ISSCC 2023

TRENDS

Conditions of Publication

PREAMBLE

The Session Overviews and Highlights to follow serve to capture the context, highlights, and potential impact, of the papers to be presented in each Session at ISSCC 2023 in February.

OBTAINING COPYRIGHT to ISSCC press material is EASY!

You may quote the Subcommittee Chair as the author of the text if authorship is required.

You are welcome to use this material, copyright- and royalty-free, with the following understanding:

- o That you will maintain at least one reference to ISSCC 2023 in the body of your text, ideally retaining the date and location. For detail, see the FOOTNOTE below.

- o That you will provide a courtesy PDF of your excerpted press piece and particulars of its placement to lcfujino@aol.com and shahriar@ece.ubc.ca.

FOOTNOTE

- From ISSCC's point of view, the phraseology included in the box below captures what we at ISSCC would like your readership to know about this, the 70th appearance of ISSCC, on February 19th to February 23rd, 2023.

This and other related topics will be discussed at length at ISSCC 2023, the foremost global forum for new developments in the integrated-circuit industry. ISSCC, the International Solid-State Circuits Conference, will be held virtually on February 19 - February 23, 2023

ISSCC Press Kit Disclaimer

The material presented here is preliminary.

As of November 4, 2022, there is not enough information to guarantee its correctness.

Thus, it must be used with some caution.

HISTORICAL TRENDS IN TECHNICAL THEMES

ANALOG SYSTEMS

ANALOG SUBCOMMITTEE

POWER MANAGEMENT SUBCOMMITTEE

DATA CONVERTERS SUBCOMMITTEE

Analog – 2023 Trends

Subcommittee Chair: Maurits Ortmanns, University of Ulm, Institute of Microelectronics, Ulm, Baden-Württemberg, Germany

At ISSCC 2023, new analog circuit techniques for advancing the performance of amplifiers, oscillators, and sensor interfaces continue to emerge.

Despite their high maturity, Class-D amplifiers continue to advance in performance, thanks to new creative solutions revealed in this year's papers. A capacitively-coupled chopper Class-D amplifier with digital input combines feedback after the output filter with other techniques to achieve the highest DR, >120dB. A chopper-stabilized amplifier uses a relaxed fill-in technique to reduce the interaction between the input signal and the chopper clock that causes IMD.The reduced input switching activity results in 25× less input current and a 15% power saving, achieved by power-cycling the fill-in OTA.

In the sensor area, significant advances are made as well. A shunt-based current sensor uses a calibration scheme based on an accurate on-chip current source to achieve ±0.2% gain error over -40°C to 125°C. Magnetic current sensors are widely used in applications where galvanic isolation and wide bandwidth (BW) are desired. By using Hall plates for low frequencies and pick-up coils for high frequencies, a hybrid magnetic sensor achieves high resolution over a wide frequency range, ±1.1% gain flatness, and significantly improved energy efficiency over previous designs.

Other papers describe significant advances in Integrated stress sensors, and temperature sensors. BJT-based high-accuracy temperature sensors with only 1 point trim are demonstrated by using capacitive biasing and continuous-time current-mode readout techniques, achieving 0.34pJ·K^2 resolution FoM with sub-1V supply and 0.85pJ·K^2 resolution FoM with 3σ inaccuracy of ±0.1°C, respectively (Figure 1), and both sensors cover the military temperature range.
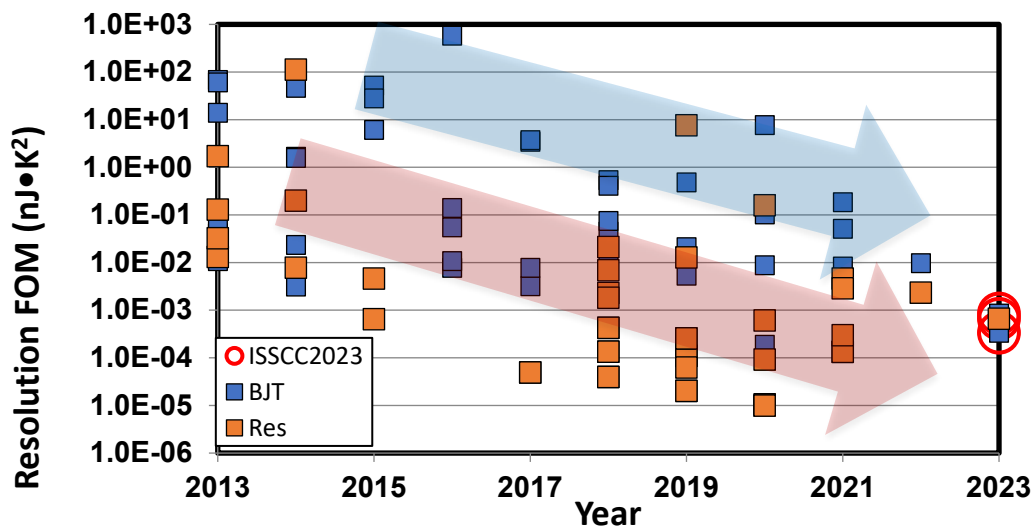

Figure 1: Trends in power efficiency of temperature sensors.

Significant advances in the aging resilience of RC frequency references and the startup time and energy of XOs are also disclosed. By periodically locking to a less-aged reference oscillator, an RC oscillator achieves ±1030ppm frequency inaccuracy from -40°C to 85°C after accelerated aging for 500 hours at 125°C. XOs reduce startup time and improve energy efficiency with new schemes that expand the tolerable injection frequency error up to 10,000ppm. A BAW-based oscillator uses a new divider to achieve a frequency stability of ±4ppm and jitter of <90fs.

Power Management – 2023 Trends

Subcommittee Chair: *Bernhard Wicht, Leibniz University, Hannover, Germany*

Power management supports a diverse application space with a range of requirements coming from automotive, computing, communications, biomedical, consumer and other applications. Some systems may operate at less than 1V, but draw energy from a high-voltage supply, others may draw power from low-voltage batteries, however, require hundreds of volts to kV to operate; some systems must respond to near-instantaneous changes in voltage and current levels; some require extreme reliability and careful design to minimize electromagnetic interference (EMI); some must transfer power and/or signals across isolation boundaries, and yet others must efficiently process and recover energy from intermittent low-power transducers. Thus, while there is no single figure of merit for the wide scope of power management systems and technologies, most of them strive to achieve high performance, minimal power loss, small size, and low cost.

Trends seen at this year's ISSCC align with the growing diversity of applications and goals outlined above. Power-conversion topologies continue to evolve, many leveraging hybrid architectures that mix switched-capacitor (SC) architectures with inductor(s). A number of these topologies focus on multiple-output designs, some using single-inductor multiple-output (SIMO) techniques, others using shared- or multi-stage designs. There are a variety of efforts to address key challenges in hybrid designs such as balancing flying capacitors and achieving safe startup, as well as other efforts to improve voltage regulation and transient response.

This year we have seen growth in universal serial bus (USB) and computing power delivery with a number of designs addressing challenging aspects (cost, density, transient response) of these applications. Among USB-focused designs, some leverage the USB cable inductance as passive energy storage element to improve these metrics and reduce heat generated in the system battery chargers while providing bi-directional operation. Some designs focus on compute power delivery and work to place a final power-conversion stage close to the load using a vertical power-delivery concept. In power delivery for computing applications, the use of high switching frequencies and meeting strict transient requirements remain a core objective and a key challenge, respectively.

Gallium nitride (GaN) designs are expanding in scope and demonstrating higher levels of integration for high-voltage and automotive applications. This year, we have seen an increase in the complexity and feature set of monolithic GaN designs with new analog circuit concepts providing functionalities that previously were reserved for silicon-only designs. There are new concepts in gate driving for GaN and SiC which achieve high common-mode transient immunity (CMTI) and provide adaptive switching for improved performance, as well as spread spectrum techniques to reduce EMI. Additional concepts include gate drive signaling across (up to 20kV) galvanic-isolation barriers which are also subject to high dV/dt transients. GaN-based power conversion designs include high conversion ratio, bidirectional buck-boost and flyback topologies.

Other areas of focus at ISSCC 2023 include wireless power transfer (WPT) and energy harvesting. In particular, energy harvesting techniques such as bias-flip piezoelectric, multi-source harvesting, and maximum power point tracking (MPPT) are explored in several papers. Wireless power transfer papers include techniques for foreign object detection and methods to improve the spatial distribution of electromagnetic energy using multi-path transmitters. Other concepts such as new receivers and rectifiers with adaptive zero current switching (ZCS) and hybrid techniques, some pushing to very high operation frequencies, are also explored.

Data Converters – 2023 Trends

Subcommittee Chair: *Jan Westra, Broadcom, Bunnik, The Netherlands*

Data converters are a critical link between the analog physical world and the world of digital computing and signal processing, prevalent in modern electronics. The need to faithfully preserve the signal across domains continues to pressure data converters to deliver more bandwidth and linearity while continuing to increase power efficiency. This year, ISSCC not only continues the trend of reporting highly energy-efficient analog-to-digital converters (ADCs), but also showcases new and exciting converter architectures, which opens new possibilities for data conversion.

Time-based quantization and hybrid pipelined-SAR architectures are expanding the speed limits of the current state-of-the-art in Nyquist converter design, while incremental converters are reaching new levels of efficiency. In noise-shaping converter design, delta-sigma and noise-shaping-SAR converters are continuing their prominent role and show their strengths in both high-efficiency, as well as high-speed data conversion. Various types of dynamic amplifiers such as ring-amp and floating inverter amplifiers are one of the key aspects in pushing the limit of power efficiency in these architectures.

The three figures below represent traditional metrics that capture the innovative progress in ADC design. The first figure plots power dissipated relative to the Nyquist sampling rate (P/f_{snyq}), as a function of signal-to-noise and distortion ratio (SNDR), to give a measure of ADC power efficiency. Note that a lower P/f_{snyq} metric represents a more efficient circuit on this chart. For low-to-medium-resolution converters, energy is primarily expended to quantize the signal; thus the overall efficiency of this operation is typically measured by the energy consumed per conversion and quantization step. The dashed trend-line represents a benchmark of 1fJ/conversion-step. Circuit noise becomes more significant with higher-resolution converters, necessitating a different benchmark proportional to the square of the signal-to-noise ratio, represented by the solid line. Designs published from 1997 to 2022 are shown in circles. ISSCC 2023 designs are shown in red stars.

The second figure plots signal fidelity vs. the Nyquist sampling rate normalized to power consumption. At low sampling rates, converters tend to be limited by thermal and flicker noises, independent of the sample rate. Higher speeds of operation present additional challenges in maintaining accuracy in an energy-efficient manner, indicated by the roll-off vs. frequency in the dashed line. The last ten years have resulted in an improvement of over 10dB in power-normalized signal fidelity, or a 10× improvement in speed for the same normalized signal fidelity. At ISSCC this year, hybrid pipelined-SAR and noise-shaping SARs are continuing the trends in the speed vs efficiency corner of the graph. Time-based conversion of signal and time-domain quantization are becoming an integral part of converters in scaled nodes. Time-interleaving multiple channels with error correction over process, voltage and temperature (PVT) as well as clock skew between channels is a necessity in these architectures to achieve robust performance in high-speed conversion rates.

The final figure plots ADC bandwidth as a function of SNDR. Sampling jitter or aperture errors coupled with an increased noise bandwidth make achieving both high resolution and high bandwidth a particularly difficult task. While ten years ago, a state-of-the-art data converter showed an aperture error of approximately 1ps$_{rms}$, in recent years, designs with aperture errors below 100fs$_{rms}$ have been published, many of which have been published at ISSCC.

Finally, this year's ISSCC presents multiple time-assisted data converters with over gigahertz sampling speed with extremely power and area efficient implementations.
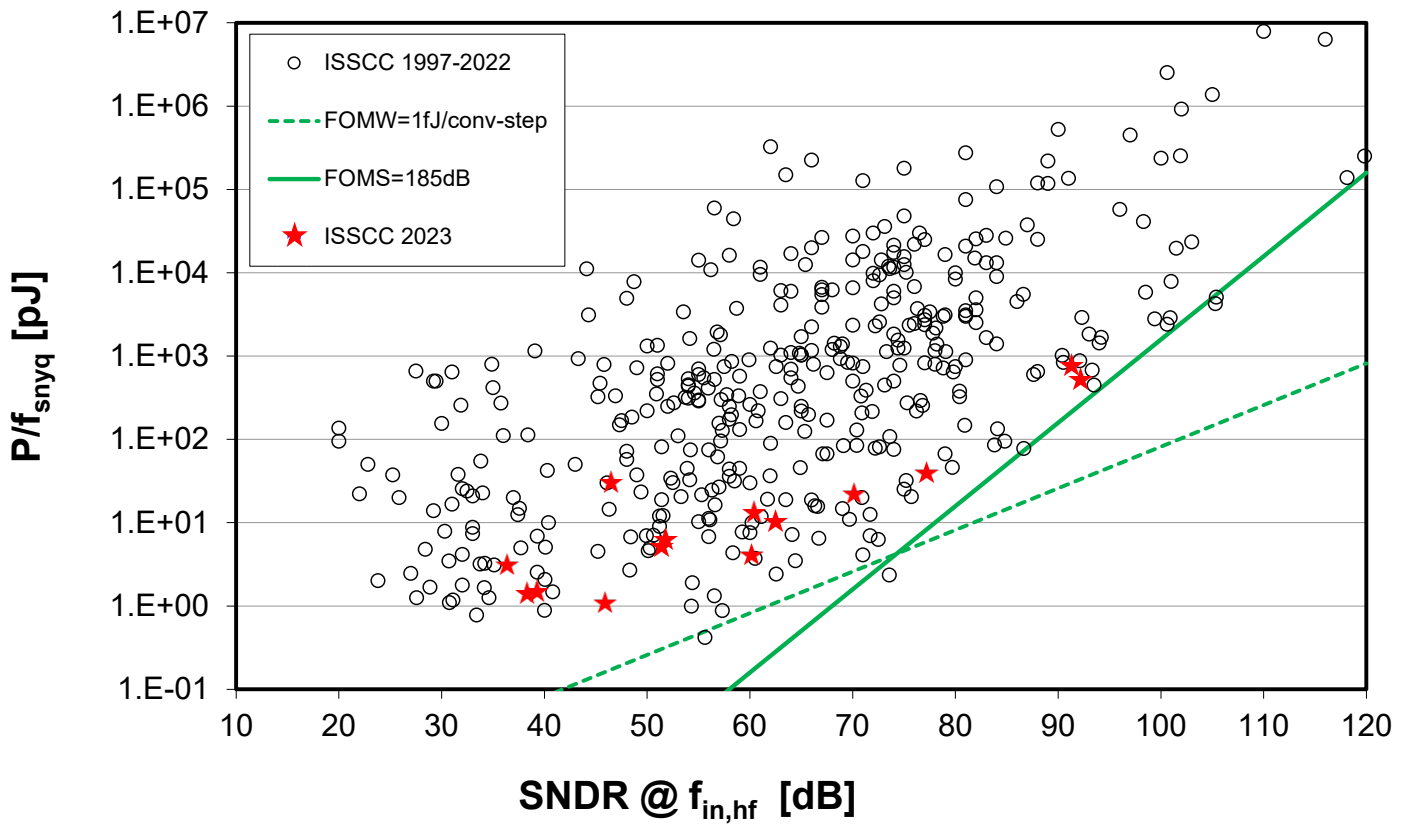
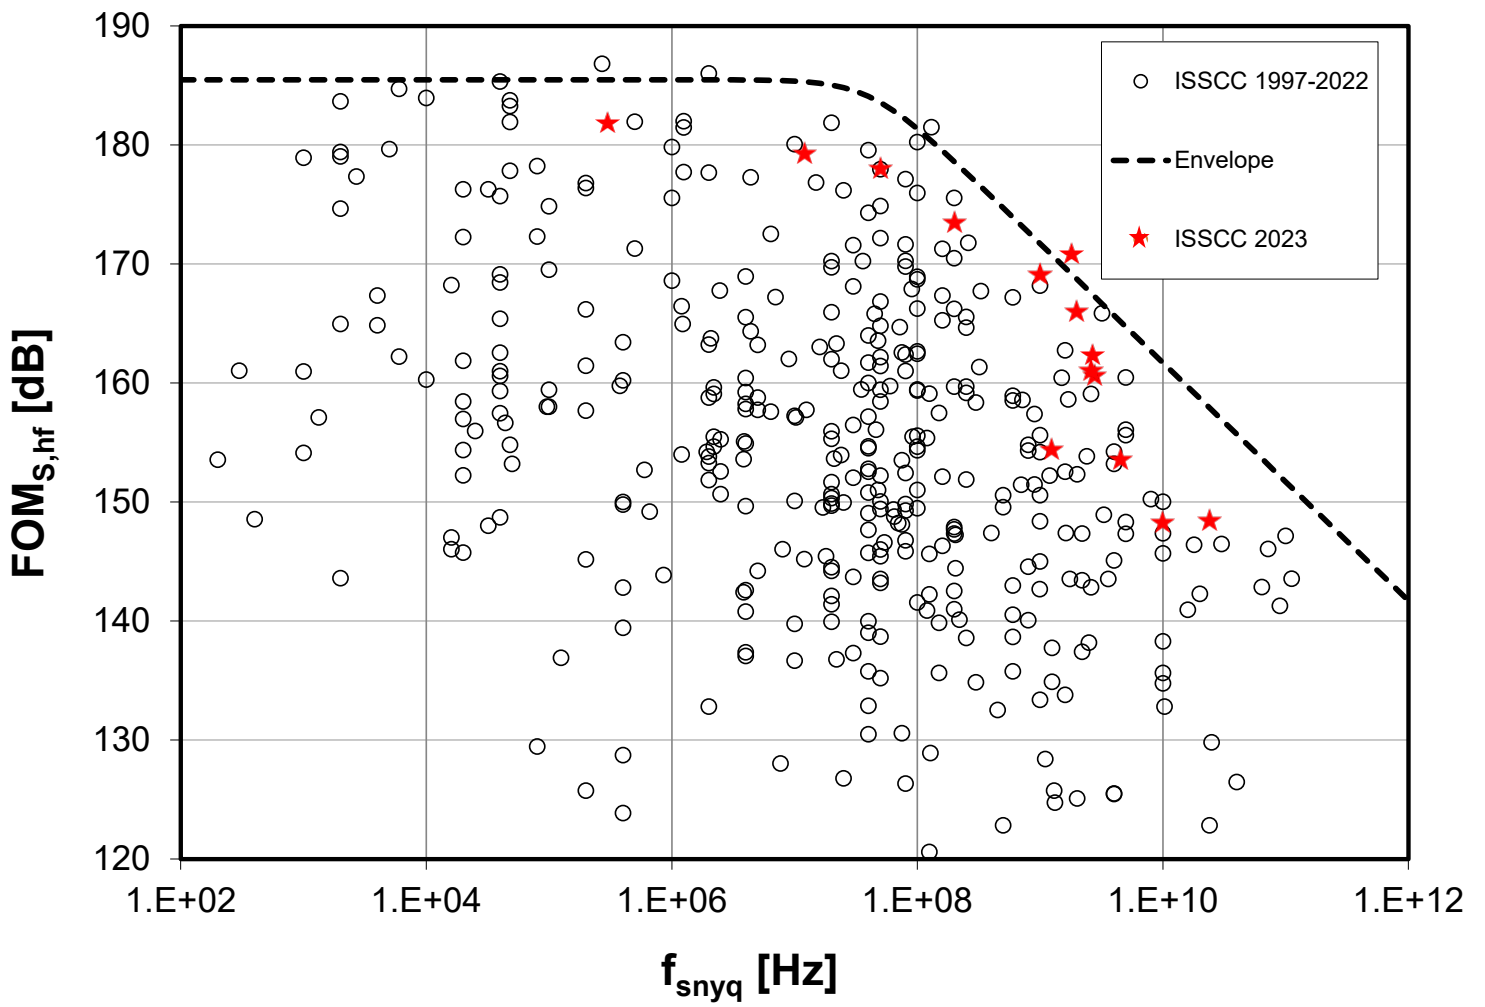Figure 1: ADC power efficiency (P/f$_{snyq}$) as a function of SNDR.

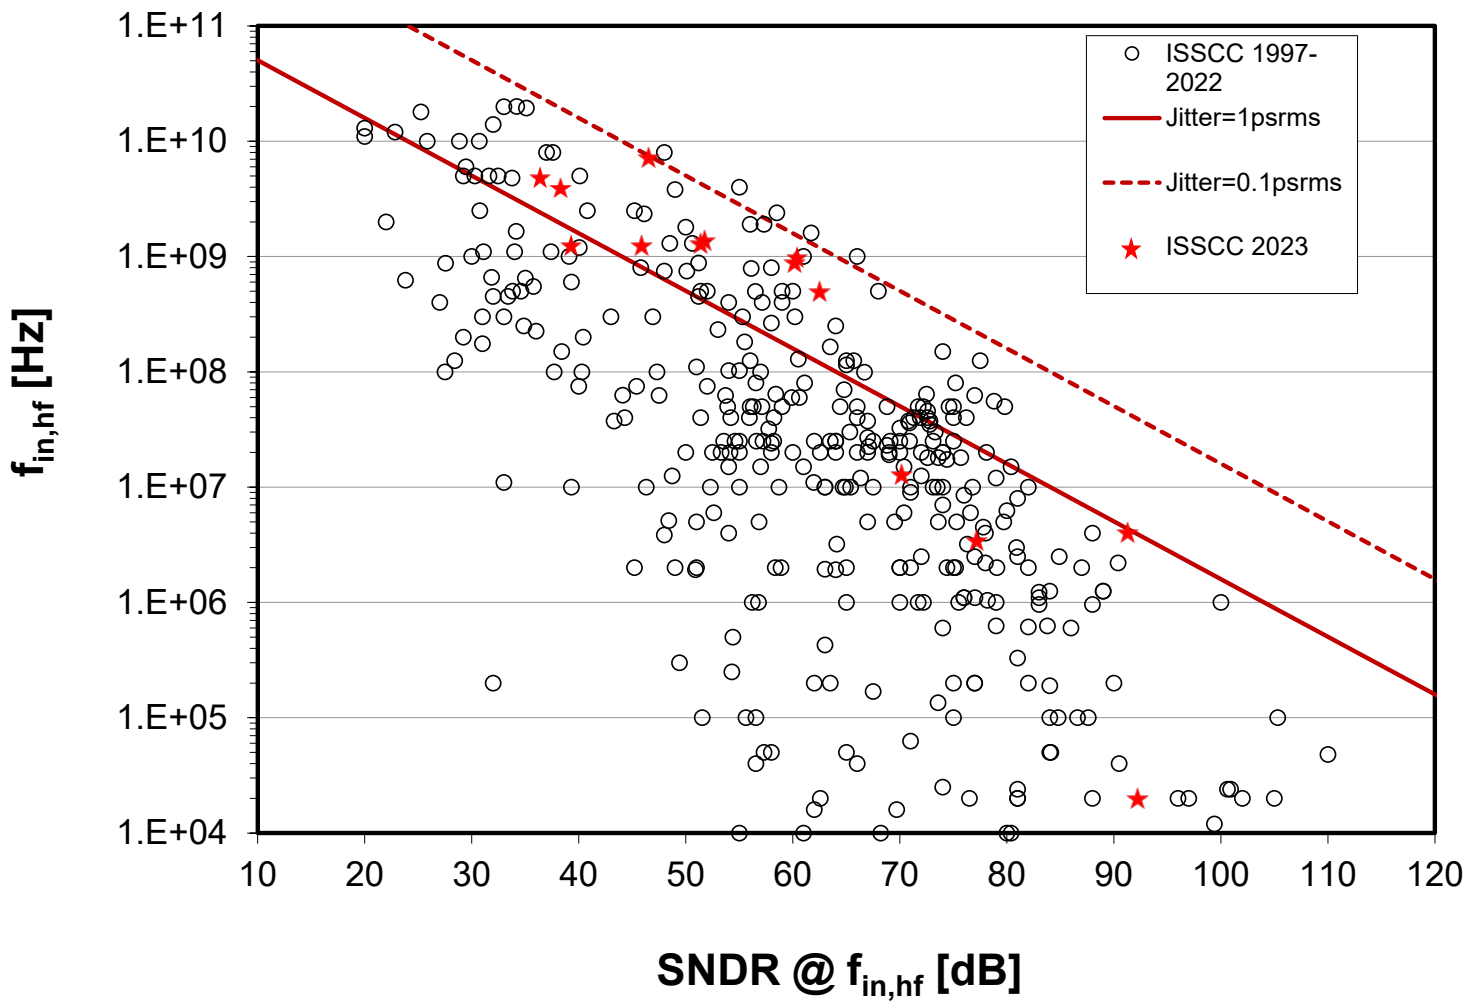Figure 2: Power normalized noise and distortion vs. the Nyquist sampling rate.

Figure 3: Bandwidth vs. SNDR.

HISTORICAL TRENDS IN TECHNICAL THEMES

COMMUNICATION SYSTEMS

RF SUBCOMMITTEE – WIRELESS SUBCOMMITTEE

WIRELINE SUBCOMMITTEE

RF Subcommittee – 2023 Trends

Subcommittee Chair: Jan Craninckx, imec, Belgium

ISSCC 2023 features record-setting advancements in phase-locked loops (PLLs), voltage-controlled oscillators (VCOs), RF transceiver building blocks, and THz signal generation driven by emerging requirements in 5G and 6G communications, the Internet of Things (IoT), radars, and imaging at RF, mm-wave, and THz frequencies. This document highlights such trends that will be presented at ISSCC 2023.

Frequency Generation: ISSCC 2023 highlights new results in voltage-controlled oscillators (VCOs) and features several approaches that demonstrate better-than 190dBc/Hz FoMs. First, a Class-F VCO with inherent common-mode-noise self-cancelation and isolation achieves 192.8dBc/Hz FoM at 1MHz offset. Second, a ~26GHz dual-path synchronized quad-core oscillator with a 193.3dBc/Hz FoM at 10MHz offset is reported. Third, a 25-to-30GHz 4- and 20-core VCOs achieving a 193.3dBc/Hz FoM at 1MHz offset is introduced. Finally, a W-band 3rd-harmonic-extraction VCO employing a multi-resonance/core/mode technique is demonstrated with excellent phase-noise-FoM balance and a 21% tuning range.

ISSCC 2023 also introduces new PLL concepts for generating RF, microwave, and mm-wave frequency carriers with several record-setting low-jitter and low-power-consumption prototypes that ultimately push the jitter-power FoM below -250dB as shown in Fig. 1. A 16GHz charge-pump PLL-based FMCW synthesizer is presented with robust duty-cycled operation achieving 1.5GHz modulation bandwidth and 41kHz rms error with below 1μs re-lock time. A 0.6-to-7.7GHz LO generator achieves 135fs rms jitter using a single LC-VCO-based with a ring-oscillator-based sub-integer-N frequency multiplier. A 9-to-11GHz fractional-N digital PLL is discussed with inverse-constant-slope DTC and FCW subtractive dithering technique. A 2.4GHz fractional-N PLL, which operates from a 32kHz reference clock and employs nonuniform-timing reference oversampling, is demonstrated with 3.94ps rms jitter. A 9.25GHz bang-bang digital PLL employs a novel multi-DTC topology with phase-shifted quantization-error sequences and achieves -60.3dBc in-band fractional spur and 77fs$_{rms}$ jitter. A W-band PLL employing novel power-gating injection-locked frequency-multiplier-based phase detector reports 47fs$_{rms}$ jitter and a -253dB FoM. Finally, a 2.4GHz ultra-low-voltage subsampling PLL operating from a 0.4V supply achieves 236fs rms jitter, -76.1dBc reference spur, and -253dB FoM. These integer-N and fractional-N PLLs continue to improve power consumption and integrated jitter to keep pace with advances in communications and sensing applications.
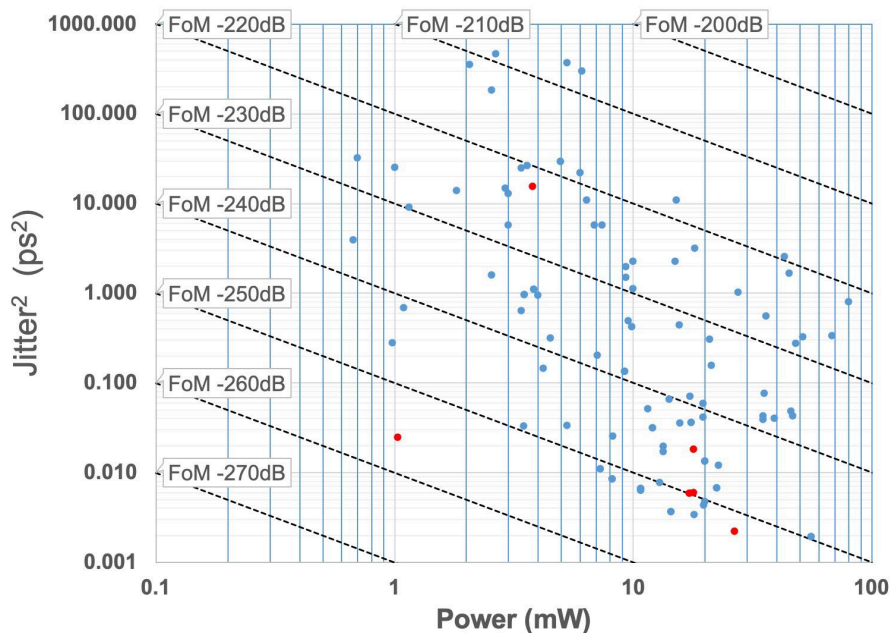

Fig. 1. PLL trends.

RF Transceiver Building Blocks: ISSCC 2023 will introduce several exciting developments in circuits applicable to GHz and mm-wave transceivers. These developments represent new benchmarks in output power and bandwidth. A quadrature digital power amplifier with eight-way power combining using IQ-reuse and Doherty techniques achieves 4.1W (36.1dBm) peak P_{out} and 33.6% peak PAE at 2.9GHz in 28nm bulk CMOS. The PA achieve the peak RF-output-power 1dB bandwidth from 2.6 to 3.2GHz and a peak PAE of more than 30% from 2.7 to 3.2GHz. A broadband mm-wave power amplifier uses asynchronously tuned coupled-resonator output match and adaptive feedback linearization to achieve 19.7-to-43.8GHz 3dB gain bandwidth and demonstrates linear performance with 5G NR FR2 waveforms from 24.24 to 43.5GHz. An E-band LNA evolves the noise-cancelling LNA with an asymmetric compensation transformer and a hybrid-phase combiner to exhibit 4.8-to-6.5dB noise figures across 70 to 86GHz in 40nm CMOS while consuming 25mW. A 4b RFDAC combined with direct-digital frequency synthesis demonstrates the generation of 4GHz-bandwidth FMCW chirps from 5 to 9GHz in 28nm CMOS with a maximum chirp slope of 800MHz/μs. An N-path filter is also presented with the center frequency tunable across 1 to 5GHz, supporting a 5-to-80MHz bandwidth, and achieving +23dBM IIP3 using a charge-pump-based clock booster and embedded frequency translation.

THz Signal Generation: ISSCC 2023 features three new contributions in signal generation from 200 to 689GHz. These developments improve peak output power, efficiency, phase noise, and tuning range. A 65nm CMOS lensless THz source radiates up to 9.1dBm output power at 675GHz from a 12×12 array while supporting a +/-45-degree H-plane beam steering, the frequency tuning range of 6.9%, and a peak DC-to-THz efficiency of 0.245%. Also, implemented in 65nm CMOS, a cascade of 2 subsampling PLLs and an H-band frequency doubler are used to generate 264 to 287GHz from a 940MHz reference. The D-band subsampling PLL adopts a dual path subsampling phase detector and achieves 75fs$_{rms}$ jitter while exhibiting -2.5dBm output power and an 8.38% tuning range. A broadband frequency doubler in 0.13μm SiGe BiCMOS uses a slotline-based transformer at the doubler input to achieve output frequency range from 200 to 350GHz, output power between 1.1 and 4.7dBm, and the peak DC-to-RF efficiency of 1.13%.

Wireless- 2023 Trends

Subcommittee Chair: *Chih-Ming Huang, MediaTek, Taipei, Taiwan*

Ultra-low-power (ULP) receivers have continued to make dramatic improvements that facilitate their widespread adoption. Selectivity, or interference rejection, is a critical metric for all receivers that operate in the presence of other incumbent transmitters. As the RF spectrum becomes more crowded, adequate SIR (signal-to-interference ratio), adjacent channel rejection (ACR), and IRR (image-rejection ratio) for heterodyne receivers are essential for scaling the number of users that can concurrently occupy a band. Figure 1 illustrates how recently published ULP receivers have steadily advanced this metric so that ULP receivers are now competitive with main connectivity radios such as WiFi and Bluetooth.

2023 in particular had an increase in UWB and back-scatter transceivers published at ISSCC. UWB, which has seen a resurgence recently, has the benefit of exceptional energy-per-bit, ability to achieve high data-rate at low power, and accurate ranging for use in asset tracking or tags. Back-scatter communication is used for communication with passive tags, eliminating the need for a battery. This year the first passive Bluetooth tag using back-scatter was published at ISSCC, demonstrating communication to a tablet with no battery in the tag.
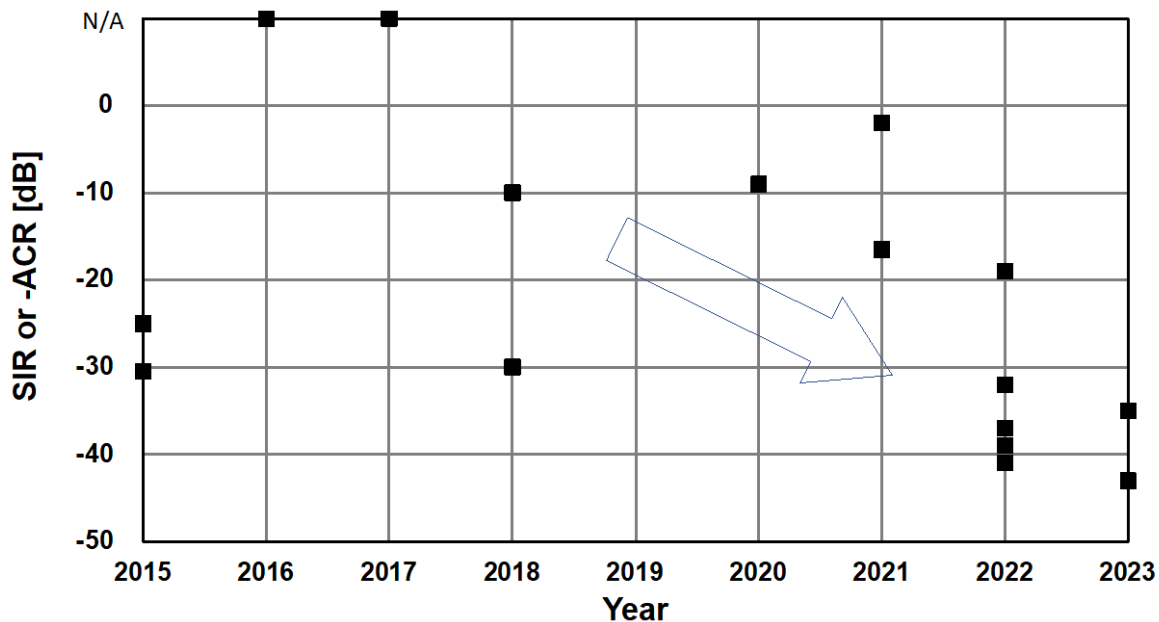

Figure 1: Plot of signal-to-interference ratio (SIR) or adjacent channel rejection (-ACR) vs. year
as reported by ULP receivers published at ISSCC.

The continuing demand for higher wireless data-rates in the context of mobile battery limitations drives the high-throughput and power-efficient transceiver development at mm-wave and sub-THz bands. Wireless transceivers continue to evolve with a higher level integration and more blocker tolerance. This year at ISSCC 2023, implemented in a 22nm FinFET technology, a 140GHz fully integrated receiver consisting of a PLL and a 16GHz ADC achieves a data-rate of 128Gb/s and consumes only 246mW. Additionally, for the blocker-rich sub-6GHz band, a receiver with harmonic-reject N-path-filter/mixer topology able to handle the 3rd and the 5th harmonic blockers as large as 10dBm and 4dBm, respectively was demonstrated with less than 1dB compression.

Figure 2 shows the trend of energy efficiency for mm-wave (<100 GHz) and sub-THz (>100 GHz) receivers presented at ISSCC. A receiver presented at ISSCC 2023 leverages a low-noise quadrature PLL, a frequency tripler and energy-efficient time-interleaved ADCs to demonstrate <1.95pJ/b system efficiency and < 1pJ/b receiver front-end efficiency while achieving 160Gb/s with the front-end.
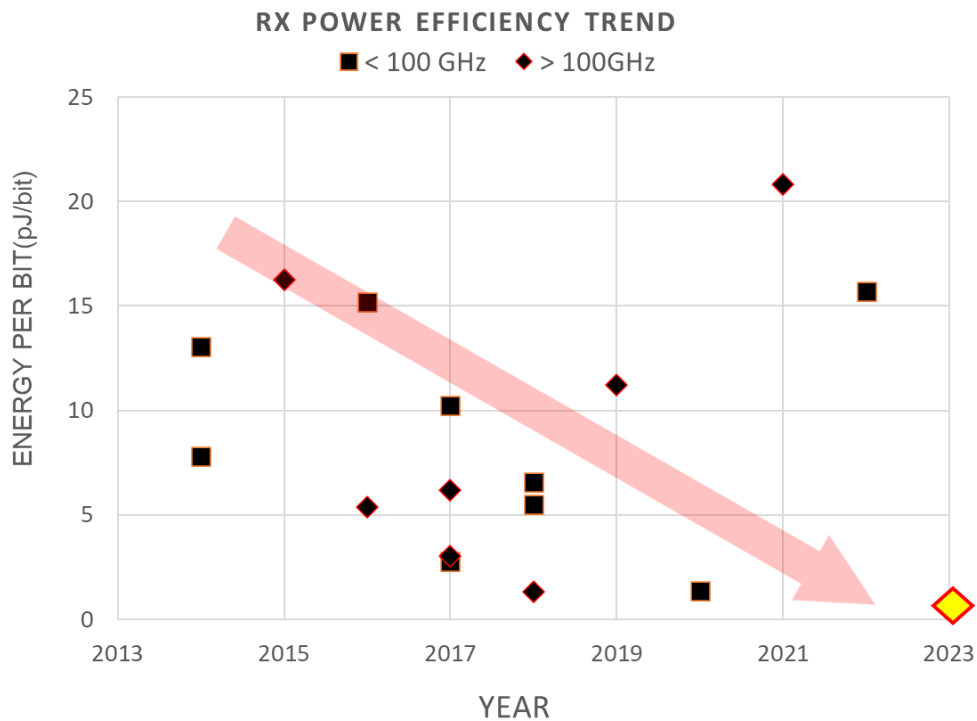

Figure 2: Receiver power efficiency trend of mm-wave and sub-THz high-speed wireless communications

Wireline – 2023 Trends

Subcommittee Chair: *Yohan Frans, AMD, San Jose, CA*

Over the past few decades, electrical and optical interconnects have been key components bridging the gap between the exponentially growing demand for data bandwidth across electronic systems and the relatively gradual increase in pin/cable density. Ranging from handheld electronics to supercomputers, wireline data bandwidth must also grow exponentially to avoid limiting the performance scaling of these systems. By increasing the data per pin or cable of various electronic devices and systems, such as memory, graphics, chip-to-chip fabric, backplane, rack-to-rack, and LAN, wireline I/O has fueled incredible technological innovation in electronic devices and systems for decades. Figure 1 shows that data rate per pin has approximately doubled every four years across various I/O standards ranging from DDR, to graphics, to high-speed Ethernet. Figure 2 shows that the data rates for published transceivers have kept pace with these standards while taking advantage of CMOS scaling. Figure 3 shows published transceiver energy efficiency vs. channel losses at the Nyquist frequency in the 40-to-50dB range. In part, this incredible improvement is enabled by the power-performance benefits of process technology scaling. However, sustaining this exponential trend for I/O bandwidth requires more than just transistor scaling. Significant advances in energy efficiency, channel equalization and clocking must be made to enable the next generation of low-power and high-performance computing systems. Papers at ISSCC this year include examples of long-reach PAM-4 copper interconnect transceivers operating up to 112Gb/s with and without ADC+DSP equalization architectures. These include:

1. A PAM-4 medium-reach electrical receiver operating at 200Gb/s with a continuous-time FFE.
2. A PAM-8 short-reach electrical transmitter operating up to 100Gb/s.
3. A 32Gb/s die-to-die chiplet NRZ transceiver with high beachfront density
4. An NRZ WDM receiver module using 7 wavelengths and 50Gb/s/wavelength to achieve 350Gb/s aggregate throughput.
5. A PAM-4 optical receiver operating at up to 106.25Gb/s
6. A short-reach optical coherent receiver operating up to 24Gb/s.

New techniques for extending data rate, power reduction, channel equalization, and clock recovery are reported. These transceivers and building blocks are implemented in CMOS technology.

Scaling Electrical Interconnects to 100Gb/s and Reaching Out to >200Gb/s

Bandwidth requirements in data centers and telecommunication infrastructure continue to drive the demand for ultra-high-speed wireline communication. Recently, complete transceivers operating up to 112Gb/s were demonstrated across a long-reach copper channel with >45dB loss. Two notable trends in these transceivers, especially for long-reach channels, are the adoption of PAM-4 modulation and a transition to DAC/ADC architectures with DSP-based equalization. Although PAM-4 provides twice the data rate at the same baud rate as conventional NRZ to relax channel loss requirements for bandwidth doubling, it also comes with more stringent requirements for linearity and noise. This trend has motivated the development of low-power data converters, digital equalization and clock recovery along with linear, high-bandwidth TX and RX analog front ends. This year, ISSCC includes two implementations of 112Gb/s PAM-4 long-reach transceivers with low power consumption. In paper 6.1, Broadcom demonstrates a non-ADC/DSP transceiver in 7nm CMOS for a wide range of data rates up to 112Gb/s for long-reach copper interconnects consuming 690mW for a channel with 43dB loss. In paper 6.2, MediaTek presents an ADC/DSP-based 112Gb/s PAM-4 long-reach transceiver in 5nm CMOS operating over a 48dB loss channel consuming 521mW. In paper 6.3, Peking University describes a 200Gb/s 5-tap coplanar-waveguide distributed-tap receiver FFE in 28nm CMOS. In paper 6.5 and 6.6, MediaTek and Korea University demonstrate low power and fast-frequency-acquisition CDR architectures operating at 32Gb/s and 52Gb/s, respectively. In paper 6.7, Peking University describes a 128Gb/s PAM-4 transmitter with improved transitions between non-adjacent levels. In paper 6.8, Hanyang University demonstrates a 100Gb/s PAM-8 transmitter with 3-tap shuffler FFE in 40nm CMOS yielding an output swing up to 1.6V_{ppd}.

In-Package Links for Chiplet Communications

As a consequence of the increasing demand for bandwidth in high-throughput systems used in AI, HPC and switch applications, multiple devices are integrated in the same package, and data is sent between chiplets on the same interposer. For these 2.5D package applications, relatively short distances have to be bridged with minimum power while targeting the highest possible throughput per millimeter of chip-edge (Gb/s/mm). Since channel attenuation and discontinuities in these links are mild, low-power

analog-oriented equalization and forwarded clock architectures are adequate. In paper 6.4, Samsung presents a 32Gb/s per lane NRZ XSR transceiver in 4nm CMOS achieving 8Tb/s/mm beach-front bandwidth density while consuming only 0.44pJ/b. It operates over 3mm silicon interposer signal traces, which are shielded against crosstalk.

Optical Links for Upcoming 400G Data Center Interconnects
The explosive growth of data and data-centric computing places stringent demands on the bandwidth and energy efficiency of data center interconnects, spurring the development of several 200-to-400G Ethernet standards. Low-power data converters and optical integration are the two key components for the development of high-performance optical pluggable modules using coherent detection. In Paper 12.1, AMD demonstrates an NRZ WDM receiver module with 350Gb/s aggregate data rate and 1e-12 BER without forward error correction. This design has 7 wavelengths in a fiber, with 1.5nm color spacing, and runs at 50Gb/s per wavelength. The module incorporates stacked 7nm CMOS and 45nm silicon photonic dies. It uses an array of thermally tuned cascaded ring resonators to distinguish laser colors. High sensitivity (-11.1dBm median) was measured under 0.96pJ/b energy efficiency. In paper 12.2, Cisco Systems and University of Illinois Urbana-Champaign demonstrate a differential TIA in an optical 106.25Gb/s PAM-4 receiver with asymmetric signal paths utilizing currents from both terminals of the PD to improve SNR, resulting in an optical sensitivity of -14dBm. In paper 12.3, University of Illinois Urbana-Champaign reports a CDR of a 24Gb/s QPSK coherent optical link in 28nm. It applies 16-phase switched-inverter-based harmonic-rejection-mixers to achieve low CDR latency.

Concluding Remarks:
Continuing to aggressively scale I/O bandwidth is essential for the industry, but the tradeoffs between bandwidth, power, area, cost and reliability are extremely challenging. Advances in circuit architecture, interconnect topologies, transistor scaling and integrated silicon photonics are changing how I/O will be done over the next decade. The most exciting and promising of these emerging technologies for electrical and optical interconnects will be highlighted at ISSCC 2023.
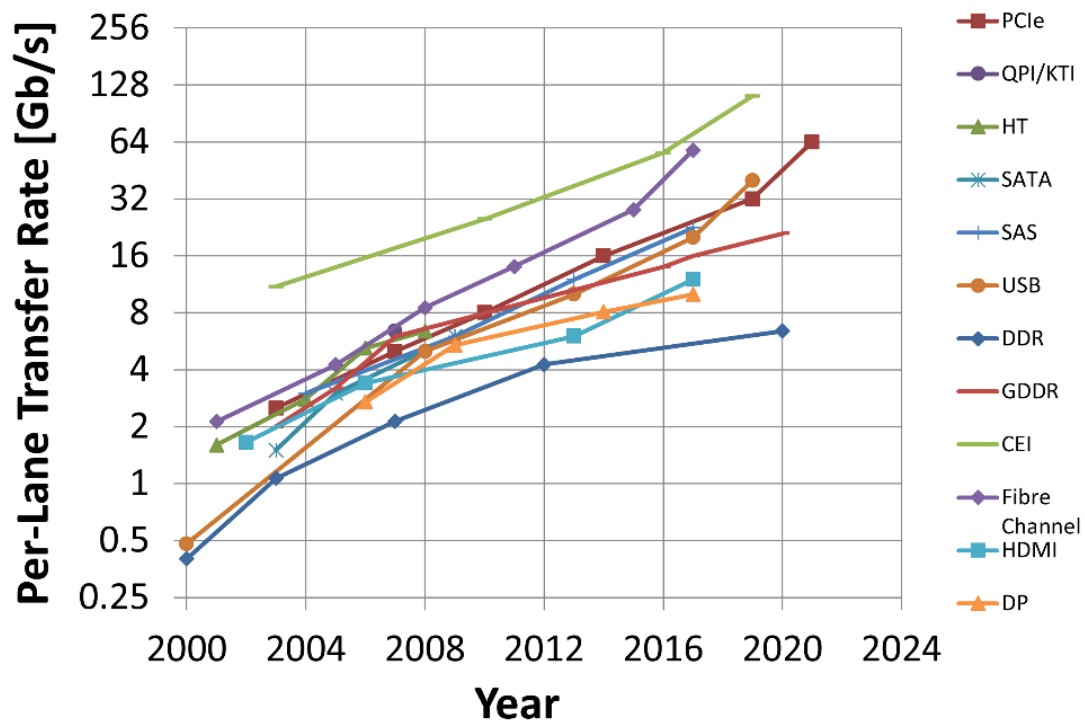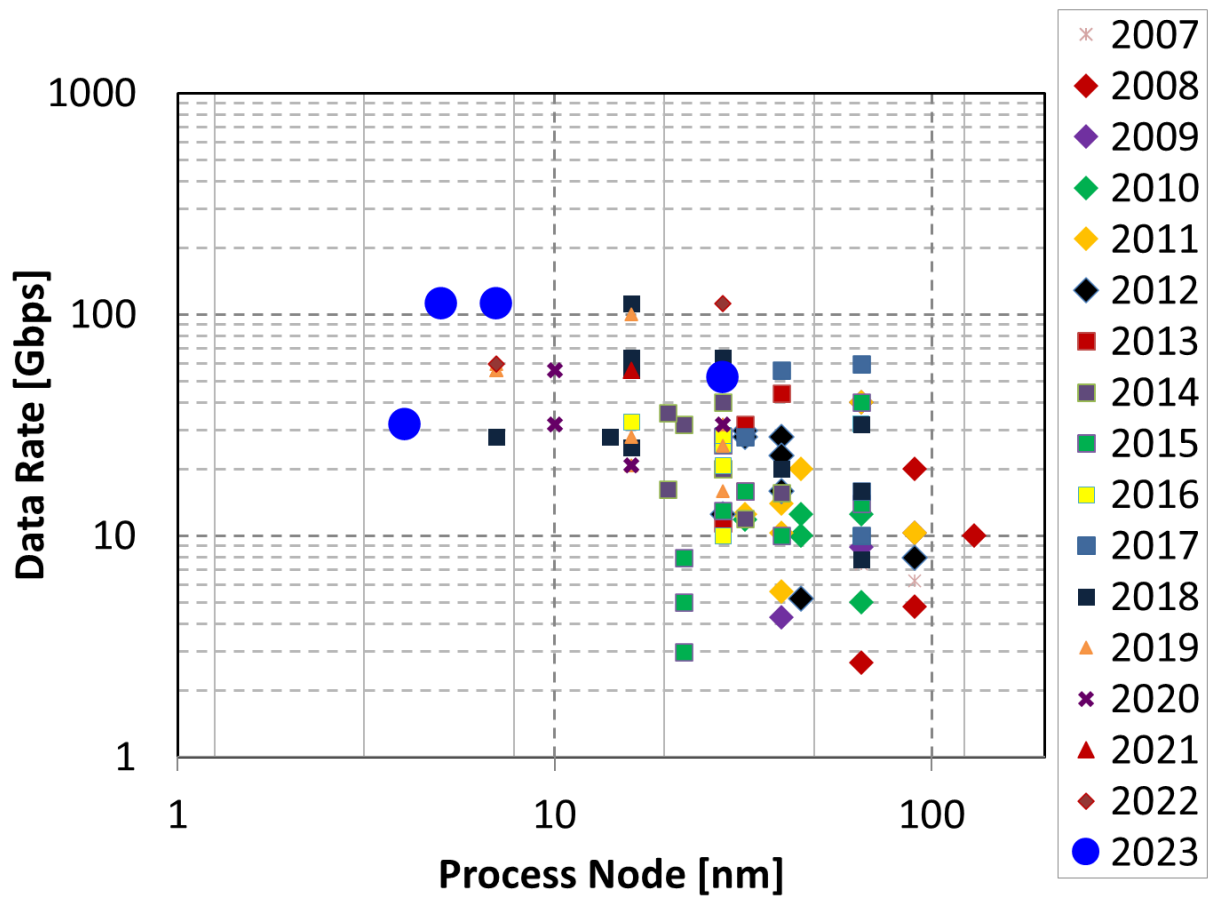

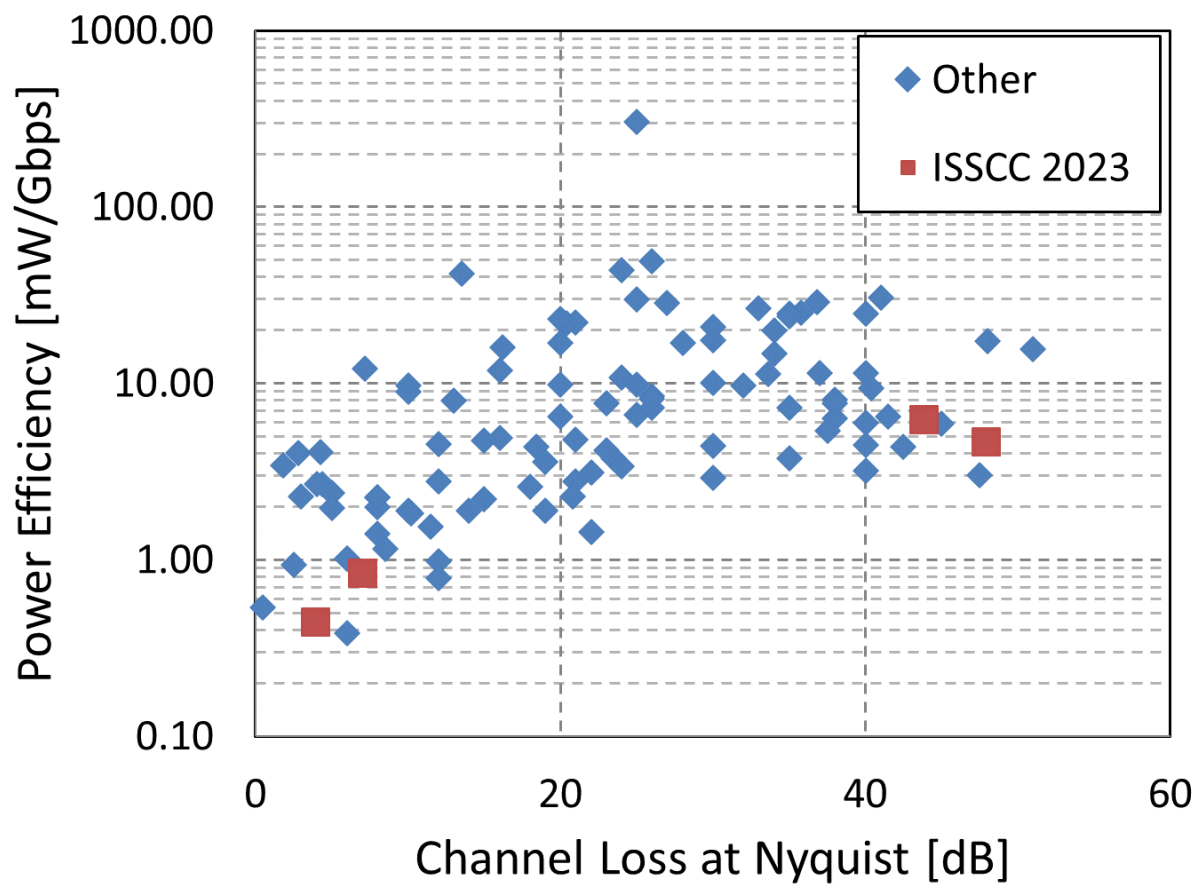Figure 1: Per-lane data rate vs. Year for a variety of common I/O standards.

Figure 2: Data-rate vs. process node and year.

Figure 3: Power Efficiency vs. channel loss and year.

HISTORICAL TRENDS IN TECHNICAL THEMES
DIGITAL SYSTEMS
DIGITAL ARCHITECTURES & SYSTEMS SUBCOMMITTEE
DIGITAL CIRCUITS SUBCOMMITTEE
MACHINE LEARNING & AI SUBCOMMITTEE
MEMORY SUBCOMMITTEE

Digital Architectures & Systems (DAS) – 2023 Trends

Subcommittee Chair: Thomas Burd, *Advanced Micro Devices, Santa Clara, CA*

This year's selection of processor papers highlights the industry adoption of the most advanced CMOS technology at the 4-5nm node. Innovative packaging technologies, including 3D stacking and direct bonding are being productized, which supports easy integration of multiple process nodes into a single socket. This has also fueled an exponential increase in on-system memory   that drives increased performance. The drive to higher clock frequencies, which having slowed, still continues to tick higher, and is supplemented by a drive to increased core counts. Bump and through-silicon-via pitches continue to scale down at a rapid rate, enabling tremendous increase in bandwidth across multiple dies. The mobile CPU continues to increase in both frequency and performance, while providing a wide range of performance and energy efficiency.
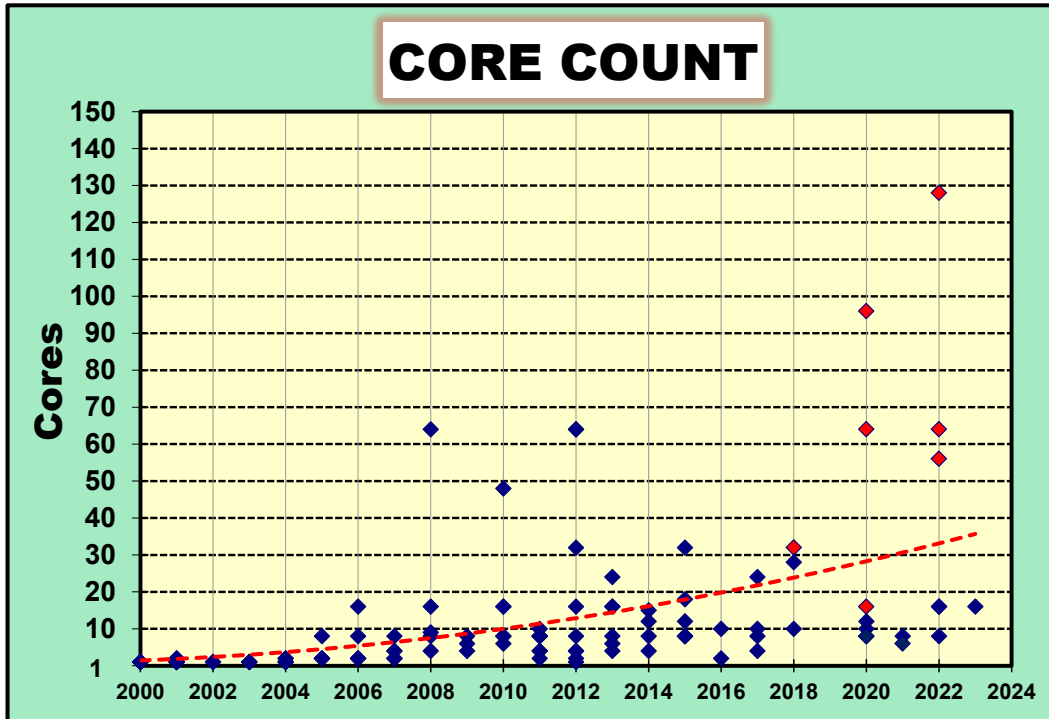

Figure 1: Core-count trends (red diamond designates multi-chip module).

Figure 2: Die counts in a system trends.

Figure 3: Chip-complexity scaling trends (red diamond designates multi-chip module).
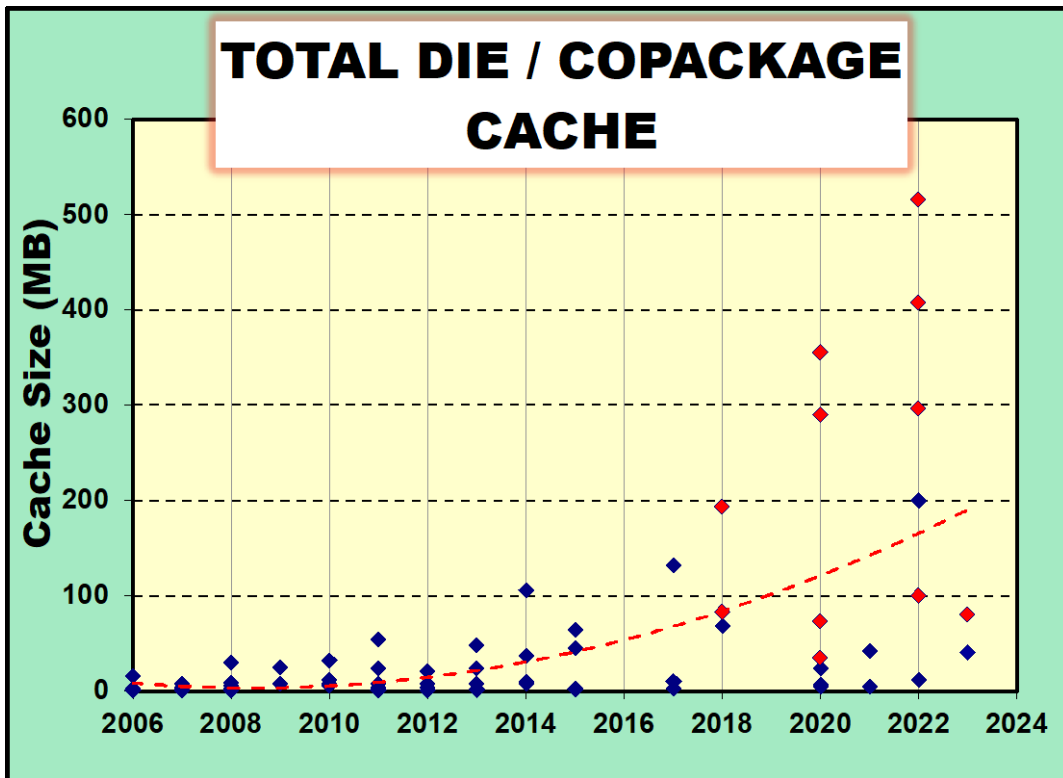
Figure 4: On-die cache-size trends (red diamond designates multi-chip module).

We see a continued push on application-processor performance and efficiency, along with video, display, and camera capabilities to fully leverage the ubiquitous connectivity smartphones provide. There are also major innovation efforts in 5G, artificial intelligence (AI), gaming and thermal management. 5G cellular technology is becoming more mature and the post-5G era is gaining more focus. The post-5G era will feature more antennas, more intelligence, and more use cases. The range of applications of neural network processing units (NPUs) is gradually expanding beyond image and speech recognition, to cellular performance improvement, SoC power and performance optimization. Therefore, not only the performance of the NPU is increasing, but also tiny NPUs for low-power operations are being applied everywhere. For better user experience and game quality, the main concern surrounding the display is moving from resolution to frame rate.

	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023
Graphics	OpenGL (ES1.1)		OpenGL/VG/MAX (ES2.0)		AR (Augmented Reality)					VR (Virtual Reality) Vulkan							
Display	VGA	WVGA @60fps	SXGA @60fps		WQXGA/WQXGA+ @60fps		WQXGA/WQXGA+ @60fpsX2 (VR)		WQXGA/WQXGA+ @120fps			4K 240fps				QHD+	
Camera	5-8M	10M	16M	20M	24M		12MxDual 360°VR		12MxDual 360°VR		48M Triple		200M Quad				Penta
Image/Video	H.264/AVC (VGA)	H.264/AVC (D1)	H.264/AVC (Full HD)		H.265/MVC H.264/SVC		H.265/VP9			H.265/VP9 HDR		AV1 HDR10+			8K @30fps		
Audio	AAC	AAC Plus	WMA Dolby 5.1		Dolby TrueHD/Digital+			DSD Dolby Atmos				TWS Truly Wireless					
Accelerator	FPU	SIMD Multicore (2~4)		Multicore (4~8)		Heterogenous Multiprocessing			Neural Net Processor		5 TOPS		15 TOPS				
Downlink [Mbps]	UMTS 0.4 ~ 2	HSPA 1.8 ~ 7	HSPA+ 7 ~ 42		LTE 100		LTE-A 150 ~ 750		LTE-A 1600		LTE-A 2000		5G 5000		5G 10000		
CPU [MIPS]	300 500	500 800	800 2400		2400 6000	6K 12K	12K 100K		13K 112K		19K 162K		22K 180K		26K 208K		

Figure 5: Application-processor trends in smartphones.

Circuits for Hardware Security: With the increasing risk and cost of information theft and safety hazards, hardware security has become a common requirement in intelligent and connected systems. Though focus on cryptographic implementation continues, cost-effective and low bit error rate PUFs (physically unclonable functions) are increasingly adopted in smart cards, sensor nodes, consumer devices, and automotive. TRNGs (true random-number generators) are also commonly required to strengthen secret key generation in cryptographic applications. Techniques to counteract side-channel attacks are enabling higher levels of security at lower design cost, thanks to the higher degree of design reuse and more digital circuit techniques. Counteraction of fault injection attacks is also becoming more common thanks to techniques ranging from logic-level fault detection to physical sensors. Quantum computers allow dramatic speed-ups in attacks on existing public-key algorithms. Standardization bodies such as NIST have started competitions to identify potential post-quantum cryptographic (PQC) schemes. Novel PQC accelerators are now being designed to efficiently and securely implement these schemes in hardware. Improvements in the efficiency of homomorphic encryption are being demonstrated to preserve both data usability and privacy in commercial cloud environments.

Figure 6 illustrates trends in area scaling in PUFs (area/bit) and TRNGs published recently at ISSCC, showing relentless area and cost reductions. With regards to techniques counteracting side-channel attacks (EM and power), Figure 7 shows the progressive improvement in the measurements-to-disclosure of cryptographic keys, as determined by the ratio of the power trace count necessary for a successful attack under protected and unprotected designs.
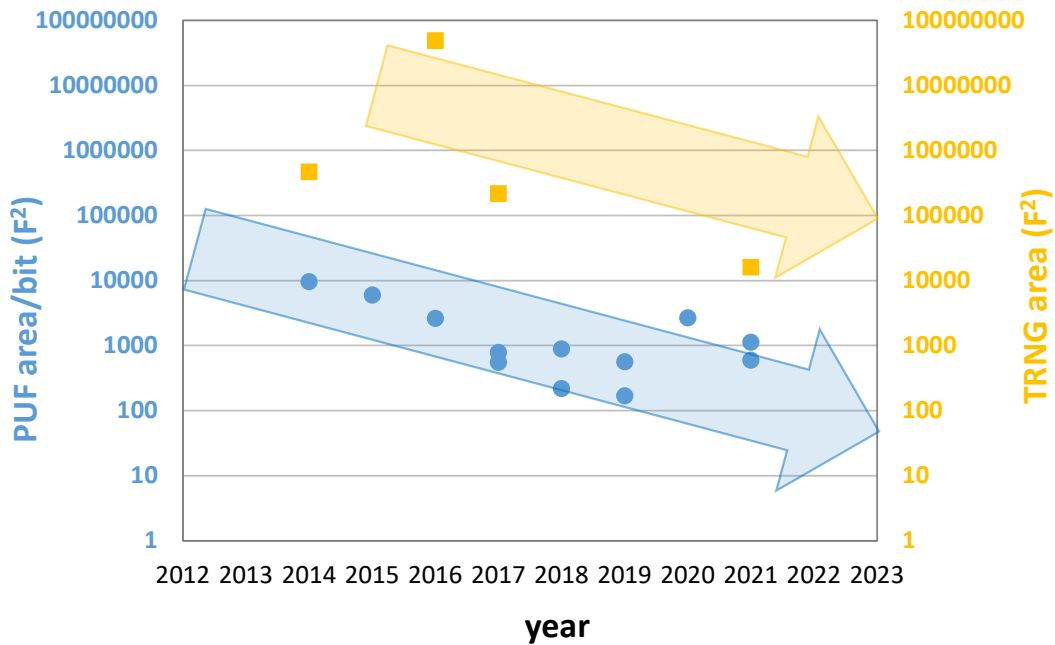
Figure 6. Area/bit trends for physically unclonable functions (PUFs) and area trends for true random number generators (TRNGs) published recently at ISSCC.
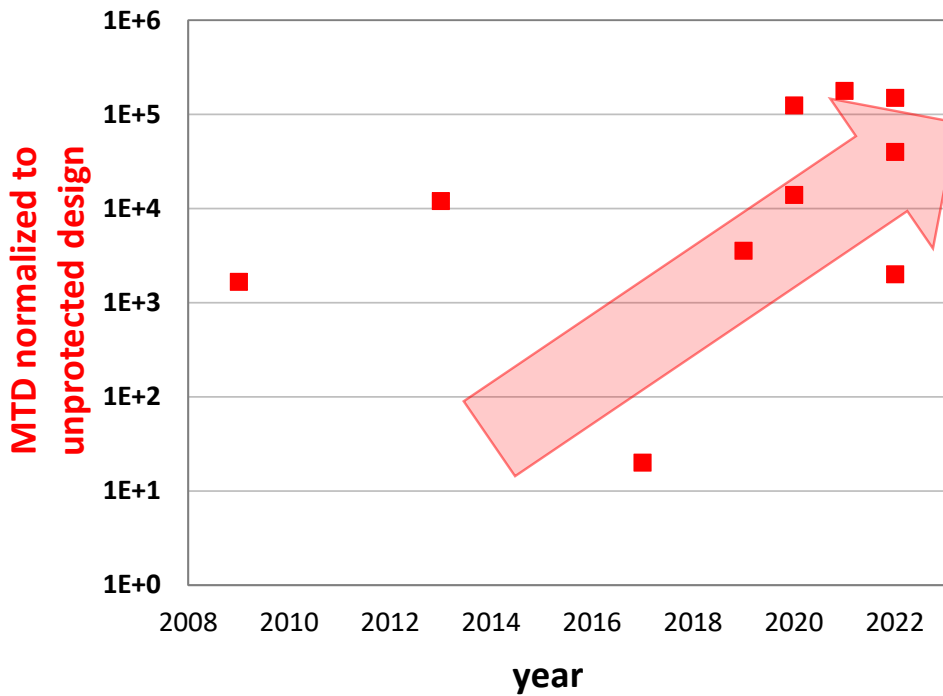

Figure 7: Improvement in measurements-to-disclosure (MTD) of cryptographic keys of side-channel counteraction techniques (normalized to unprotected design).

As shown in Figure 8, security primitives for the root of trust continue to evolve through design approaches enabling a higher level of integration with the existing silicon infrastructure (e.g., logic, memory), and merging multiple functions within unified designs. Some of these primitives are being employed in applications requiring high safety standards, such as automotive. Innovation continues in cryptographic accelerators and processors providing improvements in energy efficiency and flexible adaptation to pre- and post-quantum cryptography, as well as homomorphic encryption. The robustness against side-channel attacks continues to improve with a relentless increase in the mean traces to disclosure (MTD), which exceeds 100,000× under digital LDO- and machine learning-based protections. Countermeasures against semi-invasive and invasive attacks are devised to push physical security high enough that it

does not become the weak link, while improving security on other fronts and types of attack. Solutions from sensing to logic are demonstrated to counteract fault injection, reverse engineering and power glitching attacks. In wireless links, physical security is also being enforced through secure directional links, which allow correct data reception only across the intended transmission direction.
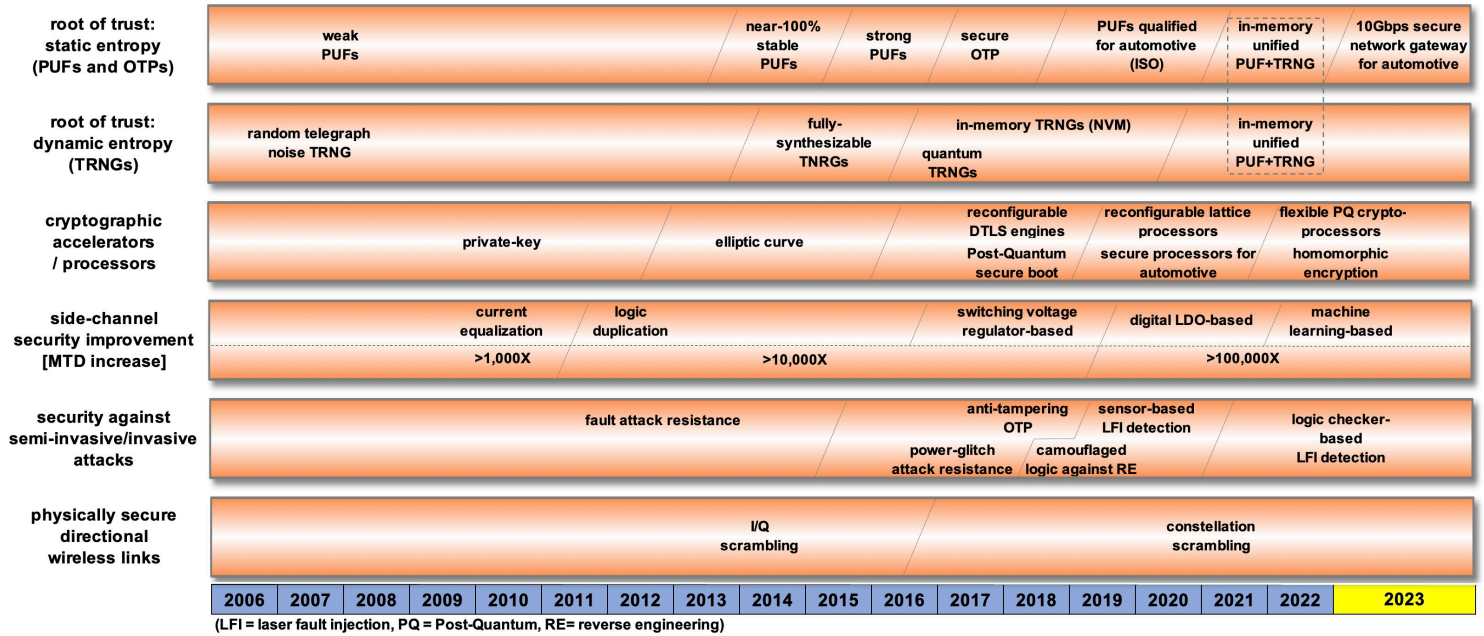

Figure 8: Security trends.

Digital Circuits – 2023 Trends

Subcommittee Chair: *Keith Bowman, Qualcomm, Raleigh, NC*

The demand for more flexible and energy-efficient platforms ranging from embedded sensors to large neural engines continues to drive innovations in CMOS digital circuits and accelerators with goals of improving performance and energy efficiency. Digital circuit innovations also benefit from emerging technologies, such as non-volatile memories or deep-trench capacitors.

A continued trend in application-specific accelerators is the development of new circuit techniques that benefit a range of emerging applications, such as optimization problems in Ising or SAT-solvers, artificial intelligence at the edge for perception, speech recognition or advanced telecom channel decoders. Many of these accelerators leverage compute-in-memory (CIM) architectures or flexible tiled processing elements, including hybrid system partitioning between mixed-signal and digital-signal processing in memory, dedicated pipelines and machine learning in recurrent or spiking neural networks, while employing non-volatility with the right combination of all these techniques to optimize energy efficiency. Furthermore, custom data representation appears as a strong trend to reduce power consumption, whether with charge-domain computation or non-conventional number representation, along with an emphasis on data sparsification to skip unnecessary computations thanks to architectural clock or power gating. On the circuit side, we observe an increasing trend for custom memory bitcells integrating computing elements or non-volatility.

In addition, continued improvements in traditional digital circuit blocks for integrated power management and clocking are permitting new usage scenarios.

Integrated Voltage Regulators: Compute-intensive digital loads have been pushing the demand for high current loads of integrated voltage regulators. Block-level regulation with digital low-dropout (LDO) linear regulators is maturing for integration into scaled process nodes to allow multiple processor cores on the same input voltage rail to operate at unique voltages according to the core workload. Concurrently, high-efficiency voltage down-conversion has driven inductor-based regulators (LCVR) and switched-capacitor voltage regulators (SCVR) at finer granularities for dynamic voltage and frequency scaling (DVFS) of individual functional blocks. With the demand for high loads, the challenge for LDOs and down-conversion regulators is to maintain good area efficiency for large-scale integration, without sacrificing energy efficiency and response time. This year's papers address flexible output voltages under the constraint of high current density, limiting the dependency between input and output voltages with adaptive linear transfer function or adaptive ganged capacitors. Figure 1 describes the achieved current density of these integrated voltage regulators across calendar years, indicating a breakthrough improvement of an order of magnitude in current density.
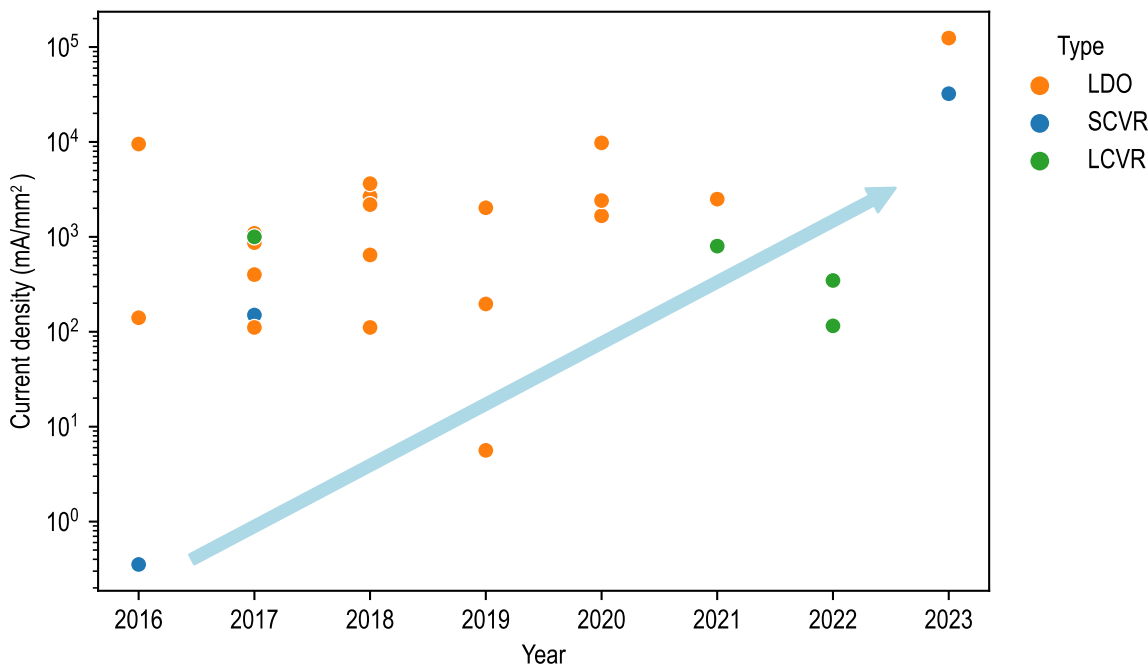

Figure 1. Integrated voltage regulator trends in current density with respect to active area.

Digital Clocking Circuits for Low-Jitter Applications: Clock generators continue an architectural migration from historical PLLs to multiplying delay-locked-loops (MDLL), injection-locked clock multipliers (ILCM) and fractional output dividers (FOD) to provide more functionality, variability management and lower design complexity at advanced nodes. Demand for compact low-jitter clocking circuits continues to increase with integer or fractional multiplication ratios with respect to a low-frequency reference. In addition, power and area reductions achieved by digital and mixed-signal PLLs allow new usage models as analog functional block drivers, but these new usages come with additional signal integrity constraints, leading to the development of digital circuit techniques for spurious-tone (Spur) cancellation due to frequency mixing with a reference or a fractional multiplier. Figure 2 describes a key figure of merit (FoM) combining jitter and power for digital clock generators across calendar years, highlighting a continued trend in FoM reduction to the point that fractional multiplier architectures now compete with integer multiplier architectures, while the second plot shows the major effort on reference and fractional spur cancellation.
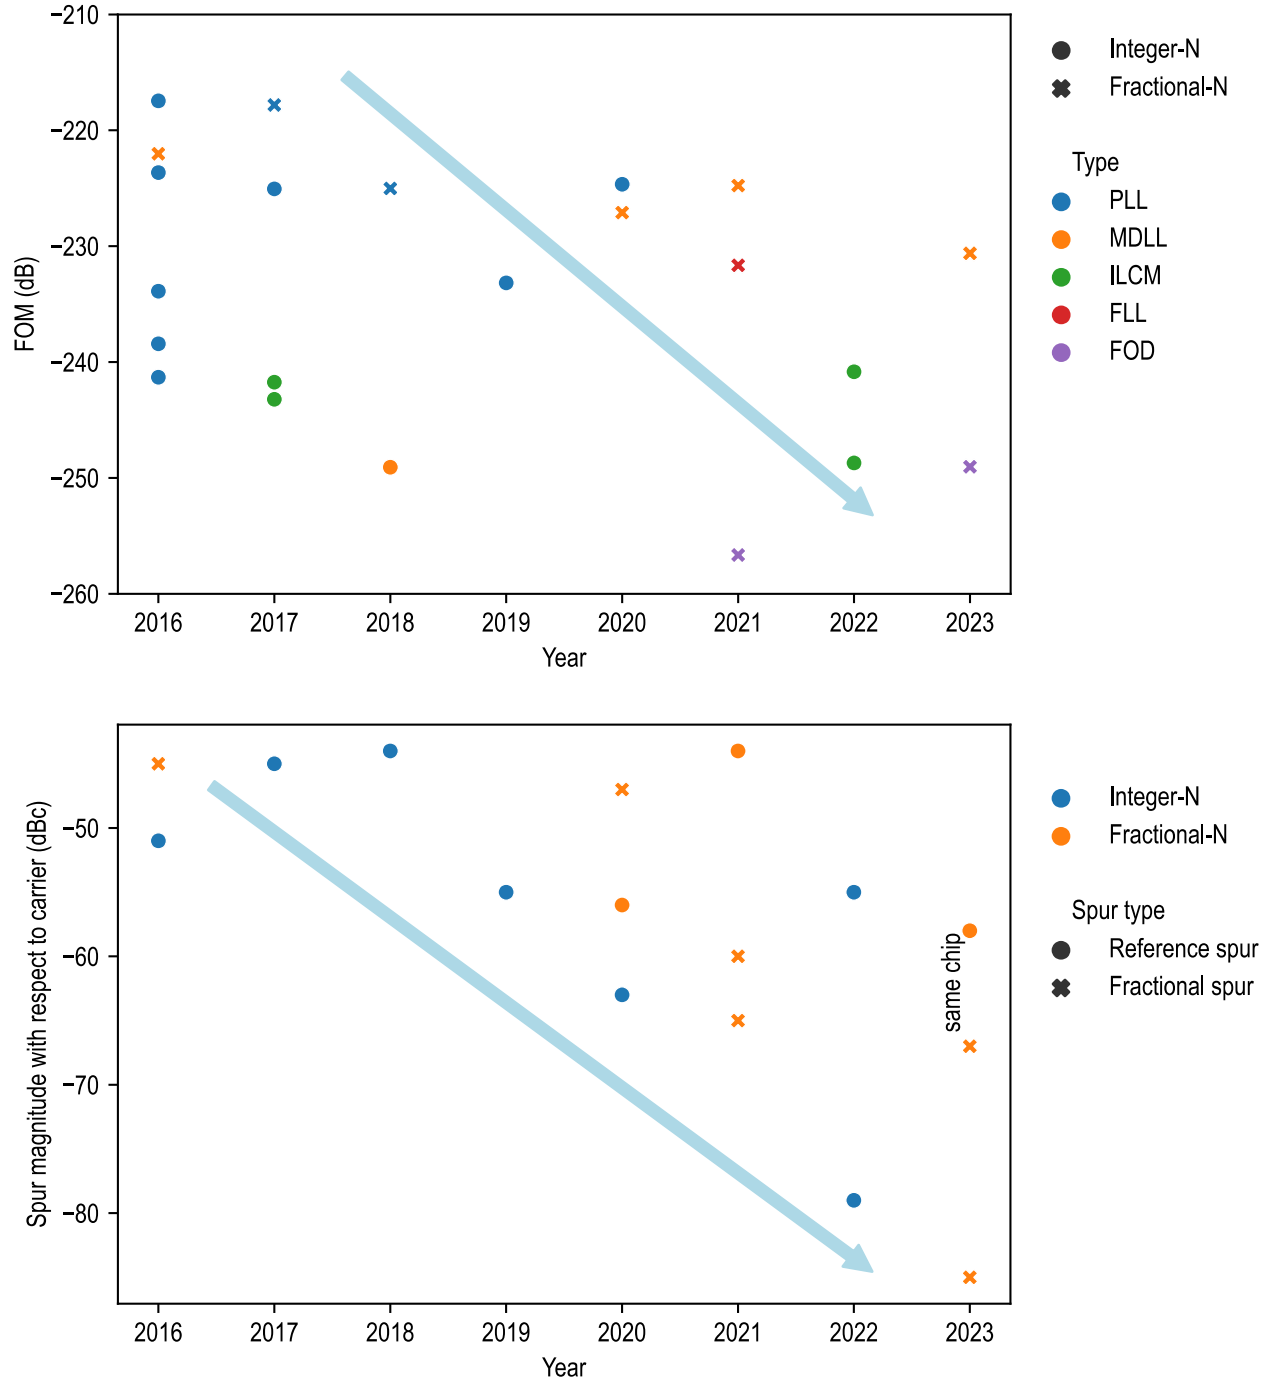

Figure 2. Top: digital clock generators key figure of merit (FoM) across recent years, defined as:
FoM = $10 \times \log_{10}\{$ (Jitter$_{RMS}$/1s)$^2 \times$ (Power/1mW) $\}$; bottom: spurious tone reduction trends.

Machine Learning (ML) & AI – 2023 Trends

Subcommittee Chair: SukHwan Lim, Samsung Electronics, Mountain View, CA

In response to the growing interest in deep learning in recent years, ISSCC established a subcommittee dedicated to machine learning & AI in 2020. As deep neural networks (DNNs) are becoming widely deployed in wearables, mobile devices, edge and servers, the size and computational complexity of these DNN models increases dramatically. This results in increasing demand for higher efficiency and performance of neural-network computing chips. This year's submissions are grouped into two sessions. First, a full session has been formed to discuss heterogeneous ML processors in support of the popular DNN models. Due to differing requirements in various applications, general-purpose processors or domain-specific accelerators are deployed. They optimize for a wide variety of applications ranging from sparse transformers, to point cloud networks to spiking neural networks (SNNs) as shown in Figure 1. Deploying heterogeneous cores allows the designs to tackle a diverse set of layers and compute needs at the expense of additional area (Figure 2). Another full session on compute-in-memory (CIM)-based processors for ML includes contributions describing CIM processors supporting different reconfigurable circuit modes (NMC/IMC, SLC/MLC), while making use of several memory technologies (RRAM, eDRAM, SRAM) for various applications (DNN training, multimodal transformer, beyond-NN applications).

It is important to note that metrics that matter at the system level are energy-per-inference (or -per-training-example), and inferences/second (or training-examples/second) on a specific task at a *given inference (or final trained) accuracy*. This year's submissions significantly push the state-of-the-art on the efficiency and throughput metrics yet again, often by combining multiple enhancement techniques within a single chip (or multiple chiplets). One important aspect to note is that increased attention is paid to non-convolutional operations, as they become the bottleneck once the convolutions are accelerated significantly.
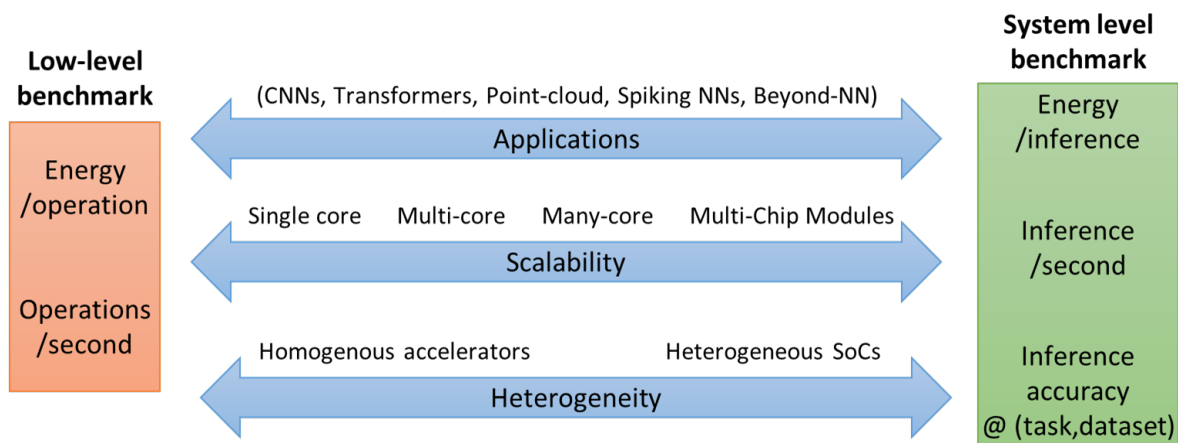

Figure 1: Various parameters impacting low-level and system-level benchmarking metrics.

Some of the key emerging trends in machine learning and AI at ISSCC are:

1. Compute-in-memory (CIM) processors are becoming more popular. They make use of SRAM, RRAM and eDRAM technologies. Digital CIM continues the trend of system-level integration to avoid excessive data movements (especially off-chip). Design directions trend towards increased efficiency, reconfigurability and flexibility by using multi-mode memory arrays, mixed NMC and IMC circuits, as well as unified fixed-point and floating-point units.

2. Exploitation of sparsity in various forms has been an important focus of both inference and training acceleration. Emerging techniques continue in the direction of workload reduction via skipping unnecessary compute that range from sparse convolution, sparsity generation using forward gradient, entropy-based early exiting and local attention use. These techniques aim to leverage sparsity in convolution kernels, forward gradients, execution paths and attention spans, respectively.

3. Incorporating the concept of dynamic adaption into ML processors is featured in several papers, which leads to better energy efficiency, while minimizing accuracy loss. Typical dynamic tuning knobs include data precision, dataflow in the pipeline, resource switching as well as body biasing.

4. Attention to SNNs has increased due to their low power consumption. An asynchronous SNN processor is proposed to realize the ultra-low power on-device training by using software-hardware co-design. A hybrid approach to both SNNs and CNNs has been presented as well.

5. Deploying a heterogeneous set of cores within an SoC has become more commonplace. As workloads such as deep convolutions are optimized heavily, the remaining operations from a diverse set of operators become the new bottleneck necessitating the need to handle other operations well. A heterogeneous architecture has been presented in ISSCC 2023, which can take the "best of both" approaches and operate at a point that is more optimal. For example, an SNN engine is combined with a CNN engine to be a complementary DNN processor for reducing both the inference and training power. It leverages task switching between the ultra-low-power SNN core and the accurate CNN core.

6. Domain-specific ML processors are covering more novel application scenarios, ranging from transformer, point-cloud network, speech enhancement, object detection and tracking, AI-IoT and beyond-NN applications. Several domain-specific architectures have been proposed using HW/SW co-design methodologies to efficiently execute specific computing and storage operations.

7. ISSCC 2023 shows machine learning processors being implemented across a wide variety of technologies (CMOS, FD-SOI), as well as with package-level innovations (MCM).
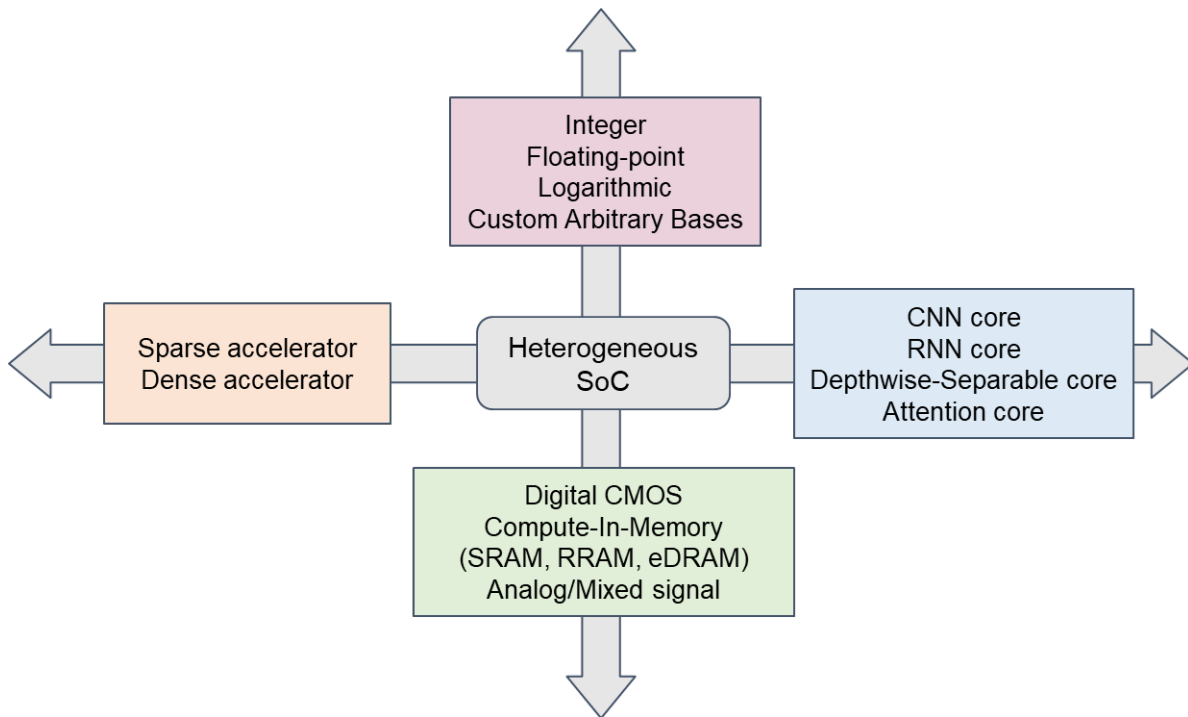

Figure 2: Heterogeneity in deep-learning accelerators.

As different accelerators or processors are often characterized on a different set of tasks, network topologies, and accuracy levels, a direct comparison of the true system-level benchmarking metrics (e.g. energy/inference or inference/s) is not always straightforward. Therefore, it is instructive to look at the reported low-level metrics of operations/s and energy/operation within the neural network. Figure 3 displays the energy efficiency vs. throughput operating points demonstrated by the accelerators presented at ISSCC 2023 (red), compared to the state-of-the-art in 2016-2020 (blue), 2021 (black), and 2022 (green). Figure 4 plots the evolution of both energy efficiency and area efficiency (throughput-per-unit-area) over the past few years.

From these graphs, the improvement in terms of low-level metrics of operations/s and energy/operation is not very apparent. ML accelerators are clearly still improving at a very fast, almost exponential pace. Yet, ISSCC attendees should keep in mind that these TOPS/W and TOPS specifications depend strongly on the level of integration of the chip, and on the particular neural-network topologies being used. We have been seeing a clear trend towards more complete integration, in which the highly efficient MAC compute arrays that were introduced over the past years are now integrated into full processing systems. Going forward, as the field matures, we believe that a common benchmarking methodology must be established which can properly account for the application context and provide

proper translation between low-level and system-level performance metrics [1]. In the meantime, clever combinations of sparsity, variable precision and in-memory computing technologies are continuing to enhance deep-learning processor efficiency and throughput. With the increase of system-level integration of machine-learning engines together with other important subsystems (imaging chips, image pre-preprocessing, audio filtering and pre-processing, etc.), these performance improvements will continue to open up new AI applications.
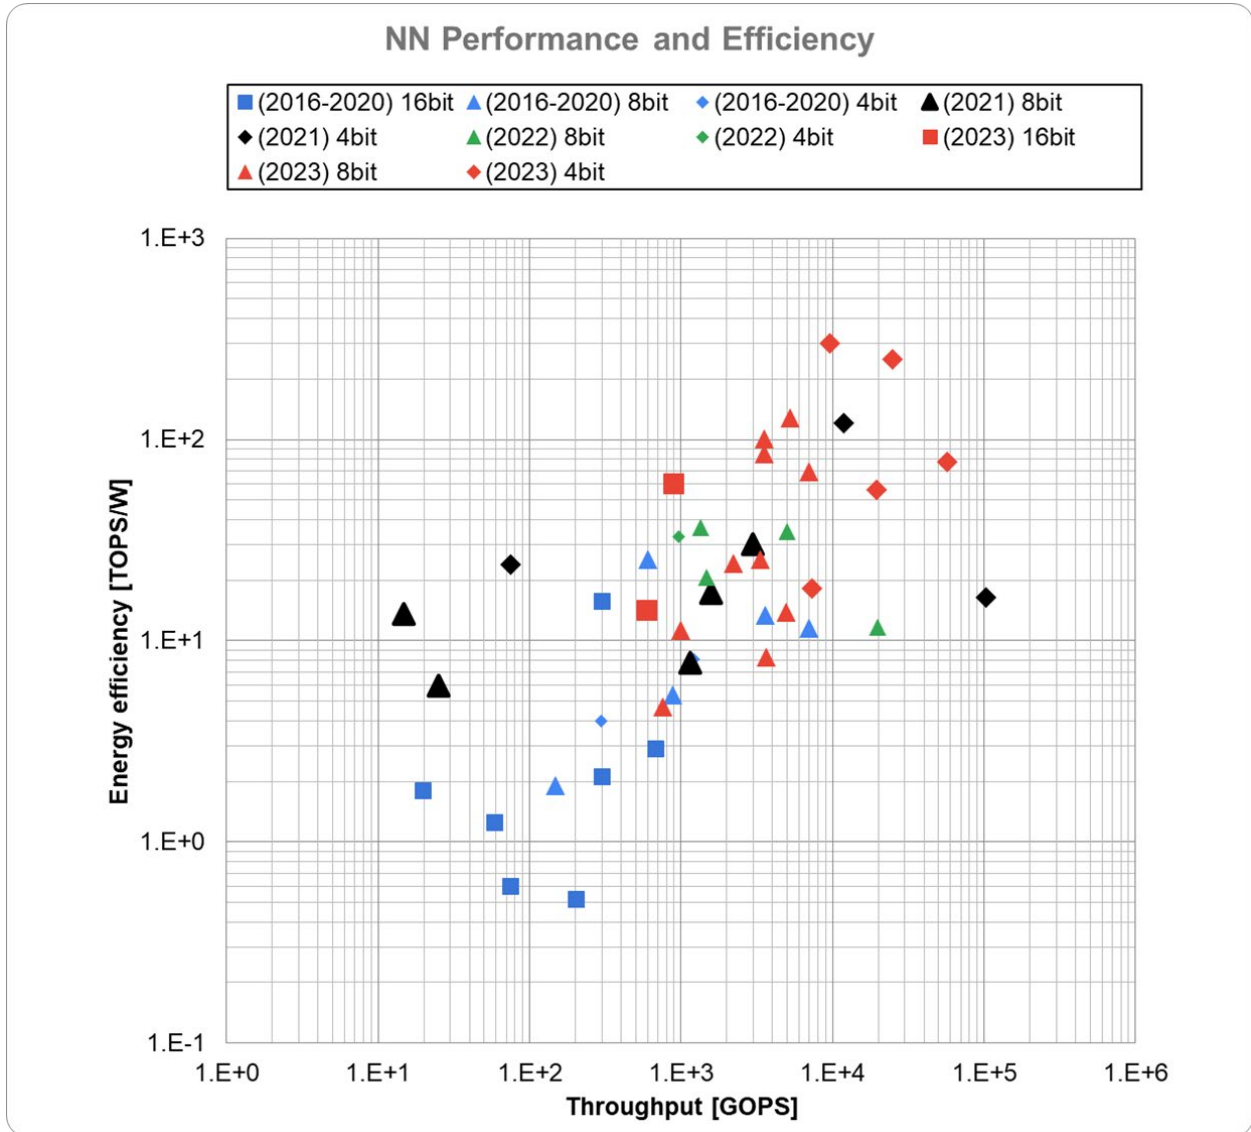

Figure 3: Deep-learning processor energy-efficiency (TOPS/W) and throughput (GOPS).
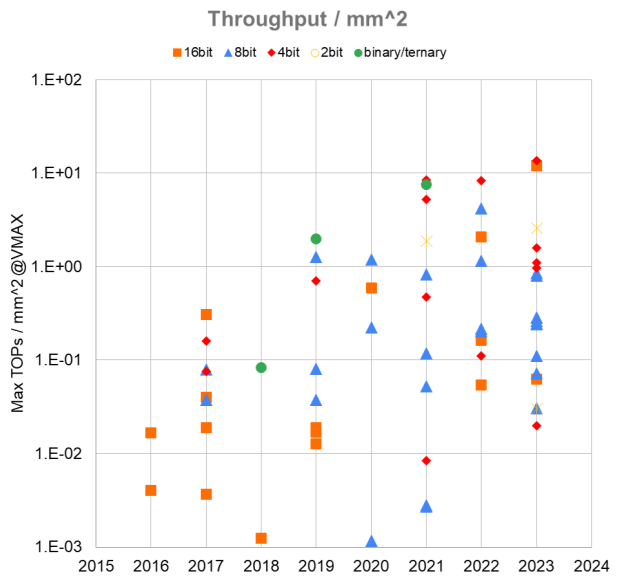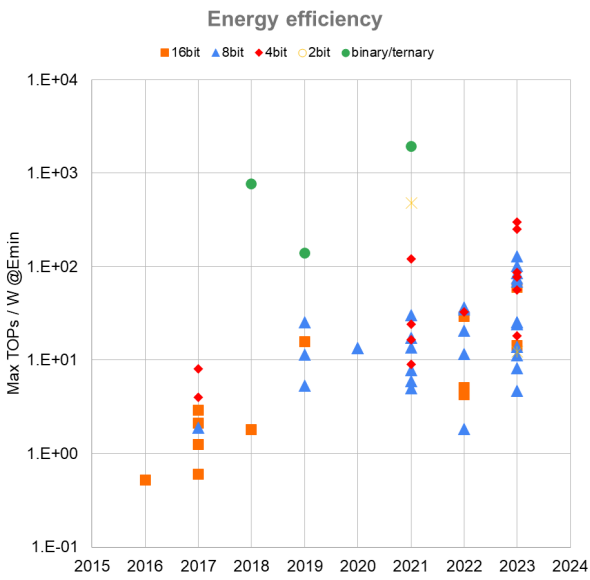
Figure 4: Evolution in energy efficiency (TOPS/W) and throughput per unit area (TOPS/mm^2) of ML inferencing processors.

[1] G. W. Burr, S. Lim, B. Murmann, R. Venkatesan and M. Verhelst, "Fair and Comprehensive Benchmarking of Machine Learning Processing Chips," in *IEEE Design & Test*, vol. 39, no. 3, pp. 18-27, 2022.

Memory – 2023 Trends

Subcommittee Chair: *Meng-Feng Chang, National Tsing Hua University, Hsinchu, Taiwan*

The demand for high-density, high-bandwidth, and low-energy memory systems continues to grow everywhere: from high-performance computing to SoC, wearables and IoT.

Innovations to 3D NAND flash, a very high density 5b/cell, will be introduced. New methods for improving reliability and quality of DRAM will be explained including probabilistic aggressor tracking against row-hammer attacks and core bias modulations to overcome process limitations. Meanwhile, the evolution of the memory high-speed interface continues by using single-ended PAM4 technology to achieve speeds of 16Gb/s/pin. Non-volatile memory continues to be applied to more advanced process nodes and continues to expand into a wider range of applications. This year a 16nm STT-MRAM that can operate in automotive environments (increased temperature range) and a 22nm near-memory-computing-macro using STT-MRAM are presented. In addition, the latest compute-in-memory developments are shown: increased energy efficiency, throughput, precision, and accuracy.

TOP PAPERS FROM ISSCC 2023 INCLUDE:

- A 1.67Tb, 5b/Cell Flash Memory Fabricated in 192-Layer Floating Gate 3D NAND Technology and Featuring 23.3Gb/mm^2 Bit Density

- A 16nm 32Mb Embedded STT-MRAM with 6ns Read Access Time, 1M Cycles Write Endurance, 20 Years Retention at 150°C and MTJ-OTP Solutions for Magnetic Immunity

- A 22nm 8Mb STT-MRAM Near-Memory-Computing Macro with 8b-precision and 46.4-160.1TOPS/W for AI-edge Devices

- A 22-nm 832-kb Hybrid-Domain Floating-Point SRAM In-Memory-Compute Macro with 16.2-70.2TFLOPS/W for High-Accuracy AI-Edge Devices

- A 4nm 16Gb/s/pin Single-Ended PAM4 Parallel Transceiver with Switching Jitter Compensation and Transmitter Optimization

- CTLE-Ising: A 1440-Spin Continuous-Time Latch-based Ising Machine with One-Shot Fully-Parallel Spin Updates Featuring Equalization of Spin States

- A 1.1V 16G DDR5 DRAM with probabilistic aggressor tracking, refresh management function, per-row hammer tracking, multi-step precharge, and core bias modulation for security and reliability enhancement

COMPUTE IN MEMORY

Memory still turns out to be the bottleneck to performance and energy not only for traditional architectures but also for non-Von-Neumann architectures: including deep learning and potentially other emerging non-conventional computing paradigms. Innovations in compute-in-memory (CIM) continue to improve energy and area efficiency while maintaining the overall network accuracy. This session showcases latest developments in SRAM-based CIM with increased energy efficiency, throughput, precision, and accuracy. Both analog and digital CIM papers are presented this year. National Tsing Hua University reports the hybrid-domain floating point SRAM CIM macro. University of California, Santa Barbara introduces a continuous-time latch-based Ising machine implementation.

HIGH-RELIABILITY AND HIGH-DENSITY DRAM

Industry requires ever-increasing DRAM performance and density for various applications . Furthermore, the ability to curtail row-hammer attacks has been a challenge to DRAM design, with improvements already discussed in other research areas. This year, this reliability enhancement is applied to a 16Gb DDR5 DRAM, and a 24Gb DDR5 with the highest density is presented in a 4TH generation 10nm DRAM technology.
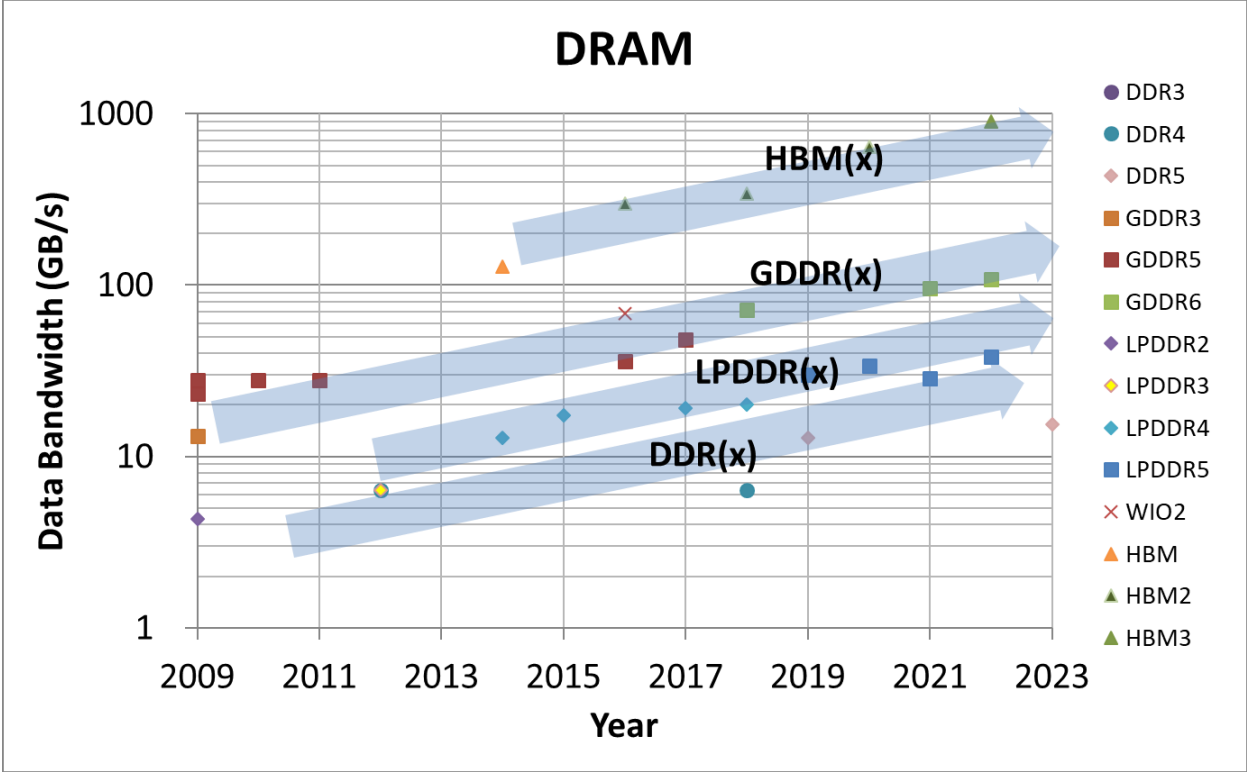
Figure 1 - DRAM data bandwidth growth

NON-VOLATILE MEMORY (NVM)

In the past decade, significant investment has been put into emerging memories to find an alternative to floating-gate-based non-volatile memory. Emerging NVMs, such as phase-change memory (PCM), ferroelectric RAM (FeRAM), magnetic spin-torque-transfer (STT-MRAM), and resistive memory (ReRAM), are showing the potential to achieve these high-cycling capabilities and lower-power-per-bit read/write operations. However, conventional flash memories are continuously improving, reaffirming them as the mainstream today and into the near future.

This year's papers report improvements in write performance (194MB/s) and read performance (34us) for conventional 3D TLC flash memories. Also reported are improvements in memory bit density for TLC (more than 20Gb/mm^2) through advancements in 3D architectures, more than 300 stacked-WL and, for the first time, 5b/cell (PLC) with the highest bit density (23.3Gb/mm^2). Figure 2 shows non-volatile memory capacity trends.
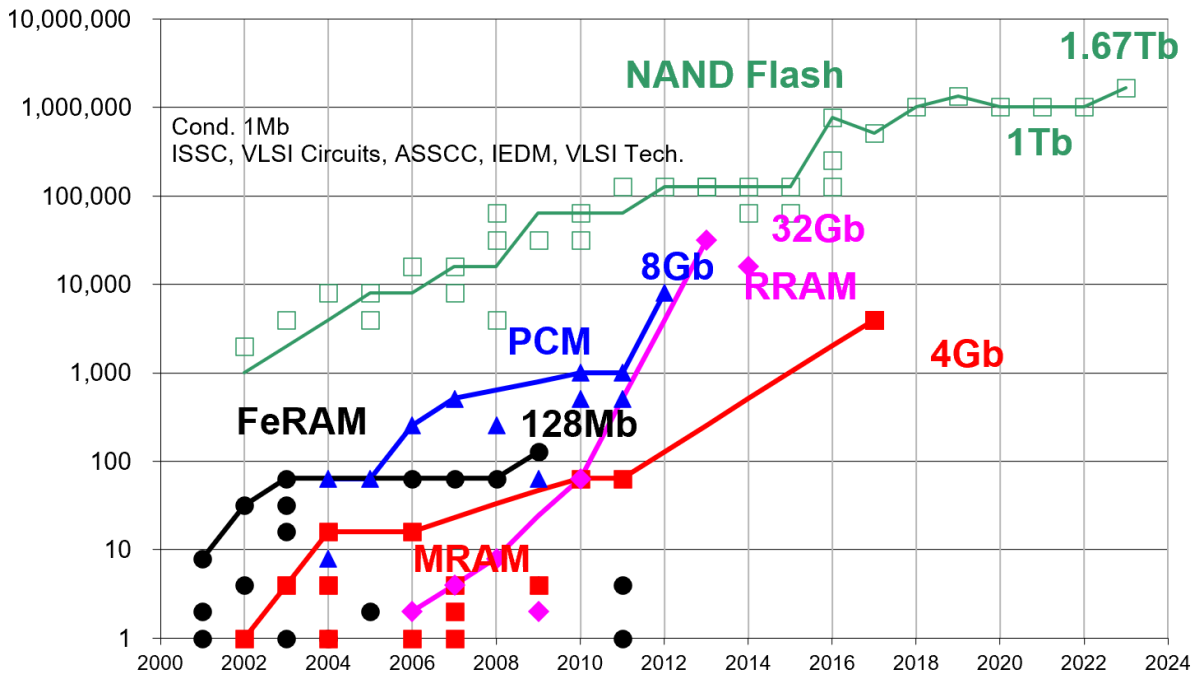

Figure 2 - Non-volatile memory capacity trend.

NAND FLASH MEMORY

NAND flash memories continue to advance towards higher density, lower power and higher performance; resulting in low-cost storage solutions that are replacing traditional magnetic hard-disk storage with solid-state disks (SSDs). The 3D memory technology is the mainstream for NAND flash memories in mass-production by semiconductor industries. Periphery-under-the-array is currently the reference architecture for TLC and QLC and PLC: it is enabling higher-bit density and multiple planes for throughput improvement.

The state-of-the-art for high performance TLC uses more than 300-stacked-WL. This year, for the first time, a 5b/cell has been presented; showing the highest bit density in the industry. Industries confirm investment to improve performance and bit density, which is expected to continue.

Figure 3 shows the observed trend in NAND Flash memory density at ISSCC over the past 20 years (and for the first time the PLC device is reported).
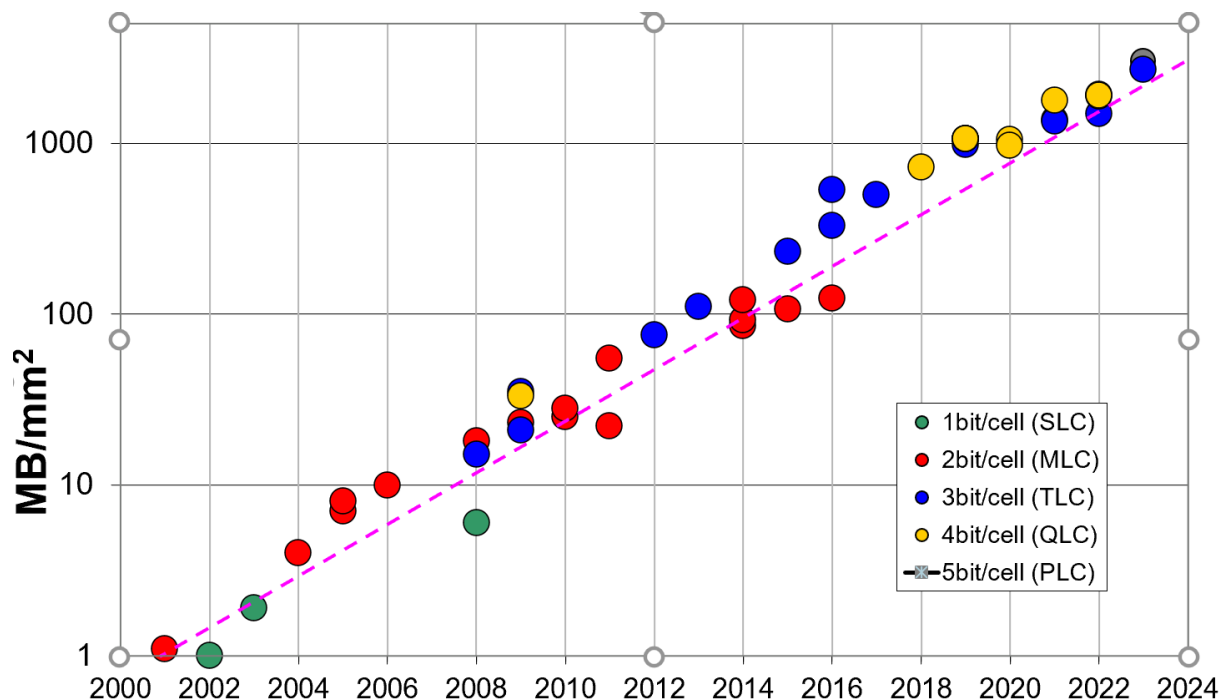

Figure 3 - NAND flash memory density trend.

HISTORICAL TRENDS IN TECHNICAL THEMES

INNOVATIVE TOPICS

IMAGERS/MEMS/MEDICAL/DISPLAYS SUBCOMMITTEE

TECHNOLOGY DIRECTIONS SUBCOMMITTEE

IMMD – 2023 Trends (Medical)

Subcommittee Chair: *Rikky Muller, University of California, Berkeley, CA*

Biomedical systems that interface with the body and nervous system in wearable and implantable applications continue to evolve toward more intelligent, multi-modal, and high-performance solutions, as well as closed-loop operation. Wearable and implantable SoCs record weak biological signals with high accuracy, stimulate neural activity, and extract key biological features under stringent power and size constraints. These new SoCs pave the way toward intelligent microdevices that enable innovations in health technologies with long-term measurement to treat on demand.

State-of-the-art biomedical integrated circuits and systems have further advanced this year at ISSCC 2023 with more intelligence as a significant trend in implantable and wearable devices, while improving dynamic range, power efficiency, and input impedance in the AFE and providing more application-oriented system-level integration towards closed-loop operation in interfacing with both the central and peripheral nervous systems. Multi-modal interfaces sense physiological signals such as PPG, BIOZ, and ExG, and provide nerve stimulation. High-dynamic-range sensing systems improve tolerance to large-amplitude interference and motion artifacts. Miniaturization combined with a high level of integration and wireless power/data transfer enables high-performance interfacing with the nervous system. Novel on-chip algorithms break the trade-off between hardware resources and real-time performance and can detect seizures on unseen patients without collecting their data for classifier training.

Innovations in intelligent classification of neurological signals (with minimal to no pre-training), enabled by on-chip or off-chip machine learning, have the potential to offer improved patient-independent neural interfaces to treat on-demand and in closed-loop. Incorporating machine learning with a high channel count of recording and stimulation will enable emerging therapeutic techniques to restore a healthy condition from chronic pain or other disorders without the side effects of drugs. These advances offer tremendous market potential in the medical market space.

IMMD – 2022 Trends (Imagers)

Subcommittee Chair: *Rikky Muller, University of California, Berkeley, CA*

In the field of image sensors, consumer applications push the evolution of fabrication processes to achieve advanced autofocusing capabilities in a small-pixel-pitch over large arrays. Samsung is further evolving the multi-photodiode pixel structure concept for all-direction autofocus with no compromises on other parameters, reporting a 50-MP CIS with 20k e- FWC, 0.98 e- random noise, and 86.2 dB dynamic range based on a quad-pixel structure. This is achieved thanks to highly customized front-side deep-trench isolation between pixels and between photodiodes in the pixel.

The race to shrink pixel size has expanded from conventional intensity imagers to event-based vision sensors, whose trend is dominated by hybrid approaches that combine intensity information with temporal contrast change detection and the 3D integration of the sensing layer with advanced processing stages. Omnivision achieves a record 4.6 Gevent/s rate on a 15Mpixel imager + 1Mpixel event vision sensor built upon a three-wafer stack. Sony addresses the pixel shrinking challenge in two 3D stacked works, both showing high dynamic range capabilities: one with a 35.6Mpixel array featuring 1.22µm pitch and 1.57e- random noise for the RGB intensity pixels and 4.88µm pitch for the event pixels, and another one that matches the resolution of the intensity and event images with a pitch of 2.97µm thanks to a shared front-end for contrast change detection.

On-chip image processing is also at the core of ultra-low-power imagers and high-dynamic-range sensors. This year, ISSCC presents a 55pW/pixel peak power imager powered by a 3.3×3.3mm^2 solar cell that can operate with no interruptions even in dim light conditions. Another work introduces a novel approach to HDR imaging that combines pixel-wise exposure coding with a binary readout scheme that compares the pixel flux to a sinusoidal reference, achieving 95 dB dynamic range.

Advances on the non-visible part of the spectrum witness the appearance of a SPAD-based X-ray detector and of the first 10Kpixel THz imager showing high-resolution images in the 3.08-to-3.86THz frequency range. The former is made of a high-frame-rate, high-dynamic-range, and low-power SPAD imager coupled to an X-ray scintillator. HDR is achieved working in a seamless global shutter mode, combining photon counting at low photon flux regimes and a time-encoded intensity estimation at high photon flux regimes. In the latter, a step-covered patch antenna and a defected ground structure achieve a high sensitivity at >3.0THz frequencies over a large pixel array

Technology Directions – 2023 Trends

Subcommittee Chair: *Ali Hajimiri, California Institute of Technology, California*

Technology innovations bring the promise of enabling new system functionalities or substantially increasing the efficiency of existing ones. Harnessing such innovations for solving tangible real-world problems requires novel system-level solutions. With a focus on envisioning the future, emerging trends in Technology Directions this year at ISSCC 2023 covers a wide range of topics including quantum engineering, emerging sensor systems, internet-of-things (IoT), optical computation and photonics. ISSCC 2023 features four sessions representing the latest technological innovations in the following areas:

Quantum computing:

Quantum computers often comprise an array of quantum devices operated at cryogenic temperatures that must be controlled by an electronic interface. To fulfill the quantum-computing promise of a disruptive computational advantage, significant research efforts are pushing the scaling of those quantum processors and, consequently, their electronic interface. ISSCC 2023 mirrors this trend by presenting SoCs that can improve system reliability and complexity by operating at cryogenic temperatures to reduce the gap between the control electronics and the cryogenic qubits. To this end, a single SoC is shown driving both single-qubit and two-qubit operations on superconducting qubits, and two papers present advances in the microwave-driver architecture for superconducting qubits to reduce power dissipation and chip area. While ad-hoc circuit techniques are introduced to improve the performance of individual cryogenic circuit blocks, such as flicker-noise reduction in the presented cryo-CMOS VCO, novel functionalities are also explored, such as a THz backscatter transceiver to avoid heat-transferring cables for data transmission between room-temperature and cryogenic electronics.

Emerging Systems and IOT:

ISSCC 2023 pushes the frontiers of integrated sensing and ultra-low-power IoT systems across a diverse space of emerging applications including biomedical, electrochemical, material and energy sensing. We see new advances in biosensing and bio-manipulation on-chip with demonstration of optics and electronics co-integration enabling chip-scale fluorescence sensor arrays, and 2D cell manipulation with on-chip electrodes. State-of-the-art performance in low-power operation and higher sensitivity are demonstrated with sub-THz electron paramagnetic resonance systems, low-power electrochemical and battery health monitoring ICs for electrical vehicles. Energy-efficient and scalable networks are needed to enable the future world of intelligent sensors. On that front, ultra-low-power IoT tags are demonstrated that can harvest energy from LTE to backscatter Bluetooth into WiFi channels.

Ideas for the Future and Ideas Outside the Box:

This year, ISSCC 2023 includes two sessions focused on presenting out-of-the-box ideas for the future, highlighting topics both familiar and less-familiar to ISSCC that push the envelope of solid-state circuits and systems that are bound to inspire the next generation of engineers and scientists. These sessions cover a wide range of topics including 3D memory, dielectric waveguides, DNA nanotechnology, bioresorbable implantable devices, 2D materials, energy harvesting systems and advanced photonics and photonic design techniques.