krantz, M. Marvian, A. Melville, B. M. Niedzielski, Y. Sung, R. Winik, J. L. Yoder, D. Rosenberg, K. Obenland, S. Lloyd, T. P. Orlando, I. Marvian, S. Gustavsson, W. D. Oliver

Sponsorship: MK acknowledges support from the Carlsberg Foundation during part of this work. AG acknowledges funding from the 2019 Google US/Canada PhD Fellowship in Quantum Computing. IM acknowledges funding from NSF grant FET-1910859. This research was funded in part by the U.S. Army Research Office Grant W911NF-18-1-0411 and the Assistant Secretary of Defense for Research & Engineering under Air Force Contract No. FA8721-05-C-0002.

The use of quantum bits to construct quantum computers opens the door to dramatic computational speedups for certain problems. The maturity of modern quantum computers has moved the field from being predominantly a quantum-device-focused research area to also include practical quantum-computing-application-focused research. Our research explores a new experimental result on a foundational aspect of how to program quantum computers. A central principle of classical computer programming is the equivalence between data and instructions about what to do with that data. In quantum computers, this equiv-

alence is broken: classical hardware is used to generate the sequence of operations to be executed on the quantum data stored in the quantum computer. Our experiment shows for the first time how the instruction-data symmetry can be restored to quantum computers. We use superconducting qubits as a platform to implement high-fidelity quantum operations enabling the so-called density matrix exponentiation algorithm to generate these quantum instructions. This algorithm provides large quantum speedups for a family of other quantum algorithms.



▲ Figure 1: Scanning electron microscope image of three superconducting qubits (the ‘+’ shapes), identical to the qubits used in our experiment.

FURTHER READING

- M. Kjaergaard, et al., "Programming a Quantum Computer with Quantum Instructions," *arXiv:2001.08838*, 2020.
- A. Greene, M. Kjaergaard, et al., "Error Mitigation via Stabilizer Measurement Emulation," *arXiv:2102.05767*, 2021.
- M. Kjaergaard, et al., "Superconducting Qubits: Current State of Play," *Annual Reviews of Condensed Matter Physics* 11, 369-395, 2020.


Silicate-Based Composite as Heterogeneous Integration Packaging Material for Extreme Environments

J. C. McRae, M. A. Smith, B. P. Duncan, E. Holihan, V. Liberman, C. Rock, D. Beck, L. M. Racz
Sponsorship: Office of the Under Secretary of Defense for Research and Engineering

Electronic microsystems are foundational to today's computational, sensing, communication, and information processing capabilities, therefore impacting industries such as microelectronics, aerospace, healthcare, and many more. Cell phones are an example of what is possible when a variety of systems can be tightly integrated into a highly portable and capable system. However, as we aim to improve our ability to interact and operate (e.g., sense, communicate, record, compute, move, etc.) in extreme environments (such as outer space or the human body), new methods and materials must be developed to manufacture such integrated systems that will endure post-processing, environmental, and operational challenges.

Typical organic-based packaging materials (e.g., polymer adhesives, coatings, and molding materials) often suffer from outgassing and leaching that can

lead to system contamination, as well as coefficient of thermal expansion (CTE) mismatches that can lead to warpage and breakage with fluctuations in system temperature during operation. This work demonstrates an alternative, by using a silicate-based inorganic glass composite as an electronics packaging material for stability in extreme environments. Combining liquid alkali sodium silicate (water glass) and nanoparticle fillers, composites can be synthesized and cured at low temperatures into chemically, mechanically, and thermally (up to 400°C) stable structures using high-throughput processing methods such as spin and spray coating. Further, this material can be processed into thick layers (10s to 100s of microns), fill high aspect ratio gaps (13:1), withstand common microfabrication processes, and have its CTE tailored to match various subs



DISTRIBUTION STATEMENT A: APPROVED FOR PUBLIC RELEASE. DISTRIBUTION IS UNLIMITED

▲ Figure 1. Common packaging and heterogeneous integration technologies

FUTHER READING

- J. C. McRae, M. A. Smith, B. P. Duncan, E. Holihan, V. Liberman, C. Rock, D. Beck, L. M. Racz, "Sodium Metasilicate-Based Inorganic Composite for Heterogeneous Integration of Microsystems," in *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 11, no. 1, pp. 144-152, Jan. 2021